Veeam Recovery Orchestrator

Version 13

Operations Guide

March, 2026

NOTE

Read the End User Software License Agreement before using the accompanying software programs. Using any part of the software indicates that you accept the terms of the End User Software License Agreement.

Contents

Contacting Veeam Software

At Veeam Software we value feedback from our customers. It is important not only to help you quickly with your technical issues, but it is our mission to listen to your input and build products that incorporate your suggestions.

Customer Support

Should you have a technical concern, suggestion or question, visit the Veeam Customer Support Portal to open a case, search our knowledge base, reference documentation, manage your license or obtain the latest product release.

Company Contacts

For the most up-to-date information about company contacts and office locations, visit the Veeam Contacts Webpage.

Online Support

If you have any questions about Veeam products, you can use the following resources:

- Full documentation set: veeam.com/documentation-guides-datasheets.html
- Veeam R&D Forums: forums.veeam.com

About This Document

The guide is designed for IT professionals who plan to use Veeam Recovery Orchestrator. It is primarily aimed at administrators who manage enterprise environments and lack a flexible scalable automation system. Veeam Recovery Orchestrator provides a comprehensive set of features to ensure easy execution, testing and documentation of DR plans.

Veeam Recovery Orchestrator is built on top of Veeam Backup & Replication and Veeam ONE, and this guide assumes that you have a good understanding of these solutions.

Accessing Orchestrator UI

IMPORTANT

Before you access the Orchestrator UI, make sure that TLS 1.2 is enabled both on the Orchestrator server and on the machine that you plan to use to access the UI.

To access the Orchestrator UI, perform the following steps:

1. In a web browser, navigate to the Orchestrator UI web address. The address consists of an FQDN of the Orchestrator server and the website port specified during installation (by default, **9898**). Note that the Orchestrator UI is available over HTTPS only.

   ```
   https://<FQDN>:<port>
   ```

   Keep in mind that Internet Explorer is no longer supported. To access the Orchestrator UI, use Microsoft Edge, Google Chrome or Mozilla Firefox.

2. If you log in for the first time, specify credentials of the local Administrator who performed Orchestrator installation. The user name must be specified in the *DOMAIN|USERNAME* or *USERNAME* format.

   To authenticate using the credentials provided when logging into the system, click **Log in as current user**.

TIP

To be able to log in as the current user, you must first do the following:

1. Complete the Initial Configuration Wizard as described in the Veeam Recovery Orchestrator Deployment Guide, section After You Install.
2. Enable the **Automatic logon with current user name and password** option in the security setting of your browser.

In future, you can configure users and roles to grant access to the Orchestrator UI. For more information, see the Veeam Recovery Orchestrator Operations Guide, section Managing User Accounts.

3. Click **Log in**.

   If the Veeam Recovery Orchestrator Operations Guide, section Enabling and Disabling Multi-Factor Authentication, Orchestrator will prompt you to enter a code to verify the user identity. In the code field, enter the temporary six-digit code generated by the authentication application running on your trusted device. Then, click **Confirm**.

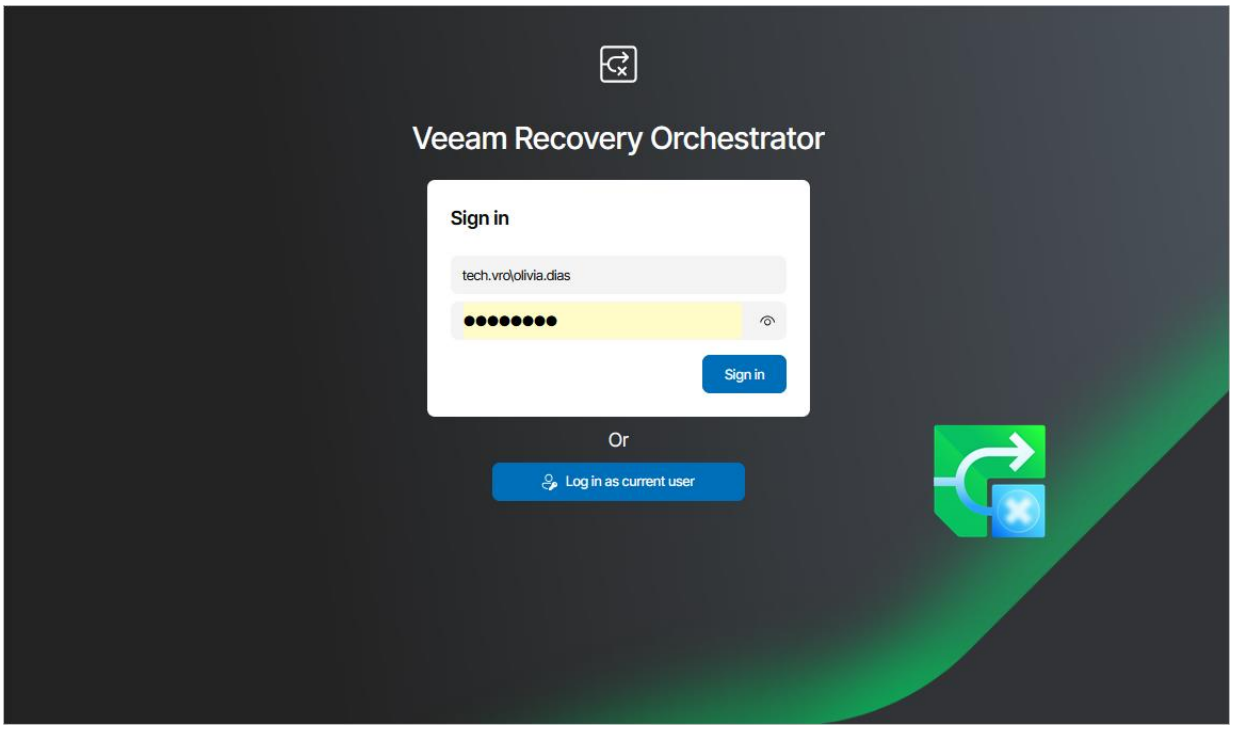

Configuring Trusted Connection

The Orchestrator UI uses SSL to ensure secure data communication between Orchestrator and a web browser.

When you install Orchestrator, you can generate or choose a self-signed certificate. In this case, when you try to access the Orchestrator UI in a web browser, the browser will display a warning notifying that the connection is untrusted (although it is secured with SSL).

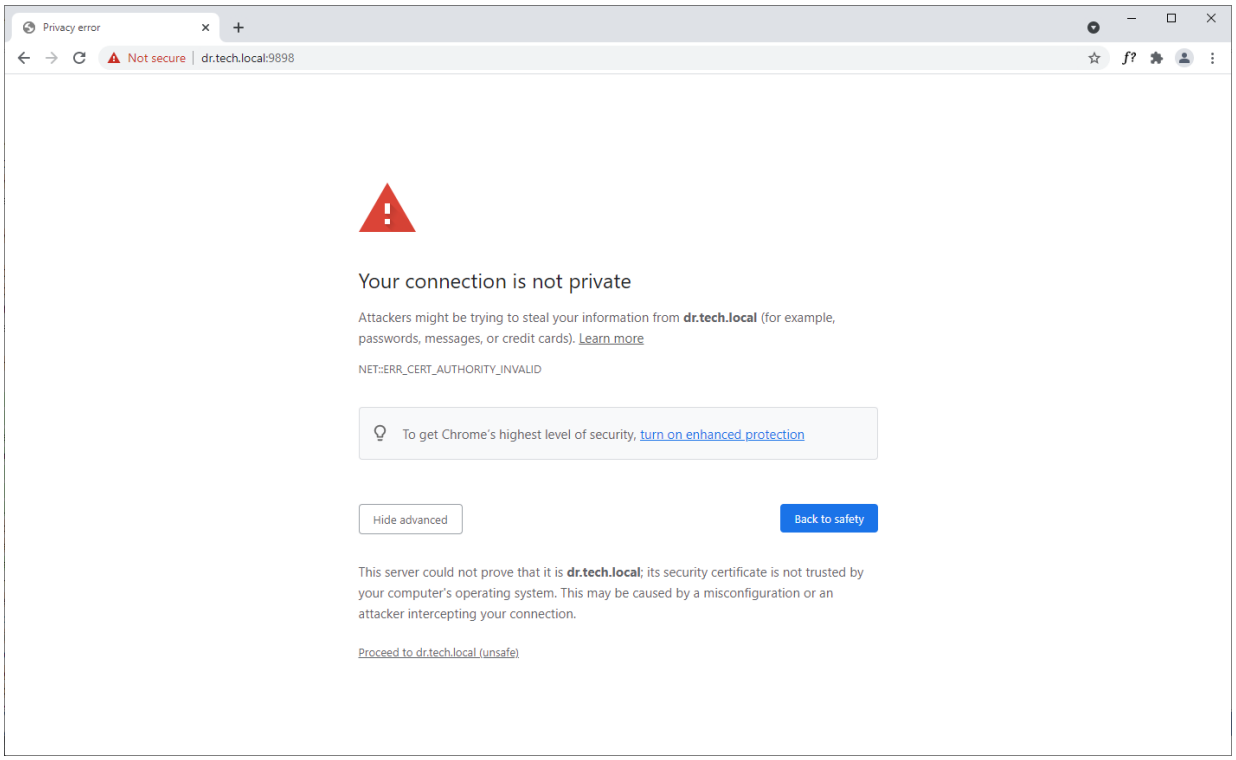

To eliminate the warning, import the self-signed certificate to client machines (the machines from which you plan to access the Orchestrator UI website). To learn how to import SSL certificates, see this Microsoft KB article.

If you want to use the certificate generated during installation, perform the following steps:

1. Log in to the machine where Orchestrator is installed.
2. Open the Microsoft Management Console snap-in.
   a. Navigate to **Certificates** > **Trusted Root Certification Authorities** > **Certificates**.
   b. Export the *Veeam Self-Signed Certificate* following the instructions provided in this Microsoft KB article.
3. Import the *Veeam Self-Signed Certificate* to client machines.

> **TIP**
>
> If you still have issues accessing the Orchestrator UI, check your browser settings to ensure that the Orchestrator UI site is included in the **Trusted Sites** list.

Logging Out

To log out of the Orchestrator UI, at the top right corner of the Veeam Recovery Orchestrator window, click the user name and then click **Logout**.

Configuring Veeam Recovery Orchestrator

To start working with Orchestrator, perform a number of steps for its configuration:

1. Configure access to Orchestrator:
    a. Create scopes to allow access to Orchestrator.
    b. For each scope, add users to specific roles.
    c. Configure MFA.
2. Connect infrastructure:
    a. Connect Veeam Backup & Replication servers.
    b. Connect vCenter Servers.
    c. Connect Microsoft Hyper-V servers.
    d. Connect NetApp and HPE storage.
3. Create recovery locations:
    a. Create a VMware vSphere recovery location.
    b. Create a storage recovery location.
    c. Create a Microsoft Azure recovery location.
    d. Create a Microsoft Hyper-V recovery location.
4. Manage credentials under which recovery plan steps will be launched.
5. Manage plan steps to be performed while running the recovery process.
6. Manage template jobs to be used to reprotect inventory groups in recovery plans.
7. Manage DataLabs to be used to test recovery plans.
8. Manage access to inventory items to be used in recovery plans.
9. Configure general settings:
    a. Specify email server settings.
    b. Specify report subscription settings.
    c. Configure report retention settings.

Managing Scopes

Orchestrator controls access to its functionality with the help of scopes. A scope defines the access that users have to inventory items such as inventory groups, plan steps, recovery locations, template jobs, credentials, DataLabs and YARA rules. These inventory items are used to build recovery plans.

For a scope, you can:

1. Assign roles to users in a scope.

   After you install Orchestrator, you will have one scope only — the *Default Scope*. By design, all items discovered by the solution are added to this scope, including items collected from connected Veeam Backup & Replication servers (such as inventory groups, credentials, template jobs and so on).

   You can remove items from the *Default Scope*; however, it is recommended that you create additional scopes to provide more granular permissions to users for enhanced security — to do that, follow the instructions provided in section Creating Scopes.

2. Limit the number of inventory items available in the Orchestrator UI to users from that scope.

   All newly-created scopes will NOT automatically have any visibility of newly discovered inventory items until those items are explicitly added to that scope. To add an item to a scope, navigate to **Scope Inventory** and follow the instructions provided in sections Managing Inventory Items and Connecting DataLabs.

Creating Scopes

To create a new scope:

1. Switch to the **Administration** page.
2. Navigate to **Scopes**.
3. Click **Add**.
4. In the **New Scope** window:
   a. Use the **Name** and **Description** fields to enter a name for the new scope and to provide a description for future reference.

      The maximum length of the scope name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

   b. Click **Apply** to save the scope.

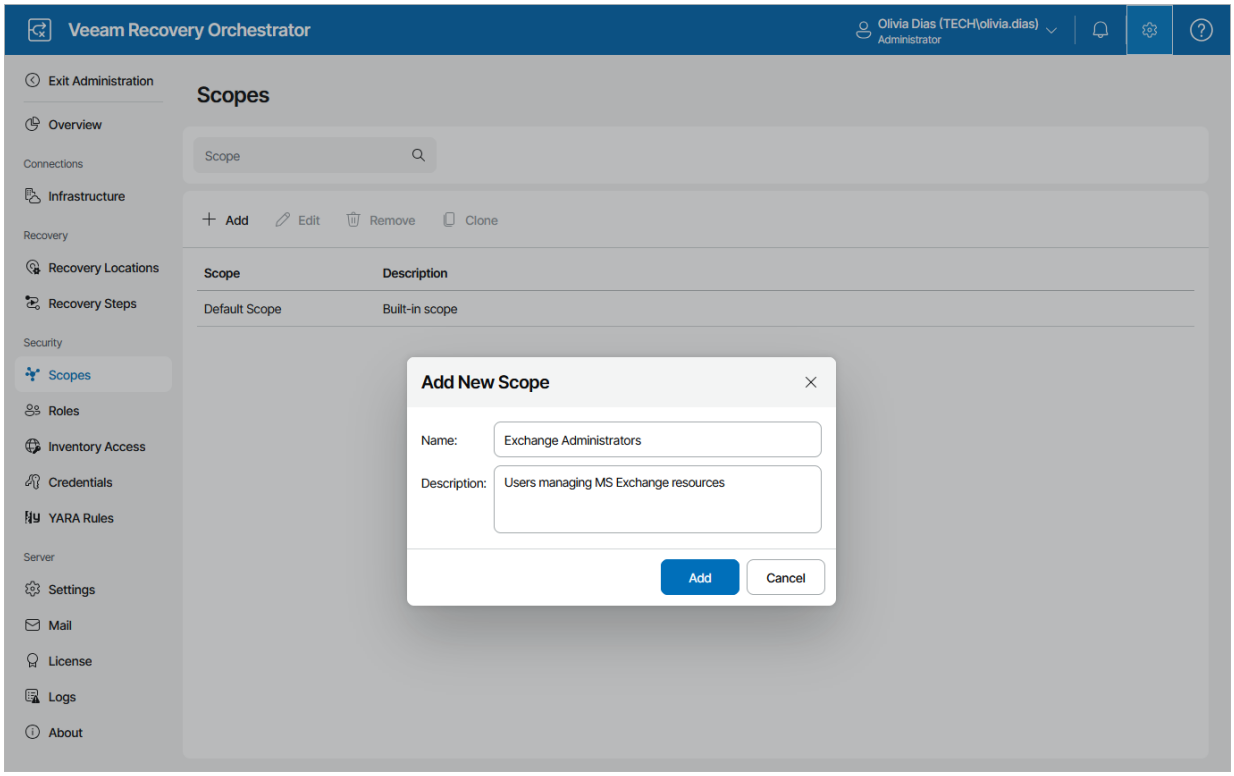

Cloning Scopes

You can also create a new scope by cloning a scope that already exists. When cloning a scope, you can copy its settings so that you do not have to configure the same settings once again for the new scope. The new scope will have the same configuration as the existing scope, which means that all items configured for the existing scope will be applied to the new scope.

To clone a scope:

1. Select an existing scope that you want to use as a template for the new scope and click **Clone**.
2. In the **Clone Scope** window:
   a. Use the **Name** and **Description** fields to enter a name for the new scope and to provide a description for future reference.

      The maximum length of the scope name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

   b. In the **Copy settings** list, select check boxes next to settings that you want to copy from the existing scope.

   c. Click **Apply** to save the scope.

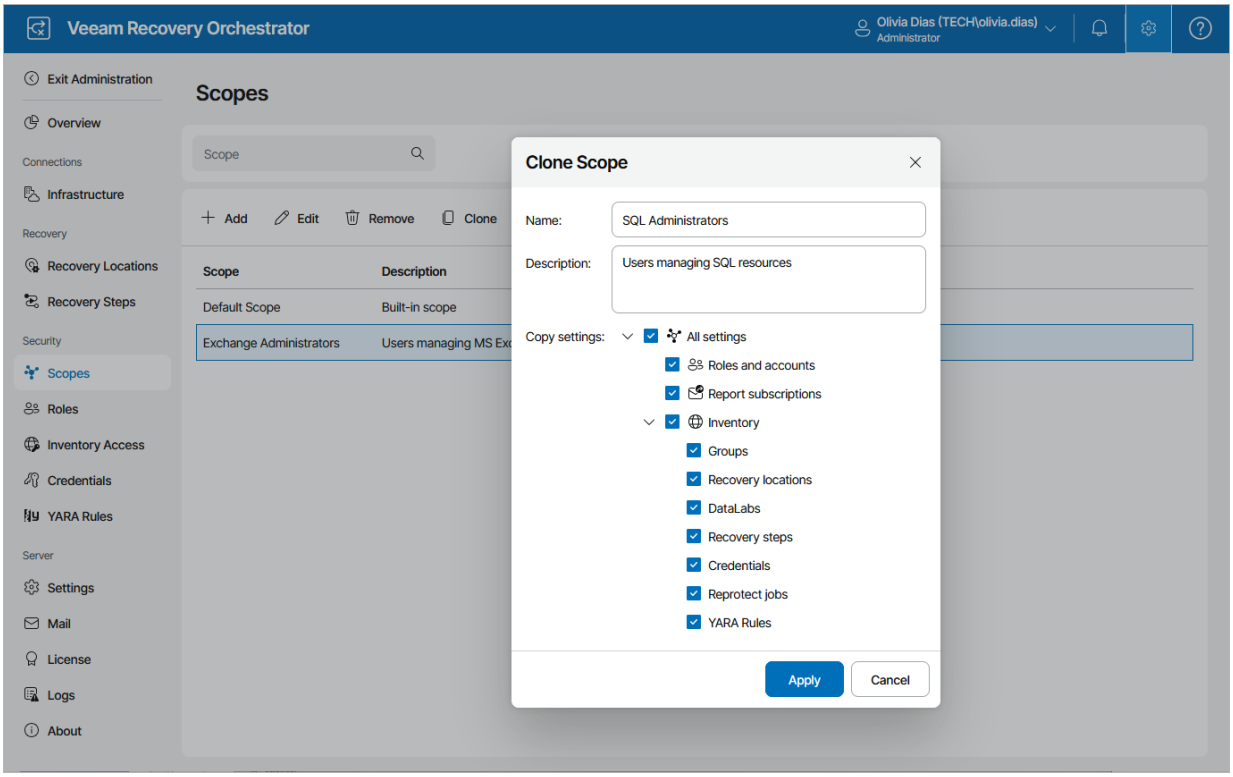

Managing User Accounts

Veeam Recovery Orchestrator controls access to its functionality with the help of user roles. A user role limits the number of operations available in the Orchestrator UI to users with that role. Role-based access is controlled by adding users and groups to the relevant role and scope pair in the Orchestrator UI.

There are 3 roles that can be assigned to users and user groups working with the Orchestrator UI: *Administrator*, *Plan Author* and *Plan Operator*. For the role descriptions, see the Veeam Recovery Orchestrator Deployment Guide, section Roles.

Adding User Accounts

To add a user or a user group:

1. Launch the Add New Account wizard.
2. Specify an account to which you want to assign a role.
3. Assign roles and scopes to your account.
4. Finish working with the wizard.

Step 1. Launch Add New Account Wizard

To launch the **Add New Account** wizard, do the following:

1. Switch to the **Administration** page.
2. Navigate to **Scope Users**.
3. Click **Add Account**.

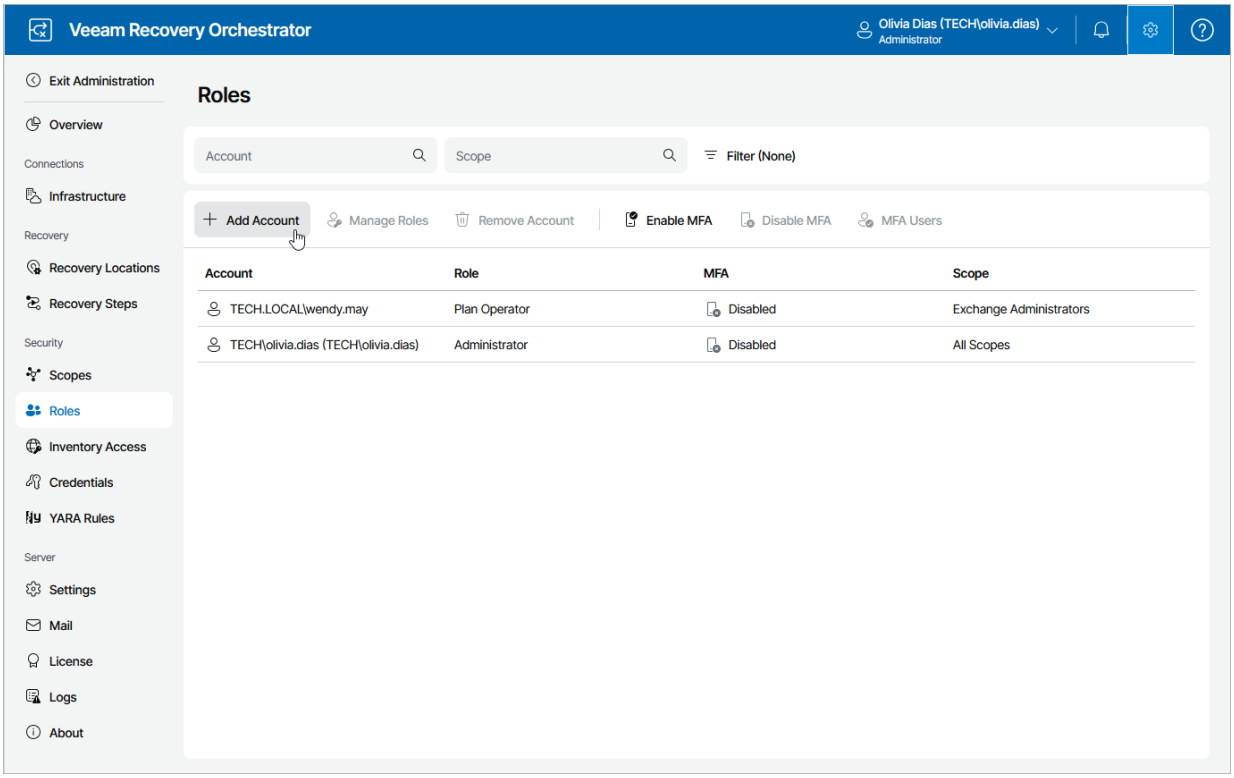

Step 2. Specify Account

At the **Accounts** step of the wizard, do the following:

1. From the **Account type** list, select *User* or *Group*.
2. Use the **Account** and **Location** fields to enter the user or group name and to select the location to which the user or group belongs – either a domain or local OS workgroup.

   For more information on the required account permissions, see the Veeam Recovery Orchestrator Deployment Guide, section Permissions.
3. Click **Add**.
4. Repeat the procedure for each user or group that you want to add, and click **Next**.

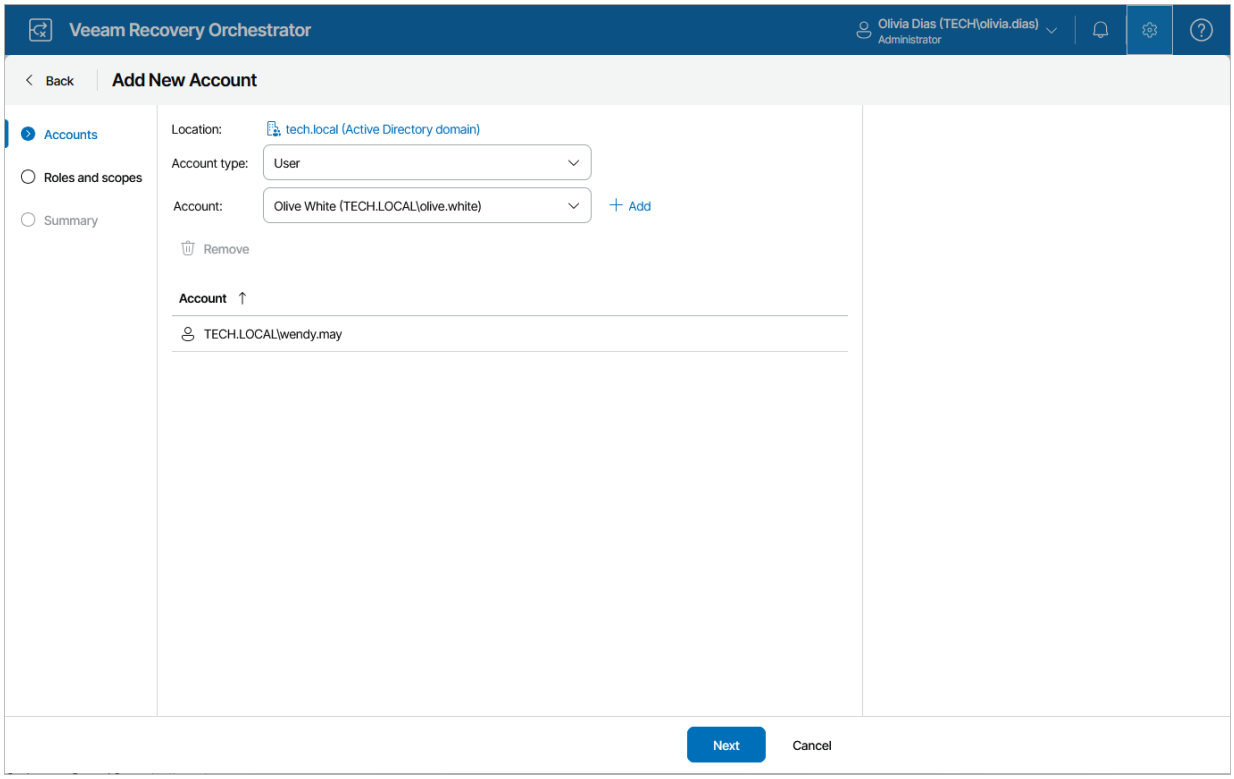

Step 3. Select Role and Scope

At the **Role and scopes** step of the wizard, do the following:

1. From the **Role** drop-down list, choose the required role. For more information on roles that can be assigned to users and user groups working with the Orchestrator UI, see the Veeam Recovery Orchestrator Deployment Guide, section Roles.
2. From the **Scope** drop-down list, choose the required scope.
3. Click **Add**.
4. Repeat the procedure to add more role and scope pairs, and click **Next**.

> **NOTE**
>
> You can neither add nor remove scopes for Orchestrator Administrators — when you add a user with the *Administrator* role, this role is automatically assigned to all the existing scopes.

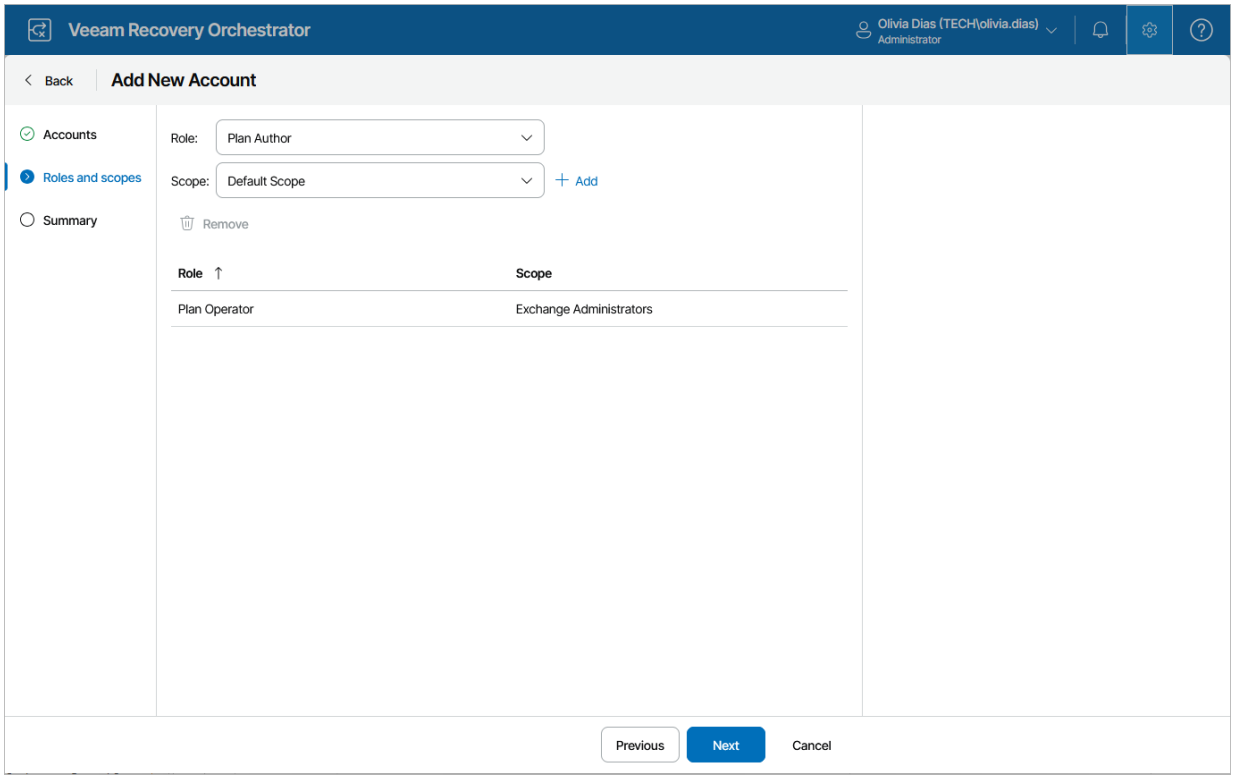

Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.

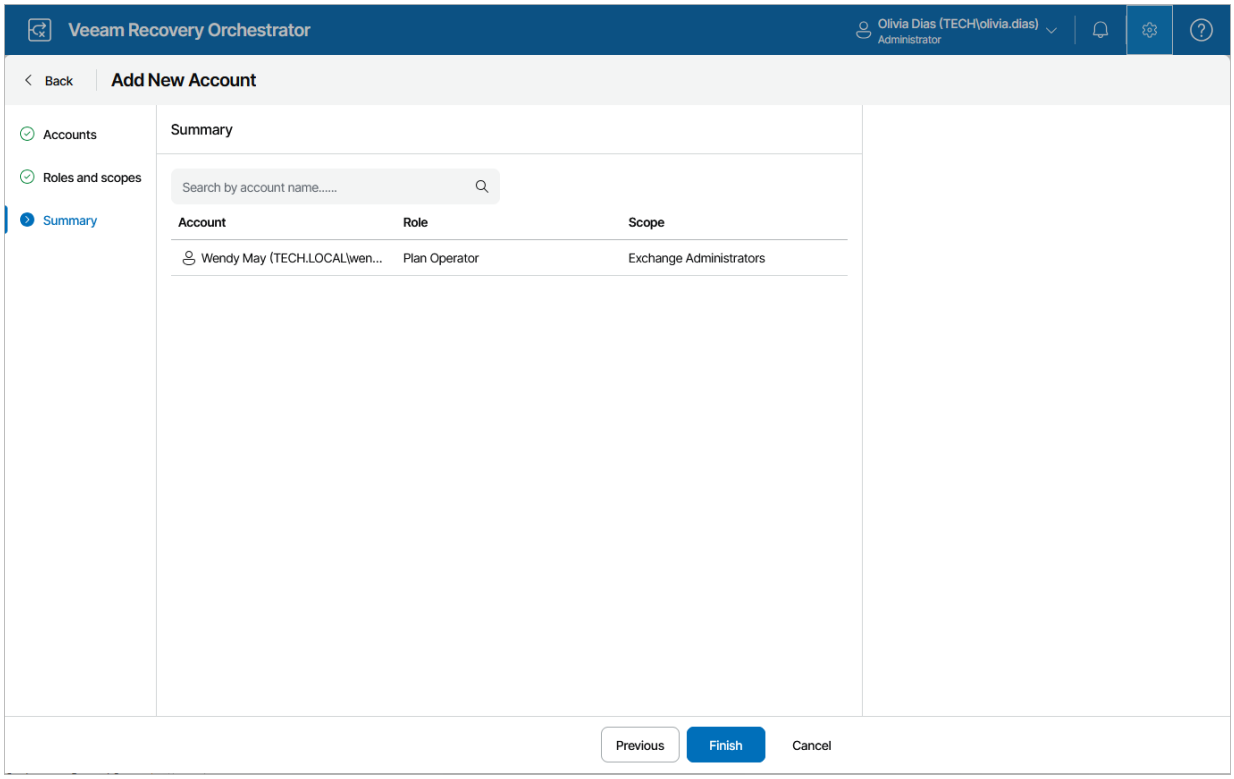

Editing User Roles and Scopes

To edit role and scope pairs assigned to the user account, do the following:

1. Switch to the **Administration** page.
2. Navigate to **Roles**.
3. Select the user account that you want to edit and click **Manage Roles**.
4. In the **Manage Roles** wizard:
    a. At the **Roles and scopes** step, modify role and scope pairs assigned to the user account.
    b. At the **Summary** step, review configuration information and click **Finish**.

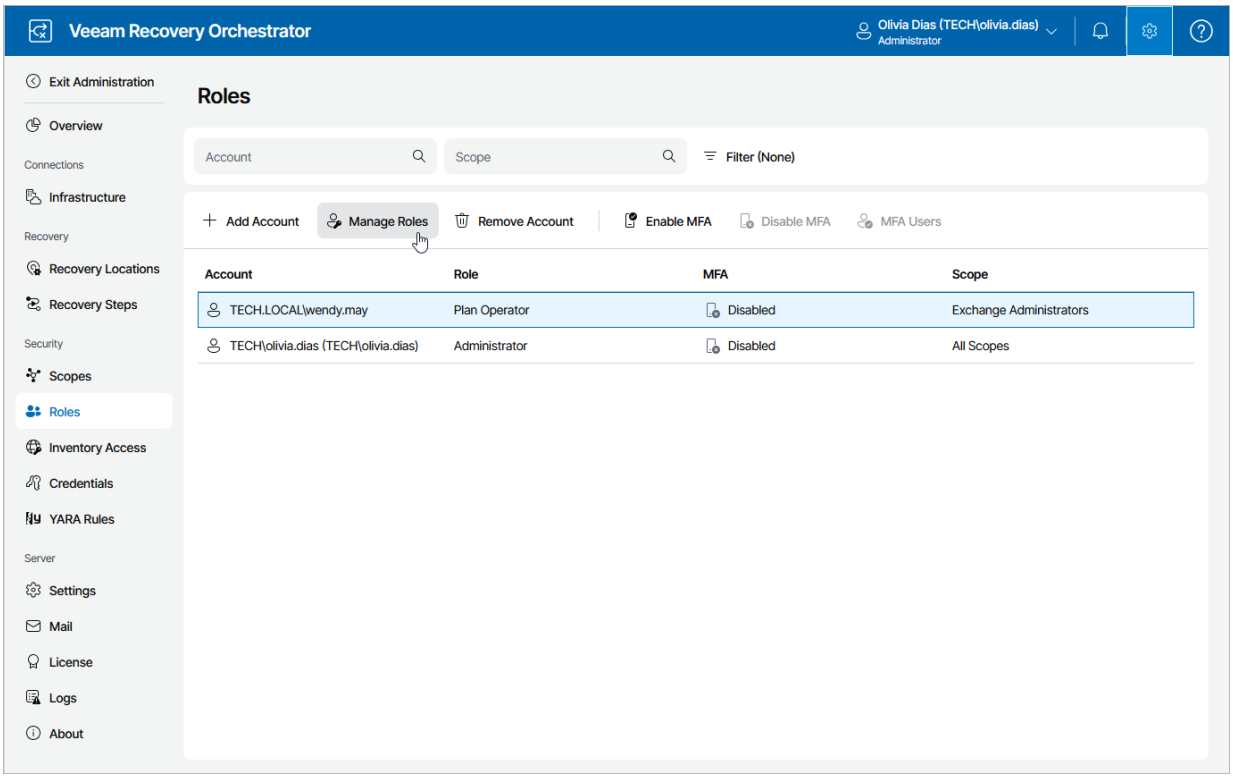

Removing User Accounts

To remove a user account from Veeam Recovery Orchestrator, do the following:

1. Switch to the **Administration** page.
2. Navigate to **Roles**.
3. Remove each role assigned to the account. To do that, select the role and click **Remove Account**.

   When you remove the last role assigned to the account, the account will be removed from the Orchestrator database automatically.

> **NOTE**
>
> There must always exist at least one user account with the *Administrator* role. This means that you cannot remove an Administrator account if you do not have any other Administrator accounts left.

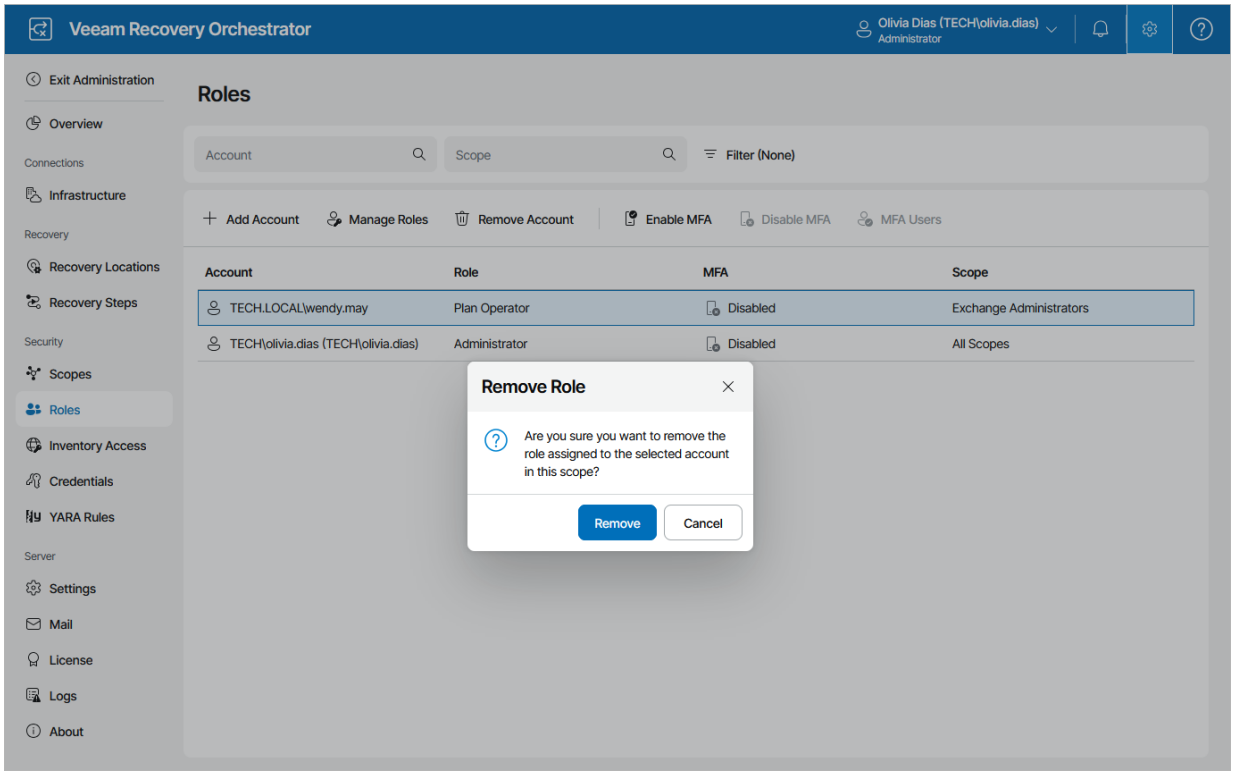

Enabling and Disabling Multi-Factoring Authentication

Veeam Recovery Orchestrator supports multi-factor authentication (MFA) for additional user verification, which allows Administrators to enable, disable and reset MFA for all users and user groups working with the Orchestrator UI. MFA in Veeam Recovery Orchestrator is based on the Time-based One-Time Password (TOTP) method that requires the user to verify their identity by providing a temporary six-digit code generated by an authentication application running on a trusted device.

To enable or disable MFA, do the following:

1. Switch to the **Administration** page.
2. Navigate to **Roles**.
3. Click **Enable MFA** or **Disable MFA**. As a result, the statuses of all user accounts added to Orchestrator will change to *Not configured* or *Disabled*.

When a user logs in to the Orchestrator UI for the first time, they must do the following:

1. Install a supported authentication application on a trusted device. Note that only Google Authenticator and Microsoft Authenticator are fully supported by Orchestrator.
2. Scan the displayed QR code using the camera of the trusted device.
3. Enter a verification code generated by the authentication application. Note that mobile push notifications are not supported — you can get a verification code only in the authentication application installed at step 1.

> **IMPORTANT**
>
> To avoid connection failures, make sure that the system time on the trusted device that you are using is synchronized with the system time on the machine where Orchestrator is installed.

4. Click **Confirm**. As a result, the status of the user account will change to *Configured*.

> **TIP**
>
> - In case a user loses access to their trusted device, an Administrator can reset MFA for this user. To do that, click **MFA Status**, choose the necessary user and click **Reset MFA**. The next time the user logs in to the Orchestrator UI, they will have to scan the QR code again using another trusted device.
> - If an Administrator loses access to their trusted device and there are no more Administrators added to Orchestrator, this Administrator will not be able to reset MFA for themselves and will no longer be able to log in to the Web UI — unless they are part of a user group with the *Administrator* role assigned.
>
> As a workaround, the Administrator can add a new user to the user group in Windows settings of their workstation or domain controller. Since Windows settings are synchronized with the configuration database of the Orchestrator server, the new user will also be added to the user group in Orchestrator and then will become able to reset MFA for the Administrator and all other users.

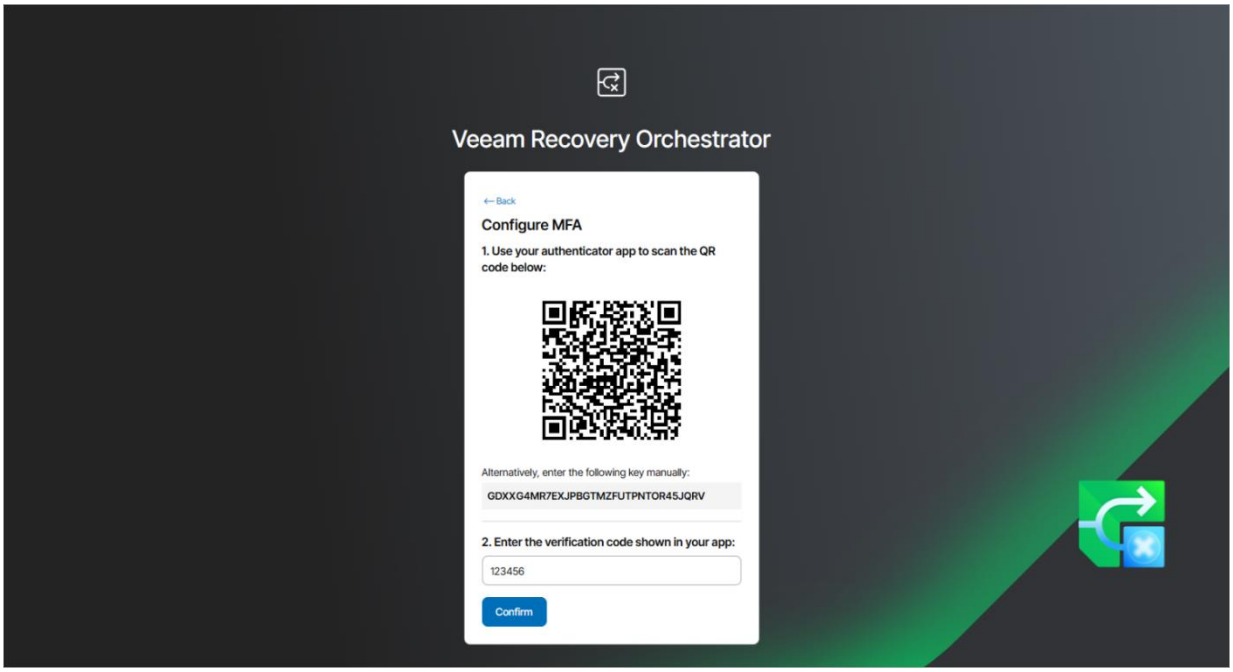
Veeam Recovery Orchestrator
← Back
Configure MFA
1. Use your authenticator app to scan the QR code below:
Alternatively, enter the following key manually:
GDXXG4MR7EXJPBGTMZFUTPNTOR45JQRV
2. Enter the verification code shown in your app:
123456
Confirm

Connecting Infrastructure

If required, you can configure connections to vCenter Servers, SCVMM servers, Hyper-V clusters and NetApp or HPE storage systems.

If you have already connected Veeam Backup & Replication servers in the **Initial Configuration** wizard, you do not need to connect them again. For more information, see the Veeam Recovery Orchestrator Deployment Guide, section After You Install.

> **NOTE**
>
> No additional connection is required for recovery to Microsoft Azure as Orchestrator will use the Microsoft Azure compute and storage credentials configured on the connected Veeam Backup & Replication servers. However, if you plan to execute custom scripts inside an Microsoft Azure machine, you must configure a direct connection to this machine as described in section Connecting Microsoft Azure Servers.

Connecting Veeam Backup & Replication Servers

To be able to orchestrate a remote Veeam Backup & Replication server, you must deploy an Orchestrator agent on the server — the agent will trigger orchestration commands on that server. For Linux-based Veeam Backup & Replication servers, you must also enable remote data collection in the Veeam Host Management console as described in the Veeam Backup & Replication User Guide, section Configuring Backup Infrastructure Settings.

If you have already connected servers during the initial Orchestrator UI configuration, you do not need to connect them again. For more information, see the Veeam Recovery Orchestrator Deployment Guide, section After You Install.

> **NOTE**
>
> - Orchestrator supports connecting remote Veeam Backup & Replication servers that run the PostgreSQL and Microsoft SQL Server database systems.
> - Orchestrator does not support connecting remote Veeam Backup & Replication servers with enabled multi-factor authentication (MFA). To work around the issue, disable MFA for the Veeam Backup & Replication service account and connect the remote backup server using the credentials of this account. For more information on how to disable MFA, see the Veeam Backup & Replication User Guide, section Multi-Factor Authentication.

To deploy an Orchestrator agent on a remote Veeam Backup & Replication server:

1. Switch to the **Administration** page.
2. Navigate to **Infrastructure** > **Veeam Data Platform**.
3. Click **Deploy Agent** and choose whether you want to deploy an Orchestrator agent for Windows or for Linux.
4. Complete the **Install Orchestrator Agent** wizard:
    a. At the **Settings** step of the wizard, specify the following connection settings:
        i. Use the **Server type** options to specify whether the server is a Veeam Backup & Replication server or Veeam Backup Enterprise Manager server.

        If you choose the **Veeam Enterprise Manager** option, Orchestrator agents will be deployed to all Veeam Backup & Replication servers managed by the Veeam Backup Enterprise Manager.

        ii. Use the **DNS name or IP** field to enter a DNS name or an IPv4 address of the server where you want to deploy the Orchestrator agent. The maximum length of the location name is 256 characters; the following characters are not supported: * : / \ ? @ [ ] ; : = + " < > | .

> **NOTE**
>
> If you want to connect a VSA high availability cluster, you must specify a DNS name or IPv4 address of this cluster.

        iii. From the **Credentials** drop-down list, choose the necessary account for connecting to the server. For an account to be displayed in the **Credentials** list, it must be added to the configuration database as described in section Adding Credentials. If you have not set up an account beforehand, click **Add** and follow the steps of the **Add Credentials** wizard. For more information on the required account permissions, see the Veeam Recovery Orchestrator Deployment Guide, section Permissions.

        The provided credentials will be also used to launch the Orchestrator agent on the server. The user name must be specified in the *DOMAIN|USERNAME* or *USERNAME* format.

NOTE

If you want to connect an Enterprise Manager server, you must specify the credentials that were used when installing Veeam Backup Enterprise Manager.

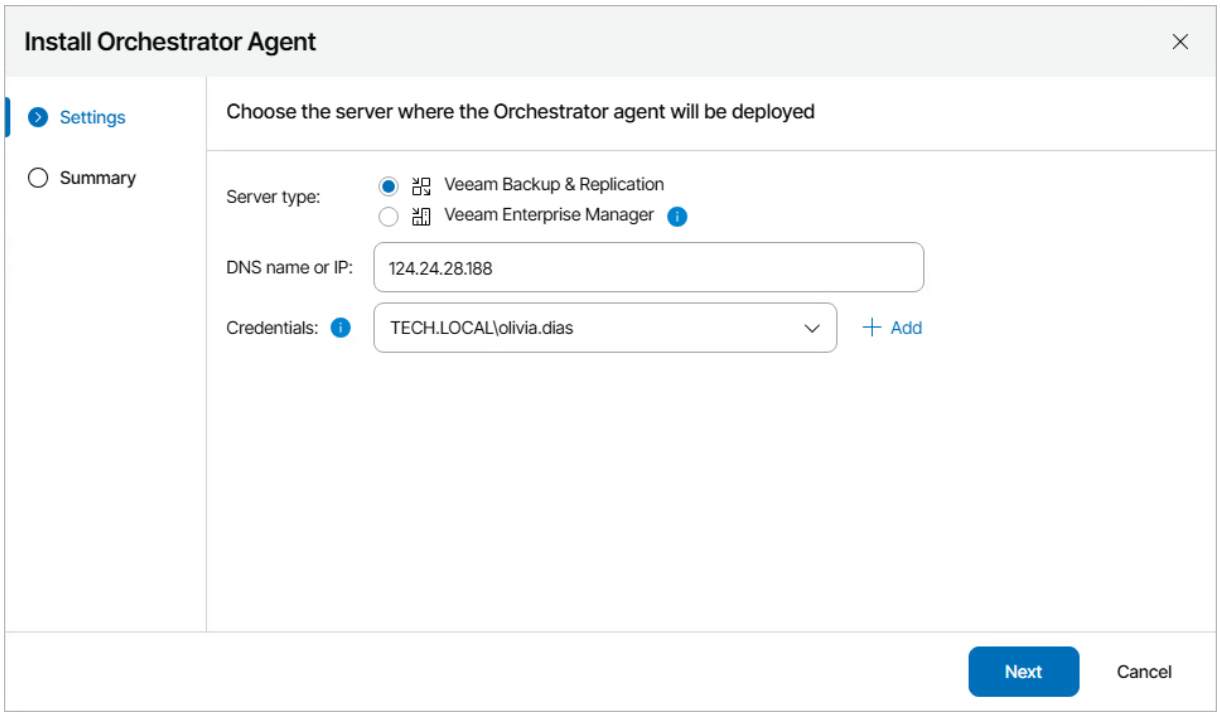

b. At the **Summary** step of the wizard, review the connection details and click **Finish**.

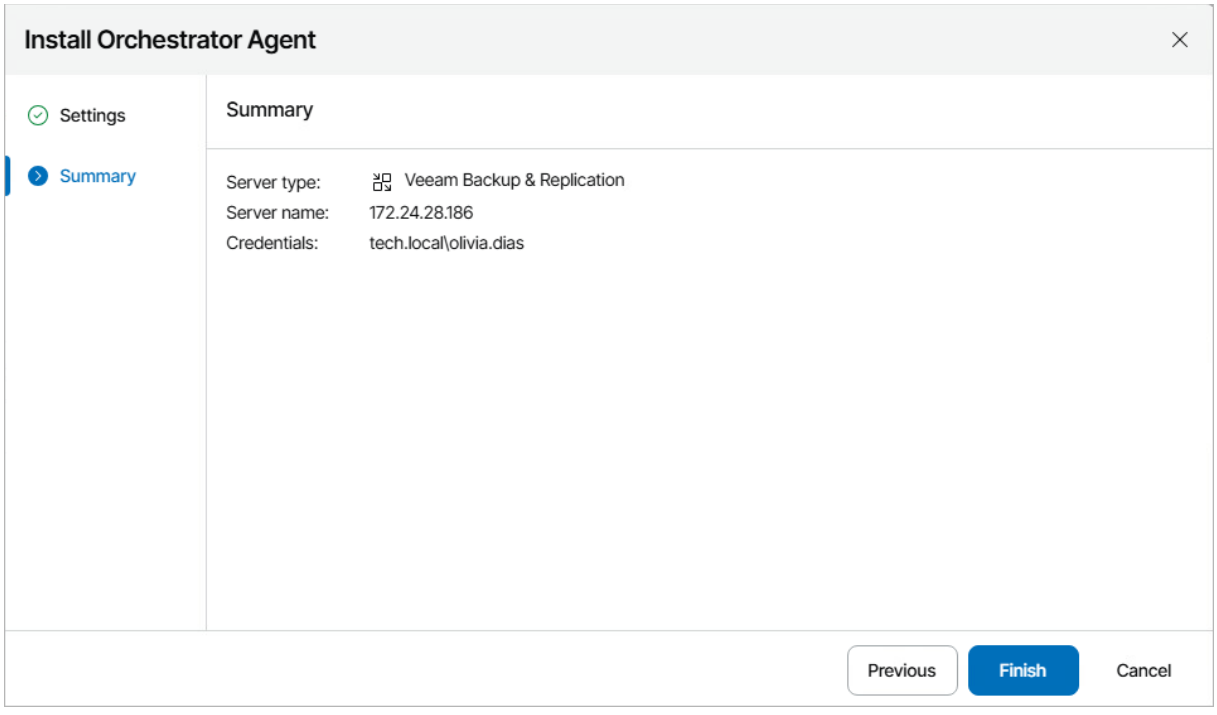

Note that after you deploy an Orchestrator agent on a remote Veeam Backup & Replication server or perform any infrastructure configuration changes, the changes may not appear in the Orchestrator UI immediately — the data synchronization process between Orchestrator and Veeam Backup & Replication may take up to 15 minutes to complete.

IMPORTANT

If you connect a Linux-based Veeam Backup & Replication server that protects Windows-based machines, you must specify a Windows-based server as the default mount backup server in Veeam Backup & Replication to be able to recover these machines. For more information, see the Veeam Backup & Replication User Guide, section Mount Servers.

Uninstalling Orchestrator Agents

If you no longer need a Veeam Backup & Replication server to be connected to Orchestrator, you can uninstall the Orchestrator agent running on the server.

1. Select the Orchestrator agent and click **Uninstall Agent**.
2. The **Dependency Check** window will inform you if any DataLabs or recovery plans are related to the Veeam Backup & Replication server.
   - If any of the items occur to be *Locked*, Orchestrator will not be able to uninstall the Orchestrator agent.

     In this case, wait until Orchestrator stops processing the items, power off plan testing in the locked DataLabs, reset the locked recovery plans — and then try uninstalling the Orchestrator agent again.
   - If none of the items are *Locked*, click **Continue** to confirm the operation.

IMPORTANT

As soon as you uninstall the Orchestrator agent from the Veeam Backup & Replication server, all its related recovery plans will no longer be able to run, and all its related DataLabs will be removed from Orchestrator.

All template jobs and credentials collected from the server will also be excluded from Orchestrator components, and the Protect VM Group steps that use any of these jobs will be removed from recovery plans as well.

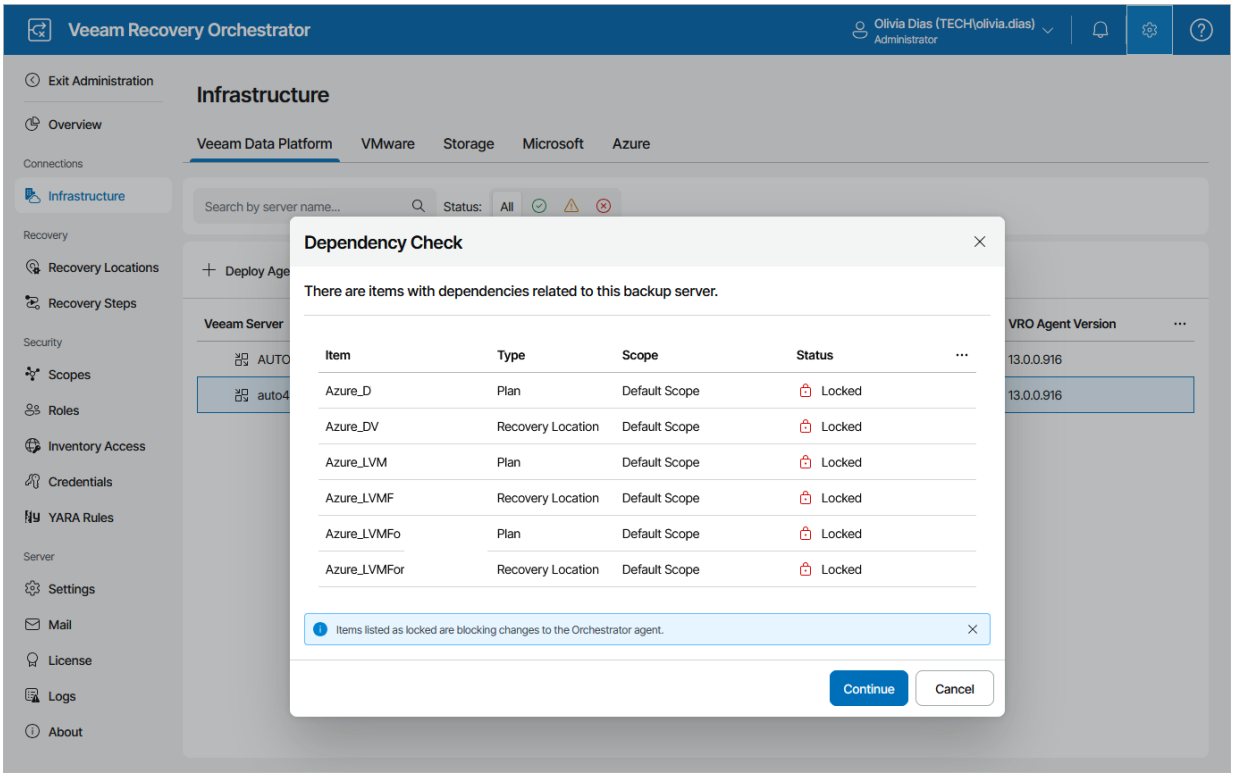

NOTE

You cannot uninstall an Orchestrator agent running on a Veeam Backup & Replication server that is used by any Microsoft Azure recovery location. Modify the settings of all the related locations to remove references to the server — and then try uninstalling the agent again.

3. Complete the **Uninstall Orchestrator Agent** wizard:

   a. At the **Settings** step of the wizard, specify the following settings:

      i. Choose whether you want to uninstall the Orchestrator agent regardless of the Veeam Backup & Replication server state.

      If Orchestrator is not able to access the server, the Orchestrator agent will be removed from the Orchestrator database but will keep running on the Veeam Backup & Replication server.

ii. From the **Credentials** drop-down list, select the necessary account for connecting to the server. For an account to be displayed in the **Credentials** list, it must be added to the configuration database as described in section Adding Credentials. If you have not set up an account beforehand, click **Add** and follow the steps of the **Add Credential** wizard. For more information on the required account permissions, see the Veeam Recovery Orchestrator Deployment Guide, section Permissions.

The provided credentials will be also used to launch the Orchestrator agent on the server. The user name must be specified in the *DOMAIN|USERNAME* or *USERNAME* format.

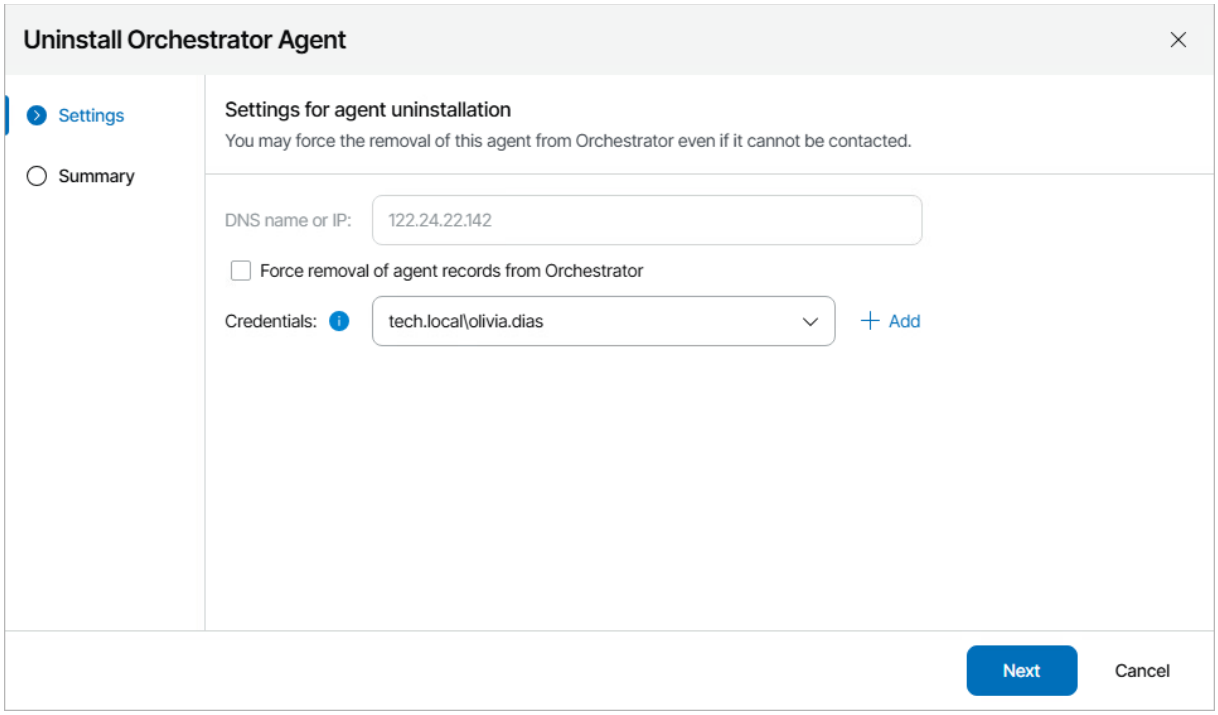

b. At the **Summary** step, review the configured settings and click **Finish**.

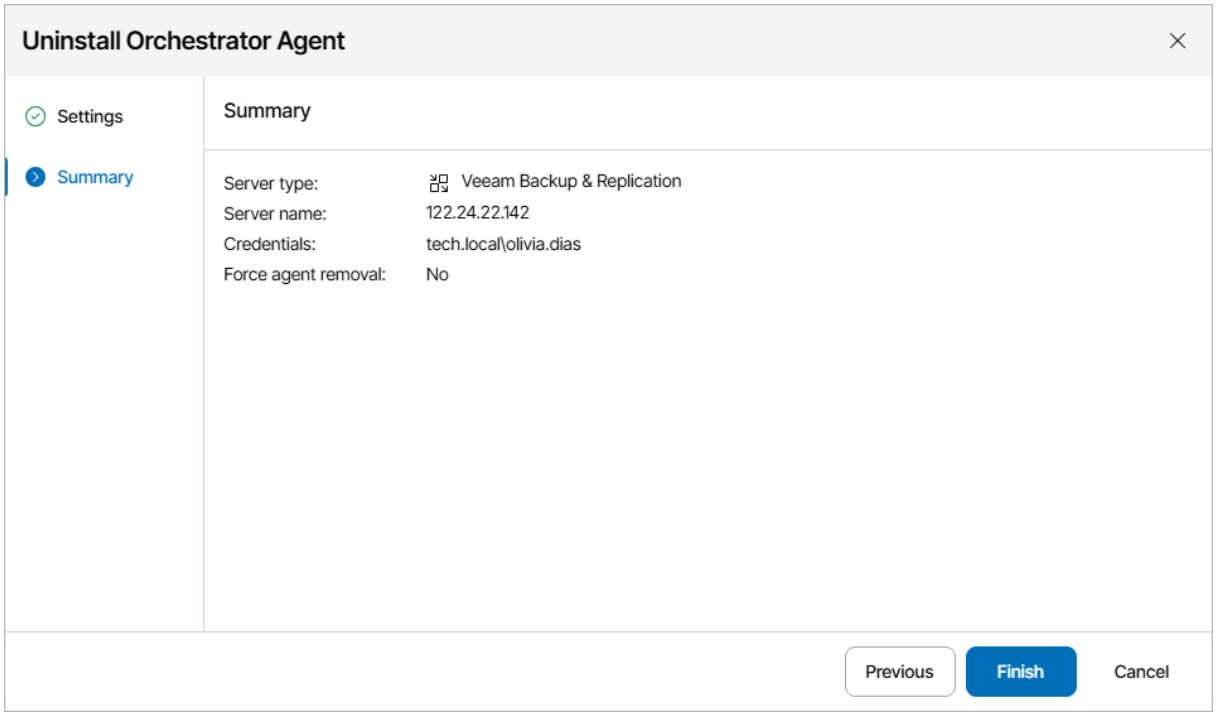

NOTE

If a Veeam Backup & Replication server is managed by Veeam Backup Enterprise Manager, you will not be able to uninstall the Orchestrator agent running on the server using the Orchestrator UI. In this case, remove the Veeam Backup & Replication server from Enterprise Manager as described in the Veeam Backup Enterprise Manager Guide. After you remove the server from Enterprise Manager, it will be automatically removed from Orchestrator as well.

Repairing Orchestrator Agents

If you want to change credentials of a user account that you specified when connecting a Veeam Backup & Replication server to Orchestrator, or if you encounter a connection issue and a Veeam Backup & Replication server acquires the *Disconnected* state, you can repair the Orchestrator agent running on the server.

NOTE

If you upgrade a remote Veeam Backup & Replication server connected to Orchestrator, you must also repair the Orchestrator agent running on the server. However, it is recommended that you wait approximately 10 minutes before repairing the agent because Veeam ONE Client has to process infrastructure changes related to the upgrade.

1. Select the Orchestrator agent and click **Repair Agent**.
2. The **Dependency Check** window will inform you if any DataLabs or recovery plans are related to the Veeam Backup & Replication server.
   - If any of the items occur to be *Locked*, Orchestrator will not be able to repair the Orchestrator agent.

     In this case, wait until Orchestrator stops processing the items, power off plan testing in the locked DataLabs, reset the locked recovery plans — and then try repairing the Orchestrator agent again.

- If none of the items are *Locked*, click **Continue** to confirm the operation.

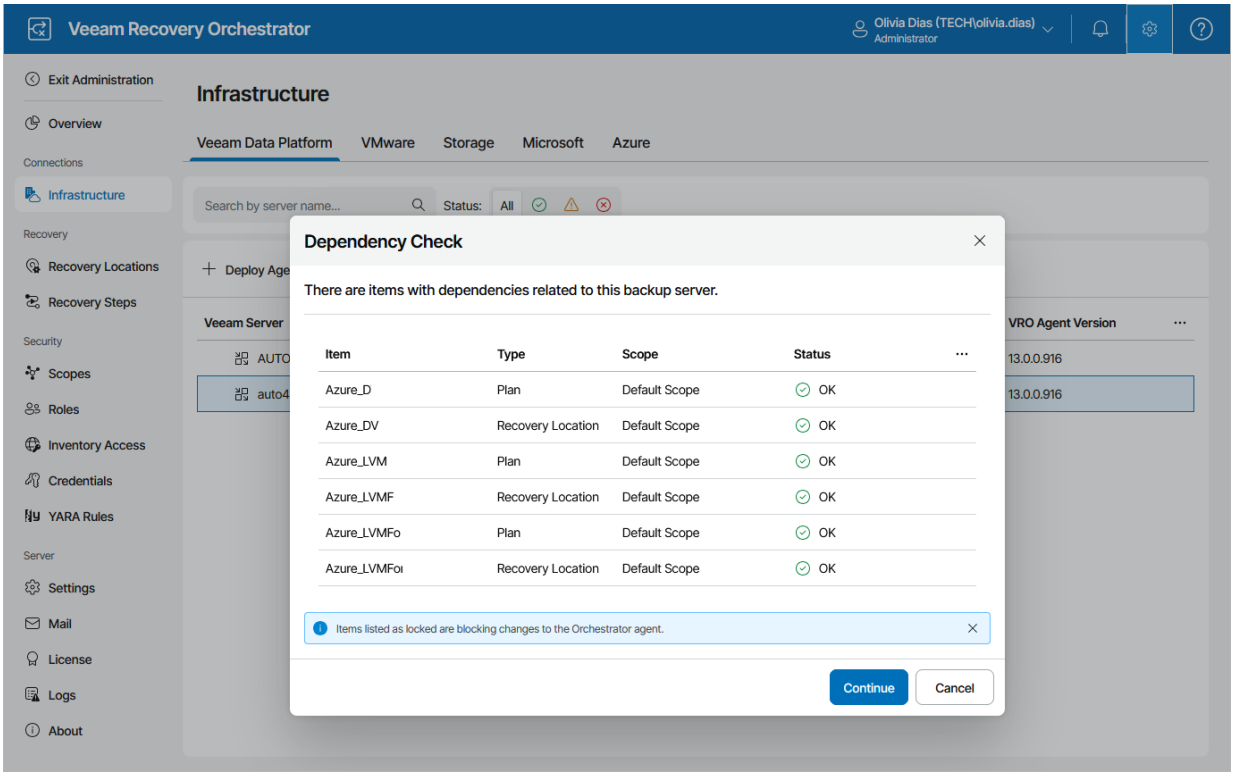

3. Complete the **Repair Orchestrator Agent** wizard:

   a. At the **Settings** step, select the necessary account from the **Credentials** drop-down list for connecting to the Veeam Backup & Replication server. For an account to be displayed in the **Credentials** list, it must be added to the configuration database as described in section Adding Credentials. If you have not set up an account beforehand, click **Add** and follow the steps of the **Add Credential** wizard. For more information on the required account permissions, see the Veeam Recovery Orchestrator Deployment Guide, section Permissions.

   The provided credentials will be also used to launch the Orchestrator agent on the server. The user name must be specified in the *DOMAIN|USERNAME* or *USERNAME* format.

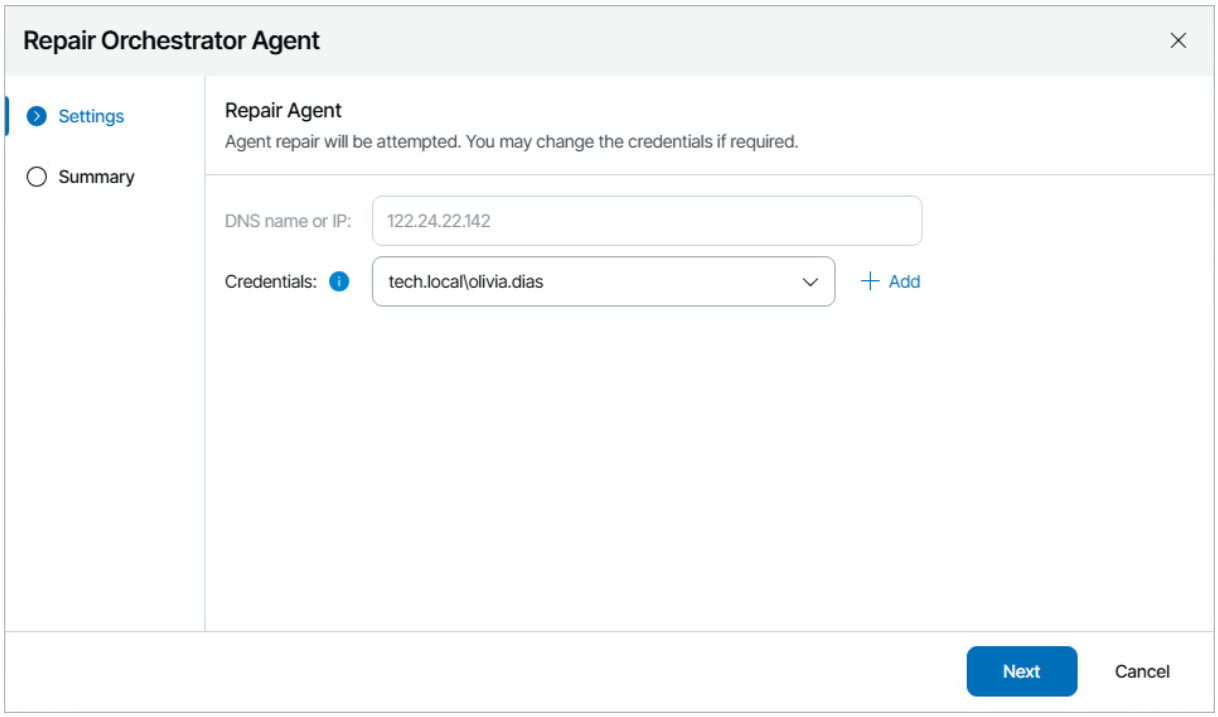

   b. At the **Summary** step, review the configured settings and click **Finish**.

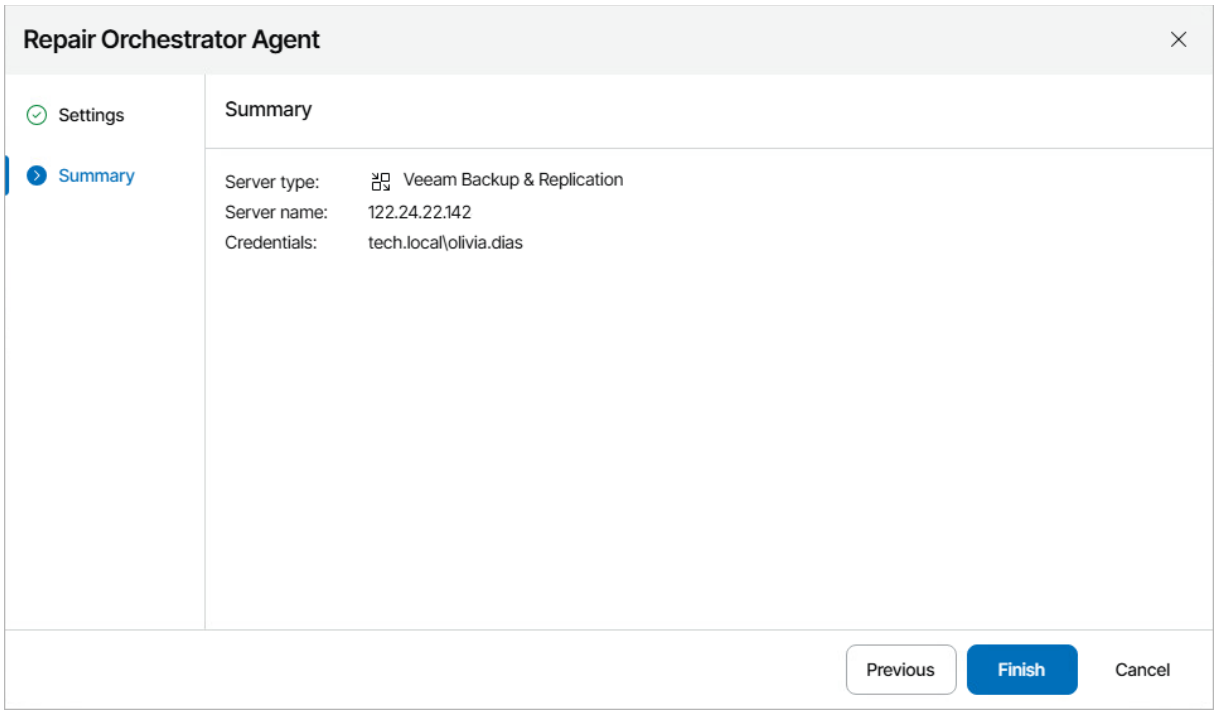

Connecting VMware vSphere Servers

To collect data about VMware vSphere infrastructure objects, you must configure connections to VMware vSphere servers. Only connections to vCenter Servers are supported.

> **IMPORTANT**
>
> To allow Orchestrator to process resource groups properly, you must connect vCenter Servers — not standalone ESXi hosts — to Veeam Backup & Replication servers orchestrated by Orchestrator. To learn how to add vCenter Servers to the backup infrastructure, see the Veeam Backup & Replication User Guide, section Virtualization Servers and Hosts.

It is required that Orchestrator has connections to the following vCenter Servers:

- All vCenter Servers managing source and replica VMs.
- The vCenter Server managing resources that will be used to recover VMs.
- All vCenter Servers managing VMs whose disks are located on the connected storage systems.

> **IMPORTANT**
>
> Starting from VMware vSphere version 7.0 Update 3, vSphere Cluster Services (vCLS) is activated by default. Before you connect to a vCenter Server version 7.0 Update 3, you must configure vCLS datastore placement in the vSphere Client as described in VMware vSphere documentation.

To configure a connection to a vCenter Server:

1. Switch to the **Administration** page.
2. Navigate to **Infrastructure** > **VMware**.
3. Click **Add**.
4. Complete the **Connect vCenter Server** wizard:
    a. At the **Settings** step of the wizard, specify the following connection settings:
        i. Use the **DNS name or IP** field to enter a DNS name or an IPv4 address of the vCenter Server that will be connected to the Orchestrator agent. The maximum length of the location name is 256 characters; the following characters are not supported: * : / \ ? @ [ ] ; : = + " < > | .

> **NOTE**
>
> If you want to add an vCenter Server that is part of a backup infrastructure already connected to Orchestrator, make sure that you add the server using the same DNS name or IPv4 address as in the backup infrastructure to avoid synchronization issues.

        ii. From the **Credentials** drop-down list, choose the necessary account for connecting to the server. For an account to be displayed in the **Credentials** list, it must be added to the configuration database as described in section Adding Credentials. If you have not set up an account beforehand, click **Add** and follow the steps of the **Add Credential** wizard. For more information on the required account permissions, see the Veeam Recovery Orchestrator Deployment Guide, section Permissions.

iii. If required, change the port number used for communication with the server.

If an untrusted security certificate is installed on the vCenter Server, you will get a security warning. You can view the certificate and click **Remember and Continue** — in this case, Orchestrator will remember the certificate thumbprint and will further trust the certificate when connecting to the server. Otherwise, you will not be able to add the server.

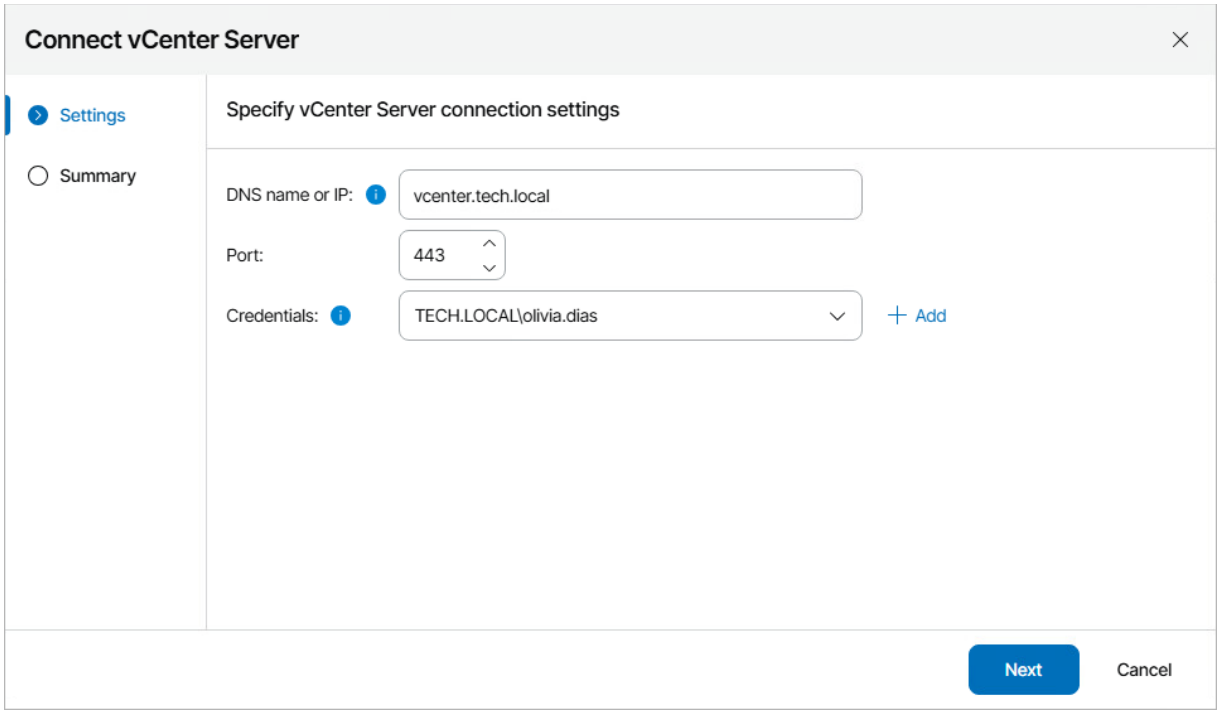

b. At the **Summary** step of the wizard, review the connection details and click **Finish**.

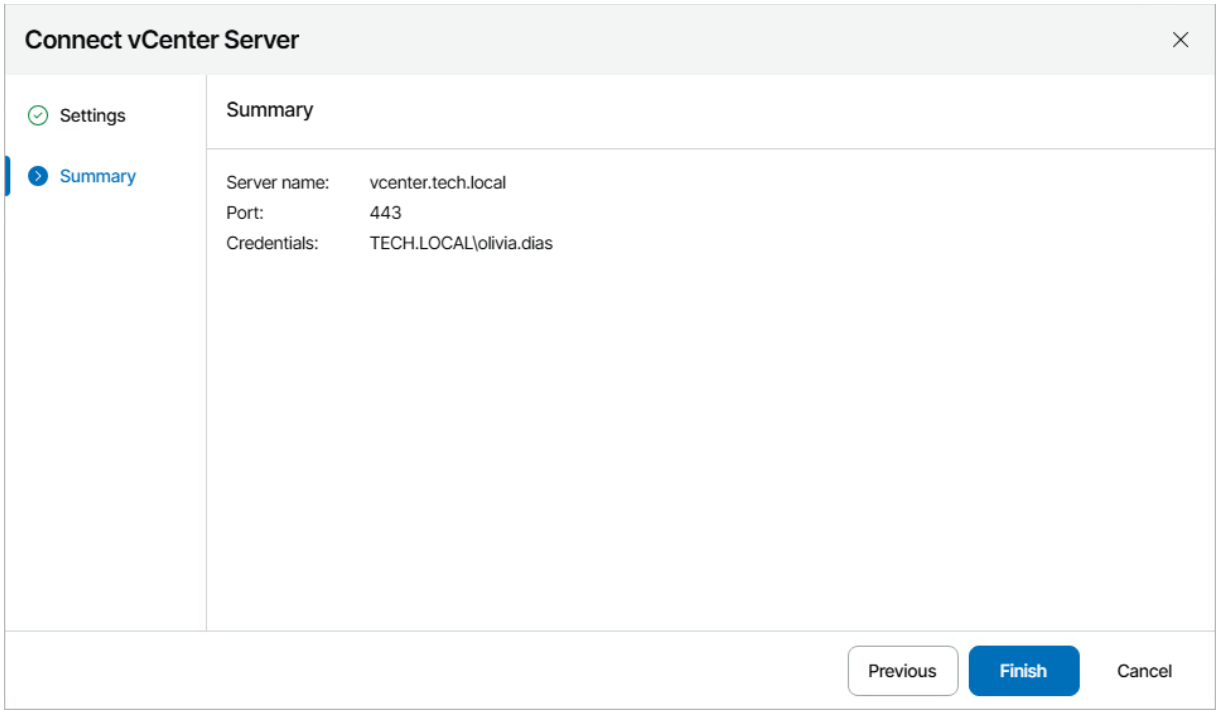

Note that after you configure a connection to a vCenter Server or perform any infrastructure configuration changes, the changes may not appear in the Orchestrator UI immediately — the data synchronization process between Orchestrator and VMware vSphere may take up to 15 minutes to complete.

Removing VMware vSphere Servers

If you no longer need a vCenter Server to be connected to Orchestrator, you can remove the server.

1. Select the vCenter Server and click **Remove**.
2. The **Dependency Check** window will inform you if any DataLabs or recovery plans are related to the vCenter Server.
   - If any of the items occur to be *Locked*, Orchestrator will not be able to remove the server.

     In this case, wait until Orchestrator stops processing the items, reset the locked recovery plans, power off plan testing in the locked DataLabs — and then try removing the vCenter Server again.
   - If none of the items are *Locked*, click **Continue** to confirm the operation.

> **IMPORTANT**
>
> As soon as you remove the vCenter Server from Orchestrator, all its related DataLabs will be removed from Orchestrator as well. All inventory groups that include VMs managed by the server will be excluded from Orchestrator components, and the VMs will be deleted from the related recovery plans.

3. Click **Yes** in the **Remove VMware vCenter Server** window.

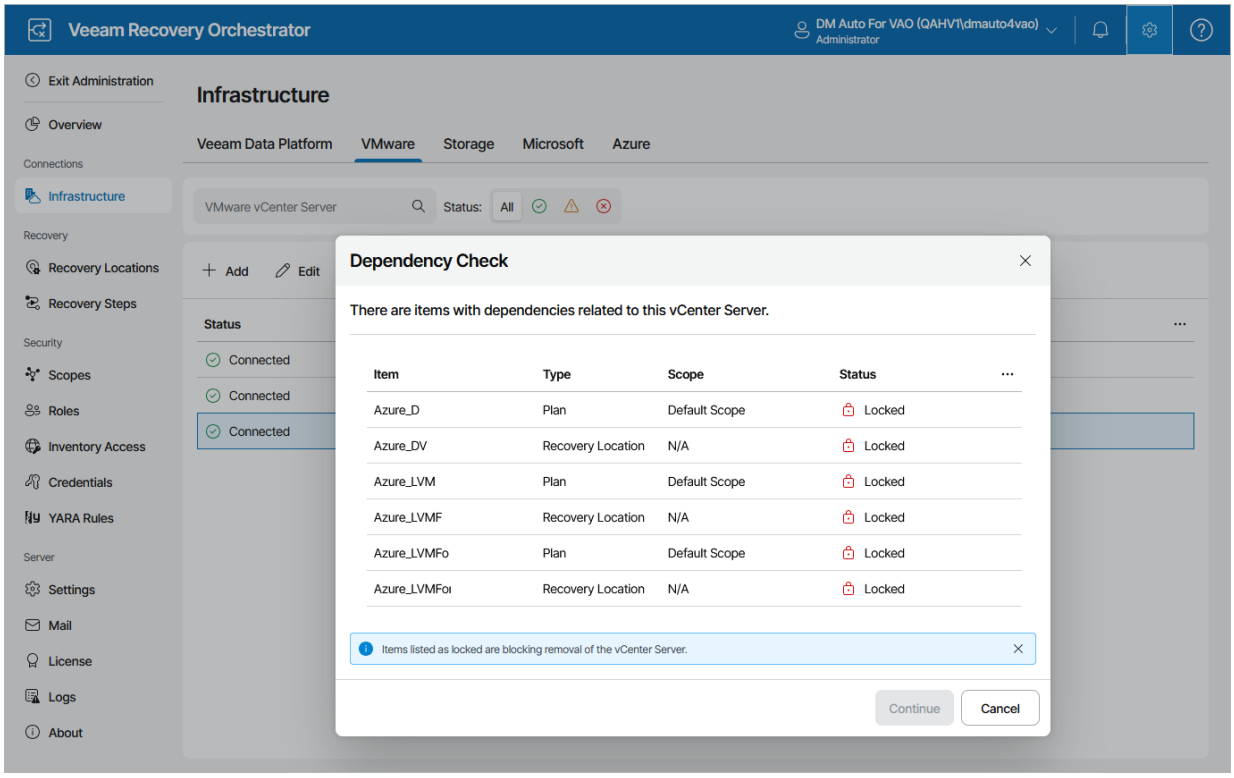

Connecting Microsoft Hyper-V Servers

To collect data about Microsoft Hyper-V infrastructure objects, you must configure connections either to Microsoft System Center Virtual Machine Manager (SCVMM) servers or standalone clusters.

Note that before you start adding a connection, it is recommended that you make sure that all the Microsoft Hyper-V infrastructure objects are online. Otherwise, if you bring these objects online after adding the connection, Veeam Recovery Orchestrator will not be able to collect their data immediately since the data synchronization process between Orchestrator and Microsoft Hyper-V may take up to 2 hours to complete.

IMPORTANT

- To allow Orchestrator to connect to an SCVMM server, you must first install the SCVMM console on the machine that runs Orchestrator as described in Microsoft Docs.

  Keep in mind that the SCVMM version of the console must match the System Center version of the server.
- If the SCVMM server or standalone cluster that you want to connect to Orchestrator is added to a backup server, you must first upgrade the server or cluster to Veeam Backup & Replication version 12.3. Otherwise, the server or cluster will no longer be able to connect to the backup server, which may affect data protection.

  To upgrade Veeam Backup & Replication to version 12.3, follow the instructions provided in the Veeam Backup & Replication User Guide, section Upgrading to Veeam Backup & Replication 12.3, and in this Veeam KB article.

Connecting SCVMM Servers

To configure a connection to an SCVMM server:

1. Switch to the **Administration** page.
2. Navigate to **Infrastructure** > **Microsoft**.
3. Click **Add** > **SCVMM**.
4. In the **Add SCVMM Connection** window, do the following:

   a. Use the **DNS name or IP** field to enter a DNS name or an IPv4 address of the SCVMM server that will be connected to the Orchestrator agent. The maximum length of the location name is 128 characters.

NOTE

If you want to add an SCVMM server that is part of a backup infrastructure already connected to Orchestrator, make sure that you add the server using the same DNS name or IPv4 address as in the backup infrastructure to avoid synchronization issues.

   b. From the **Credentials** drop-down list, choose the necessary account for connecting to the SCVMM server. For an account to be displayed in the **Credentials** list, it must be added to the configuration database as described in section Adding Credentials. If you have not set up an account beforehand, click **Add** and follow the steps of the **Add Credential** wizard. For more information on the required account permissions, see the Veeam Recovery Orchestrator Deployment Guide, section Permissions.

   d. Click **Save**.

Note that after you configure a connection to an SCVMM server or perform any infrastructure configuration changes, the changes may not appear in the Orchestrator UI immediately — the data synchronization process between Orchestrator and Microsoft Hyper-V may take up to 15 minutes to complete.

Connecting Hyper-V and Azure Local Clusters

> **IMPORTANT**
>
> Each cluster must be added to the Orchestrator infrastructure only once — either as a direct connection or as part of an SCVMM hierarchy.

To configure a connection to a standalone Hyper-V or Azure Local cluster:

1. Switch to the **Administration** page.
2. Navigate to **Infrastructure** > **Microsoft**.
3. Click **Add** > **Hyper-V Cluster**.
4. In the **Add Hyper-V Cluster** window, do the following:
   a. Use the **DNS name or IP** field to enter a DNS name or an IPv4 address of the cluster that will be connected to the Orchestrator agent. The maximum length of the location name is 128 characters.

> **NOTE**
>
> If you want to add a Hyper-V or Azure Local cluster that is part of a backup infrastructure already connected to Orchestrator, make sure that you add the cluster using the same DNS name or IPv4 address as in the backup infrastructure to avoid synchronization issues.

   b. From the **Credentials** drop-down list, choose the necessary account for connecting to the cluster. For an account to be displayed in the **Credentials** list, it must be added to the configuration database as described in section Adding Credentials. If you have not set up an account beforehand, click **Add** and follow the steps of the **Add Credential** wizard. For more information on the required account permissions, see the Veeam Recovery Orchestrator Deployment Guide, section Permissions.
   d. Click **Save**.

Note that after you configure a connection to a standalone cluster or perform any infrastructure configuration changes, the changes may not appear in the Orchestrator UI immediately — the data synchronization process may take up to 15 minutes to complete.

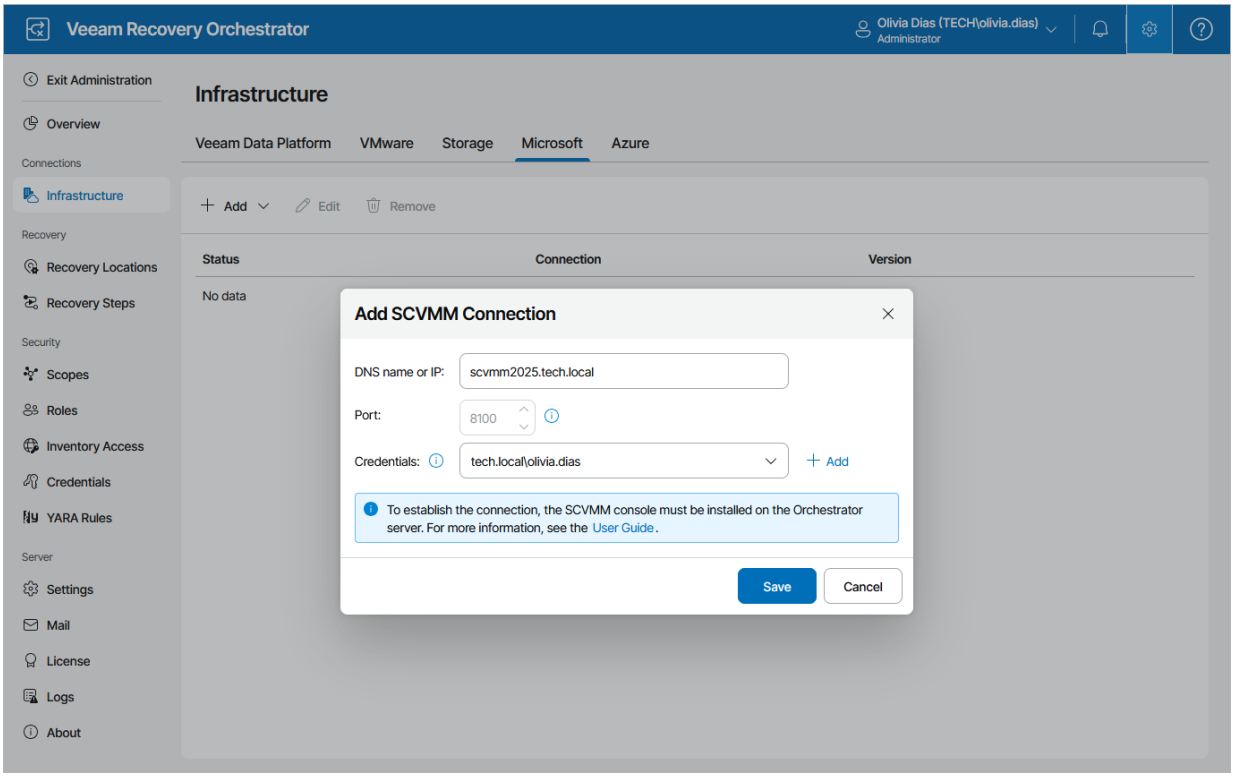

Removing Microsoft Hyper-V Servers

If you no longer need an SCVMM server or a standalone cluster to be connected to Orchestrator, you can remove it.

1. Select the SCVMM server or cluster and click **Remove**.
2. The **Dependency Check** window will inform you if any recovery locations or recovery plans are related to the connection.
   - If any of the items occur to be *Locked*, Orchestrator will not be able to remove the connection.

     In this case, wait until Orchestrator stops processing the items, remove this SCVMM server or cluster from the locked recovery locations, reset the locked recovery plans — and then try removing the connection again.
   - If none of the items are *Locked*, click **Continue** to confirm the operation.

> **IMPORTANT**
>
> As soon as you remove the connection from Orchestrator, all VMs will be deleted from the related recovery plans.

3. Click **Remove** in the **Remove Connection** window.

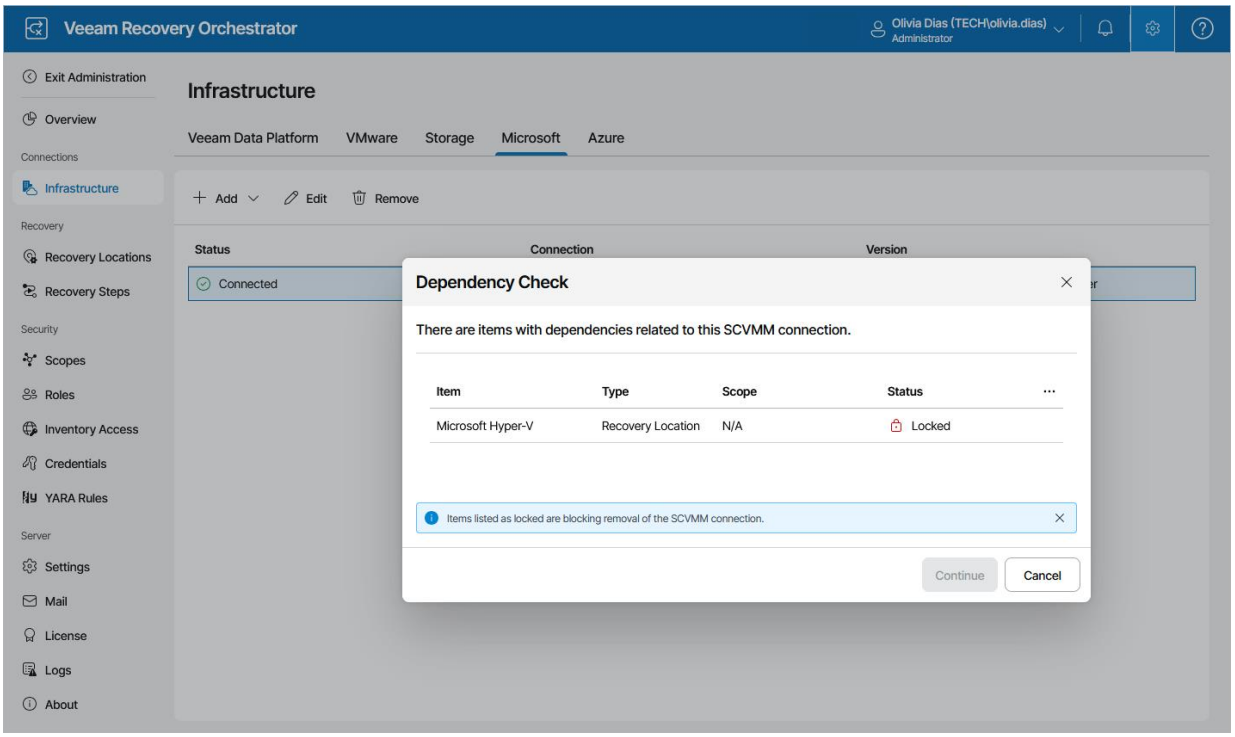

Connecting Storage Systems

The following NetApp storage systems must be connected to Orchestrator:

- Any active storage virtual machine (SVM) that will be the source of the datastore disaster recovery relationship.
- Any active SVM that will be the destination of the SVM disaster recovery relationship.

> **IMPORTANT**
>
> Make sure that you have NetApp volumes and virtual volumes protected by storage replication. If you use authentication configured to control access to iSCSI targets, make sure that you define a list of initiators and their authentication methods for all hosts managed by the target vCenter Server. For more information on iSCSI initiator security management, see the NetApp ONTAP Documentation Center. For more information on configuring CHAP parameters for iSCSI adapters, see VMware Docs.

The following HPE storage systems must be connected to Orchestrator:

- Any active storage system that will be the primary system of the remote copy configuration.
- Any active storage system that will be the secondary system of the remote copy configuration.

> **IMPORTANT**
>
> Consider the following:
>
> - To connect a NetApp storage system, the ONTAPI REST API (ZAPI) must be enabled.
> - To connect an HPE storage system, you must first enable the HPE 3PAR Web Services API (WSAPI) server as described in section Enabling HPE 3PAR Web Services API Server.
> - Connecting HPE storage systems paired in synchronous long distance (SLD) Remote Copy configurations is not supported. For more information on SLD configurations, see the Hewlett Packard Enterprise Support Center.

To configure a connection to a storage system:

1. Switch to the **Administration** page.
2. Navigate to **Storage**.
3. Click **Add**.

4. Complete the **Connect Storage System** wizard:

   a. At the **Storage Vendor** step of the wizard, choose whether you want to connect a NetApp or an HPE storage system.

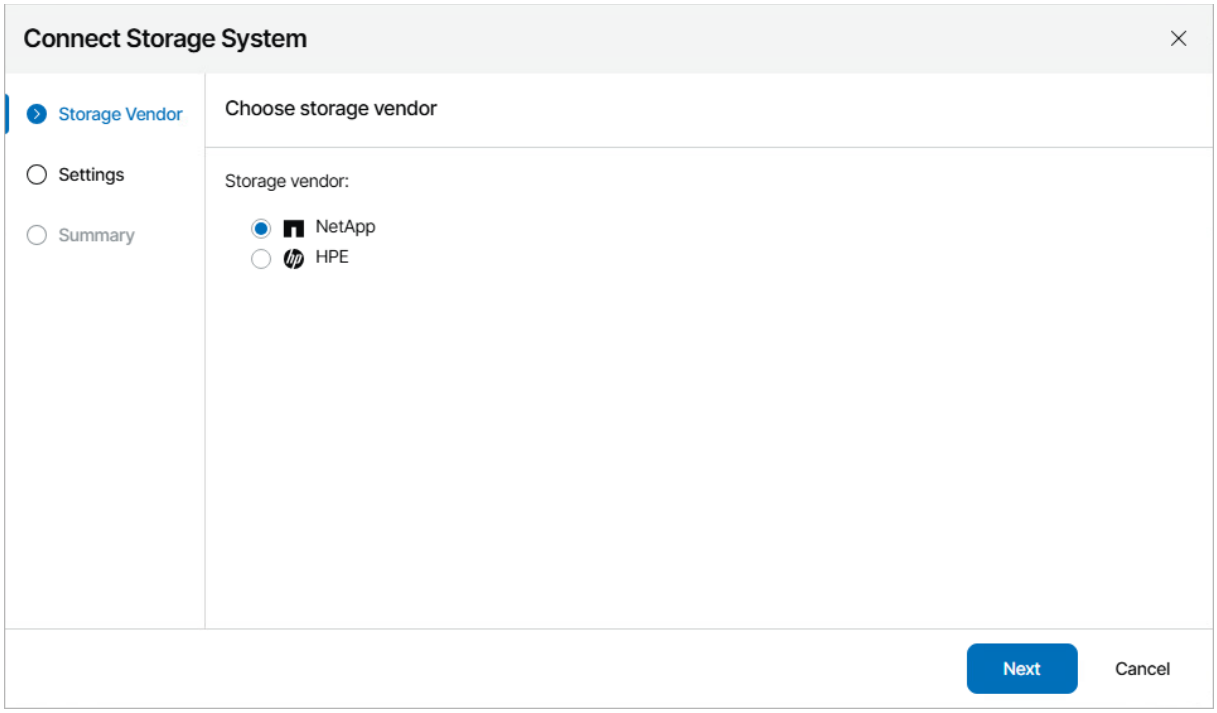

   b. At the **Settings** step of the wizard, specify the following connection settings:

      i. Use the **DNS name or IP** field to enter a DNS name or an IPv4 address of the storage system that will be connected to the Orchestrator server. The maximum length of the location name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

> **NOTE**
>
> If you want to add a storage system that is part of a backup infrastructure already connected to Orchestrator, make sure that you add the storage system using the same DNS name or IPv4 address as in the backup infrastructure to avoid synchronization issues.

      ii. From the **Credentials** drop-down list, choose the necessary account for connecting to the storage system.

      For an account to be displayed in the **Credentials** list, it must be added to the configuration database as described in section Adding Credentials. If you have not set up an account beforehand, click **Add** and follow the steps of the **Add Credential** wizard. For more information on the required account permissions, see the Veeam Recovery Orchestrator Deployment Guide, section Permissions.

iii. If required, change the port number used for communication with the system.

If an untrusted security certificate is installed on the storage system, you will get a security warning. You can view the certificate and click **Remember and continue** — in this case, Orchestrator will remember the certificate thumbprint and will further trust the certificate when connecting to the storage system. Otherwise, you will not be able to proceed with the wizard.

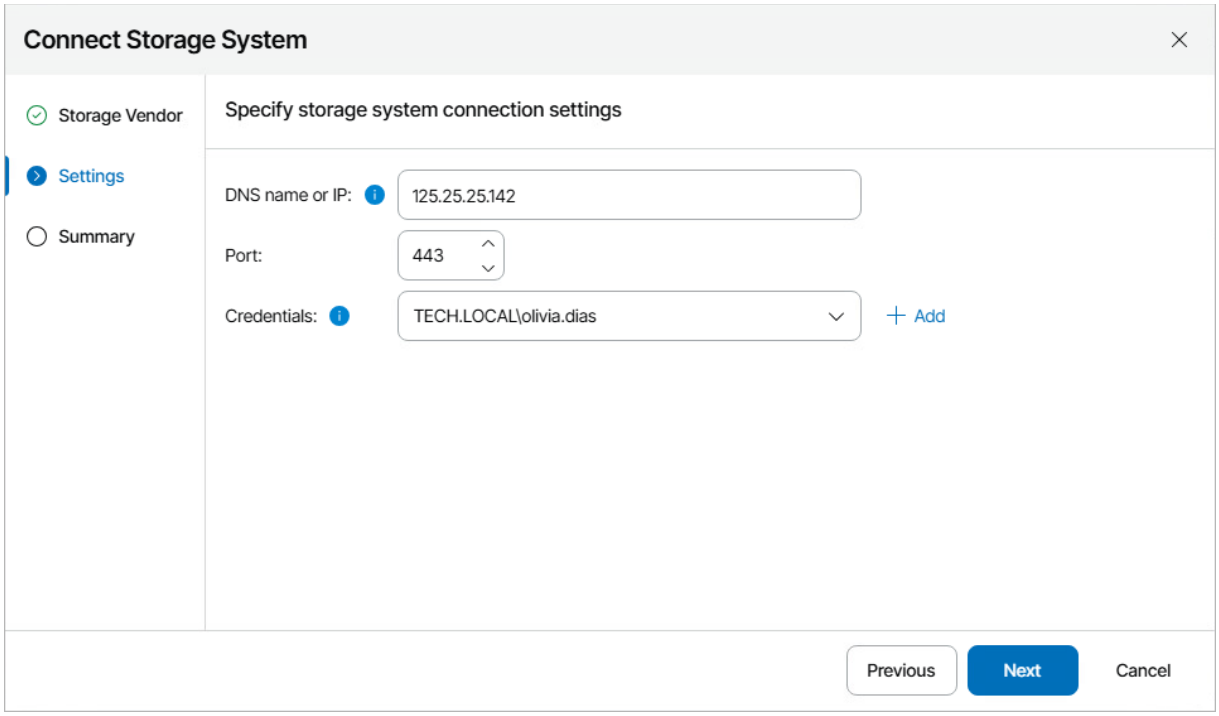

c. At the **Summary** step of the wizard, review the connection details and click **Finish.**

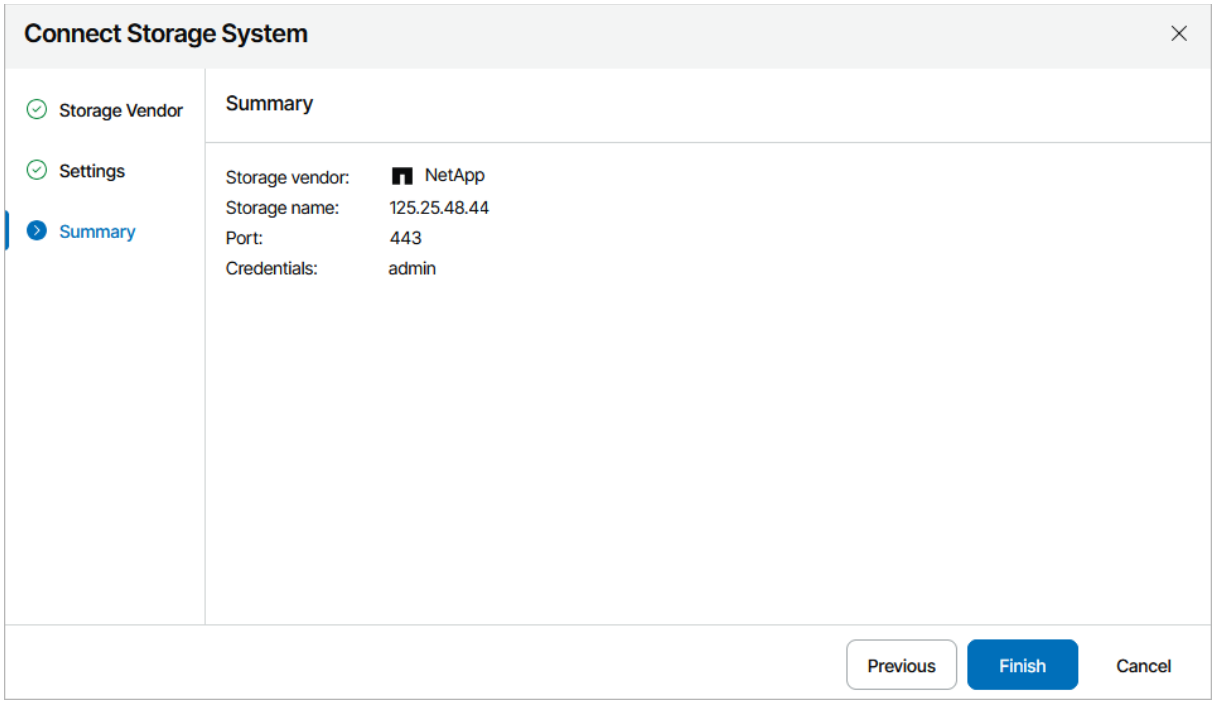

Note that after you configure a connection to a storage system or perform any infrastructure configuration changes, the changes may not appear in the Orchestrator UI immediately — the data synchronization process between Orchestrator and the storage infrastructure may take up to 15 minutes to complete.

IMPORTANT

After you connect a storage system to Orchestrator, you must include the system in a storage recovery location (either new or already existing) so that Orchestrator can use this storage system when executing and testing storage plans. To learn how to include target storage systems in storage recovery locations, see Adding Storage Recovery Locations.

Enabling HPE 3PAR Web Services API Server

Orchestrator uses the HPE 3PAR Web Services API (WSAPI) server to communicate with HPE Primera and HPE 3PAR storage systems. When you add an HPE Primera or HPE 3PAR storage system to Orchestrator, it does not enable the WSAPI server automatically because it does not have enough privileges to perform this operation. This means that you must enable the server manually before connecting the storage system to Orchestrator.

To enable the WSAPI server running on an HPE Primera or HPE 3PAR storage system:

1. Log on to the Processor with administrator privileges:

```
#ssh <administrator account>@<SP IP Address>
```

2. Run the following command to view the current state of the WSAPI server:

```
#showwsapi
-- -State- -HTTP_State-
HTTP_Port -HTTPS_State- HTTPS_Port -Version-
Enabled   Active Enabled       8008
Enabled       8080         1.1
```

3. If the WSAPI server is not running, run the following command to start it:

```
#startwsapi
```

If the HTTPS port is disabled, run the following command to enable it:

```
#setwsapi -https enable
```

Connecting Microsoft Azure Servers

No additional connection is required for recovery to Microsoft Azure as Orchestrator will use the Microsoft Azure compute and storage credentials configured on the connected Veeam Backup & Replication servers.

However, if you plan to execute custom scripts inside an Microsoft Azure machine managed by the connected Veeam Backup & Replication server, you must configure a direct connection to the Microsoft Azure compute account added to this server. To do that:

1. Switch to the **Administration** page.
2. Navigate to **Infrastructure** > **Azure**.
3. Choose the required connection and click **Configure connection**.
4. In the **Configure Direct Connection to Azure** window, do the following:
   a. Set the **Direct connection** toggle to *On*.
   b. Choose one of the following options:
      - If you want to use password-based authentication, select the **Secret** option and specify the secret ID of the connected Microsoft Azure compute account.
      - If you want to use certificate-based authentication, select the **Certificate** option, browse to a local folder to locate the self-signed public certificate and provide a password for this certificate.

      For more information on Microsoft Azure authentication, see Microsoft Docs.
   c. Click **Save**.

> **TIP**
>
> By default Orchestrator, verifies the connection to Microsoft Azure once a day to ensure that the provided certificate or secret ID is valid. However, you can also check the connection manually. To do that, choose the connection and click **Check connection**.

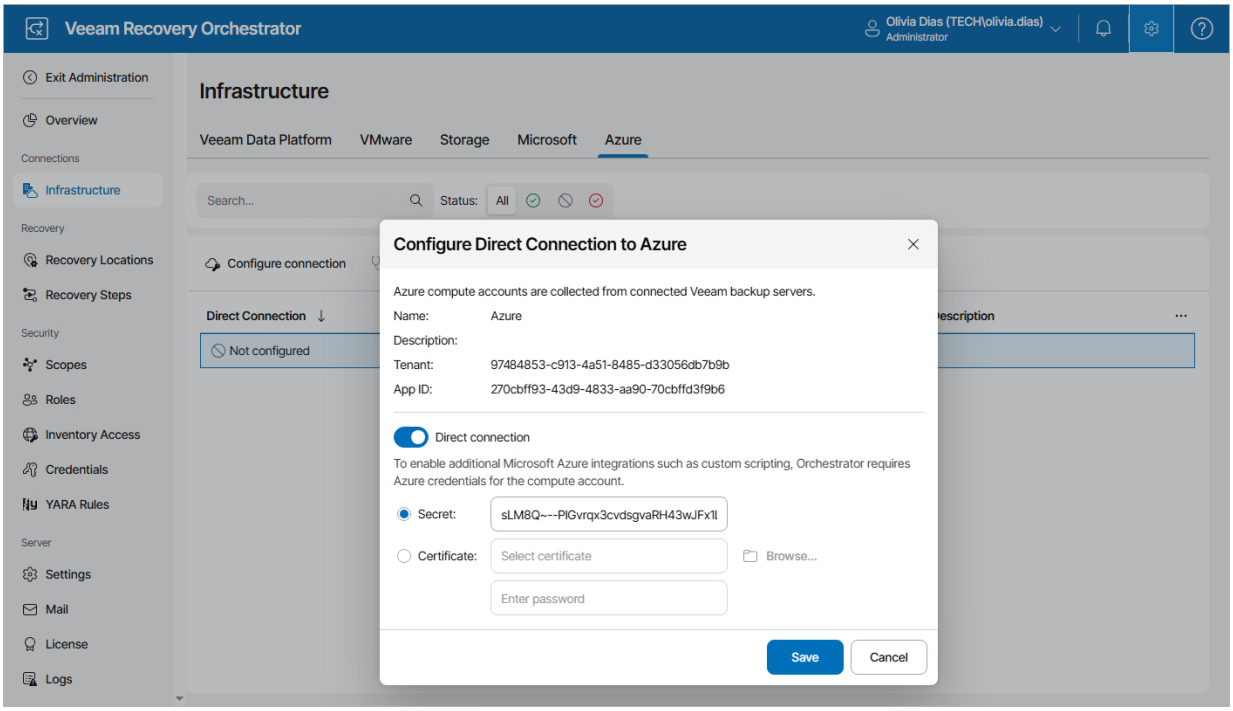
Veeam Recovery Orchestrator
Olivia Dias (TECH\olivia.dias)
Administrator
Exit Administration
Overview
Connections
Infrastructure
Recovery
Recovery Locations
Recovery Steps
Security
Scopes
Roles
Inventory Access
Credentials
YARA Rules
Server
Settings
Mail
License
Logs
Infrastructure
Veeam Data Platform
VMware
Storage
Microsoft
Azure
Search...
Status:
All
Configure connection
Direct Connection
Not configured
Configure Direct Connection to Azure
Azure compute accounts are collected from connected Veeam backup servers.
Name:
Azure
Description:
Tenant:
97484853-c913-4a51-8485-d33056db7b9b
App ID:
270cbff93-43d9-4833-aa90-70cbffd3f9b6
Direct connection
To enable additional Microsoft Azure integrations such as custom scripting, Orchestrator requires Azure credentials for the compute account.
Secret:
sLM8Q~~~PlGvrqx3cvdsgvaRH43wJFx1l
Certificate:
Select certificate
Browse...
Enter password
Save
Cancel

Configuring General Settings

The Orchestrator UI allows you to specify general settings for automated delivery of Orchestrator reports and documents.

Step 1. Specify Email Server Settings

To connect an SMTP server that will be used for sending email notifications:

1. Switch to the **Administration** page.
2. Navigate to **Mail**.
3. Click **Edit**.
4. In the **SMTP Server** window:
   a. In the **SMTP Server** field, enter a DNS name or an IPv4 address of the SMTP server. All email notifications (including test messages) will be sent by this SMTP server.
   b. In the **Port** field, change the SMTP communication port if required. The default SMTP port is **25**.
   c. In the **From address** field, enter an email address of the notification sender. This email address will be displayed in the **From** field of notifications.
   d. For an SMTP server with SSL/TLS support, select the **Use secure connection** check box to enable SSL data encryption.
   e. If your SMTP server requires authentication, select the **Requires authentication** check box, and specify authentication credentials in the **Username** and **Password** fields.
   f. The Orchestrator UI allows you to send a test message to check whether you have configured all settings correctly. To do that, click **Send test email**.
   g. Click **Save**.

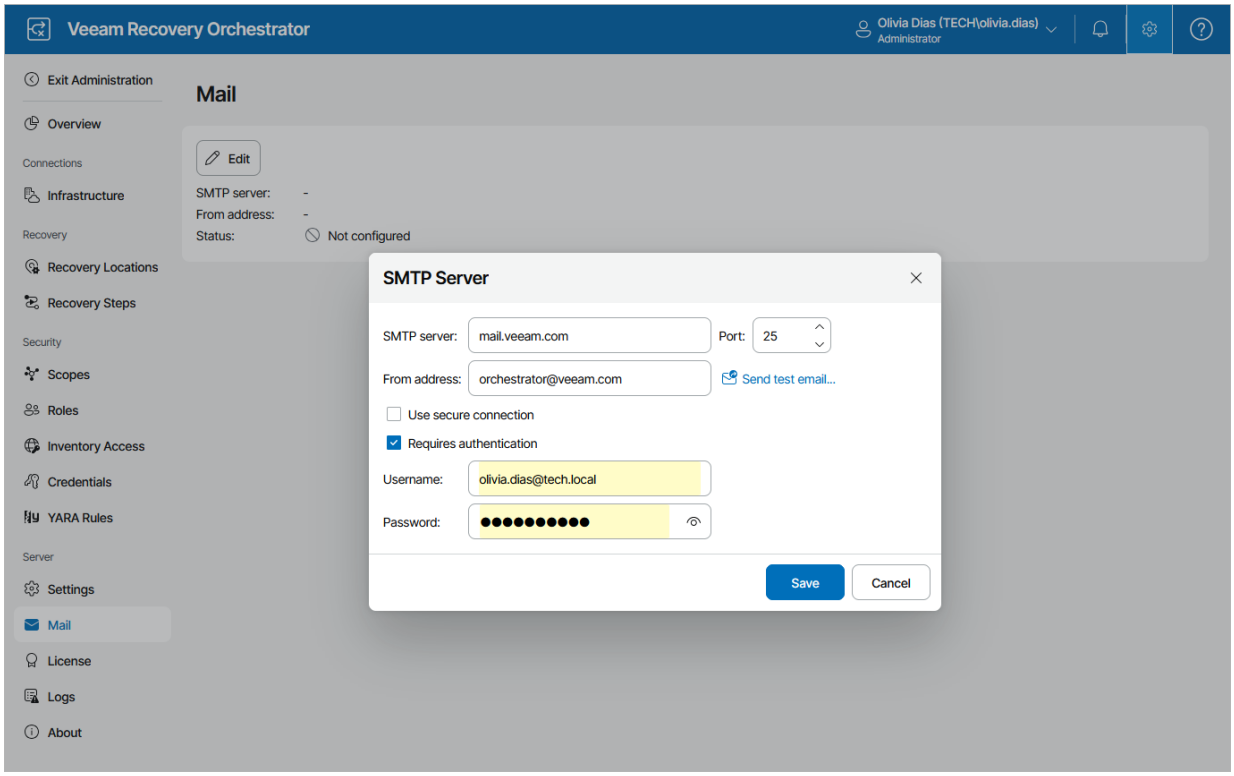

Step 2. Specify Report Subscription Settings

Add email addresses that will receive report notifications for your Orchestrator server. These addresses will be available for subscription for all created scopes.

To add recipients and subscribe them to the desired reports:

1. Navigate to **Reporting** > **Subscriptions**.
2. Click **Add**.
3. In the **Add Subscription** window:
   a. In the **Email address** field, enter an email address of a recipient.
   b. In the **Description** field, enter a short description for the recipient, if required.
   c. From the **Scope** drop-down list, select a scope for which you want to create subscriptions.
   d. Select check boxes next to the report types this address will be subscribed to.
   e. Select whether you want to subscribe to a summary or full report.
   f. The Orchestrator UI allows you to send a test message to check whether you have configured email settings correctly. To do that, click **Send Test Email**. A test email will be sent to the specified email address.
   g. Click **Apply**.

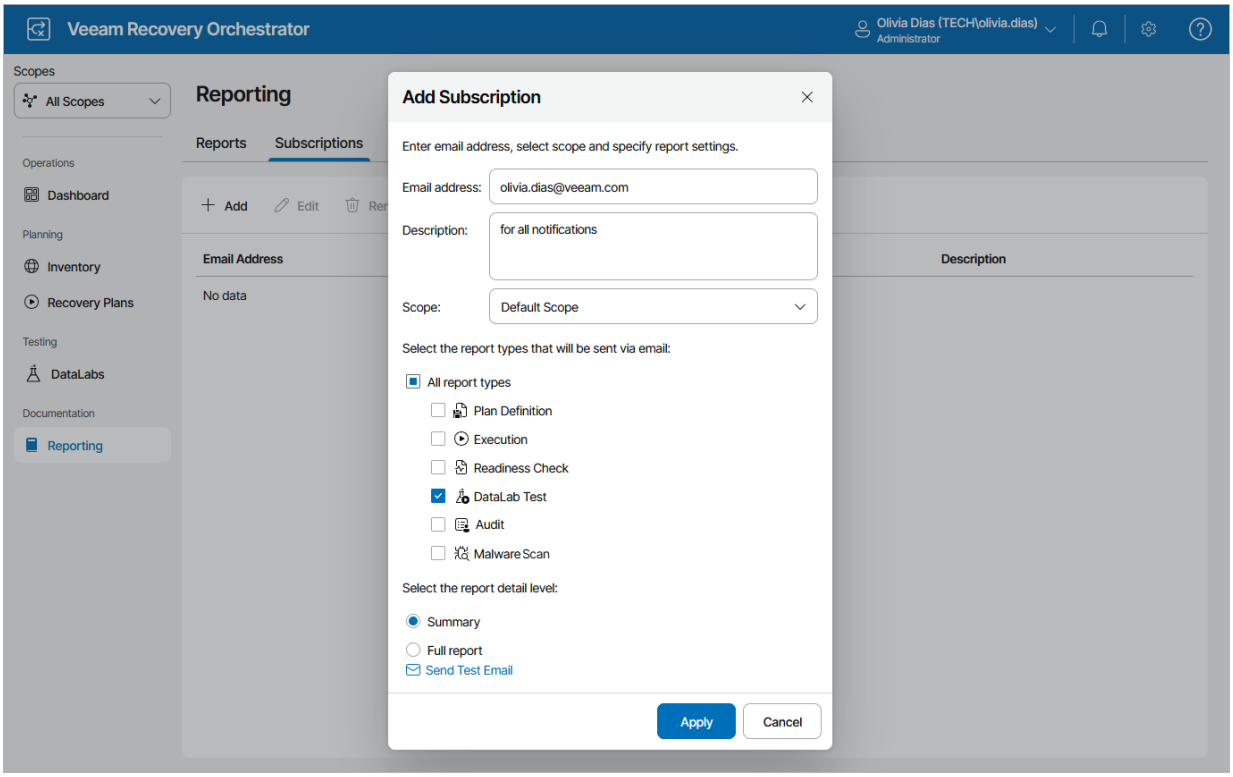

Step 3. Configure Report Retention Settings

By default, Orchestrator retains full reports for 1 year and summary reports for 3 years in its database. However, you can specify any other time period, which allows you to consume less storage space by deleting reports that are older than the specified period:

1. Switch to the **Administration** page.
2. Navigate to **Settings**.
3. In the **Report retention** section, specify for how long you want full and summary reports to be retained in the Orchestrator database.
4. Click **Save**.

> **IMPORTANT**
>
> The specified retention settings will apply to all Orchestrator reports in all scopes.

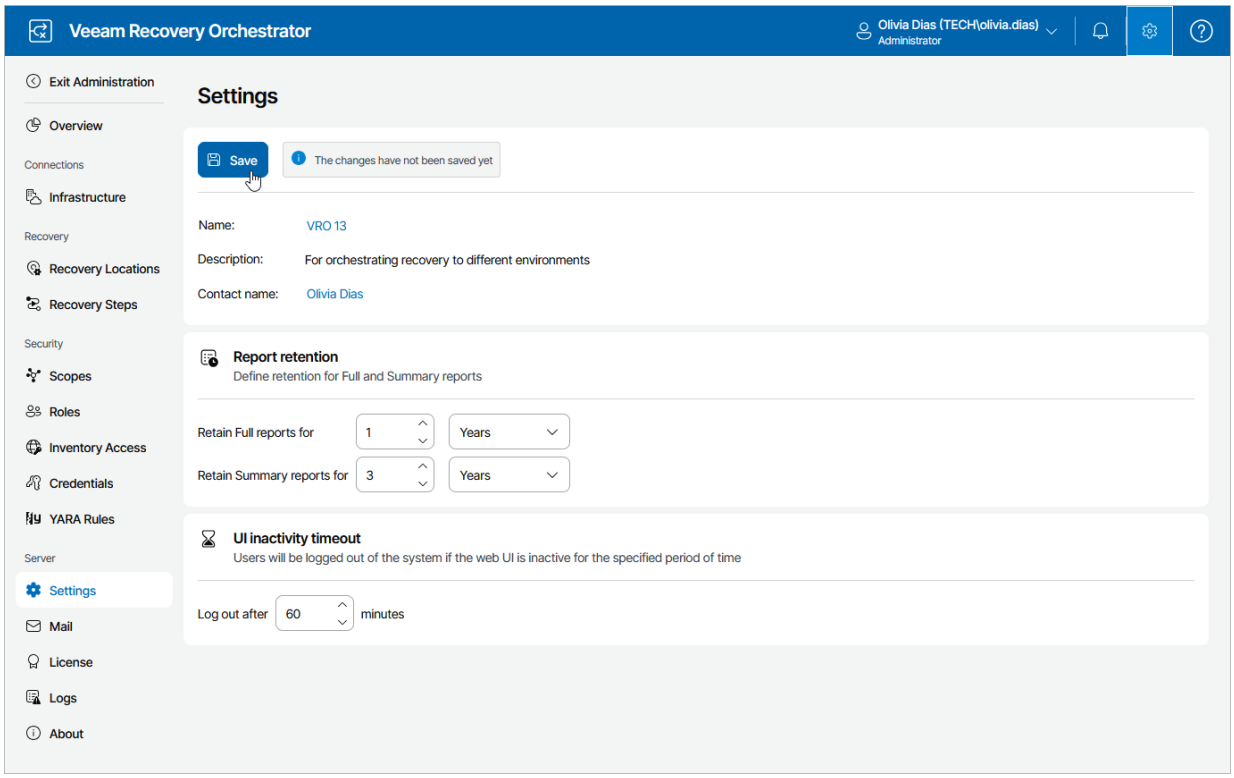

Managing Recovery Locations

Recovery locations consist of resource groups (that is, infrastructure objects such as vSphere hosts) used as target locations when orchestrating recovery. Resource groups are managed by the embedded Veeam ONE server installed on the Orchestrator server.

Orchestrator provides 4 types of recovery locations:

- VMware vSphere recovery locations that are used to define compute and storage resources in a VMware vSphere environment required when running restore plans and performing failback operations.
- Microsoft Hyper-V recovery locations that are used to define cluster and storage resources in a Microsoft Hyper-V environment required when running restore plans.
- Storage recovery locations that are used to define target storage systems and compute resources required when running storage plans.
- Microsoft Azure recovery locations that are used to define cloud resources required when running cloud plans.

Adding Recovery Locations

All resource groups created based on vCenter Server tags will automatically become available for the creation of recovery locations in the Orchestrator UI. However, after you assign a tag to an object in the vSphere inventory, the object may not be available in the Orchestrator UI immediately — the data synchronization process between Orchestrator and Veeam ONE occurs every 3 hours.

> **TIP**
>
> You can speed up the data synchronization process using Veeam ONE Reporter installed as part of the embedded Veeam ONE server. To do that:
>
> 1. In a web browser, navigate to the Veeam ONE Reporter web address.
>
>   The address consists of an FQDN of the Orchestrator server and the website port specified during installation (by default, **1239**). Note that Veeam ONE Reporter is available over HTTPS.
>
>   `https://<FQDN>:1239/`
>
> 2. Switch to the **Configuration** page.
> 3. Navigate to **Data Collection**.
> 4. From the **Advanced Actions** menu, select **Start report data collection**.

To add a recovery location to be used when performing a failback operation, follow the instructions provided in section Adding VMware vSphere Recovery Locations. To add a recovery location to be used when running a restore plan, follow the instructions provided in sections Adding VMware vSphere Recovery Locations and Adding Microsoft Hyper-V Recovery Locations. To add a recovery location to be used when running a storage plan, follow the instructions provided in section Adding Storage Recovery Locations. To add a recovery location to be used when running a cloud plan, follow the instructions provided in section Adding Microsoft Azure Recovery Locations.

Adding VMware vSphere Recovery Locations

To add a VMware vSphere recovery location:

1. Switch to the **Administration** page.
2. Navigate to **Recovery Locations**.
3. Click **Add**.
4. Complete the **New Recovery Location** wizard:
   a. Specify a recovery location name and description.
   b. Choose a recovery location type.
   c. Choose recovery options.
   d. Specify compute resources.
   e. Specify storage resources.
   f. Configure network mapping and re-IP rules.
   g. Finish working with the wizard.

Step 1. Specify Recovery Location Name and Description

At the **Location Name** step of the wizard, use the **Name** and **Description** fields to enter a name for the new location and to provide a description for future reference. The maximum length of the location name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

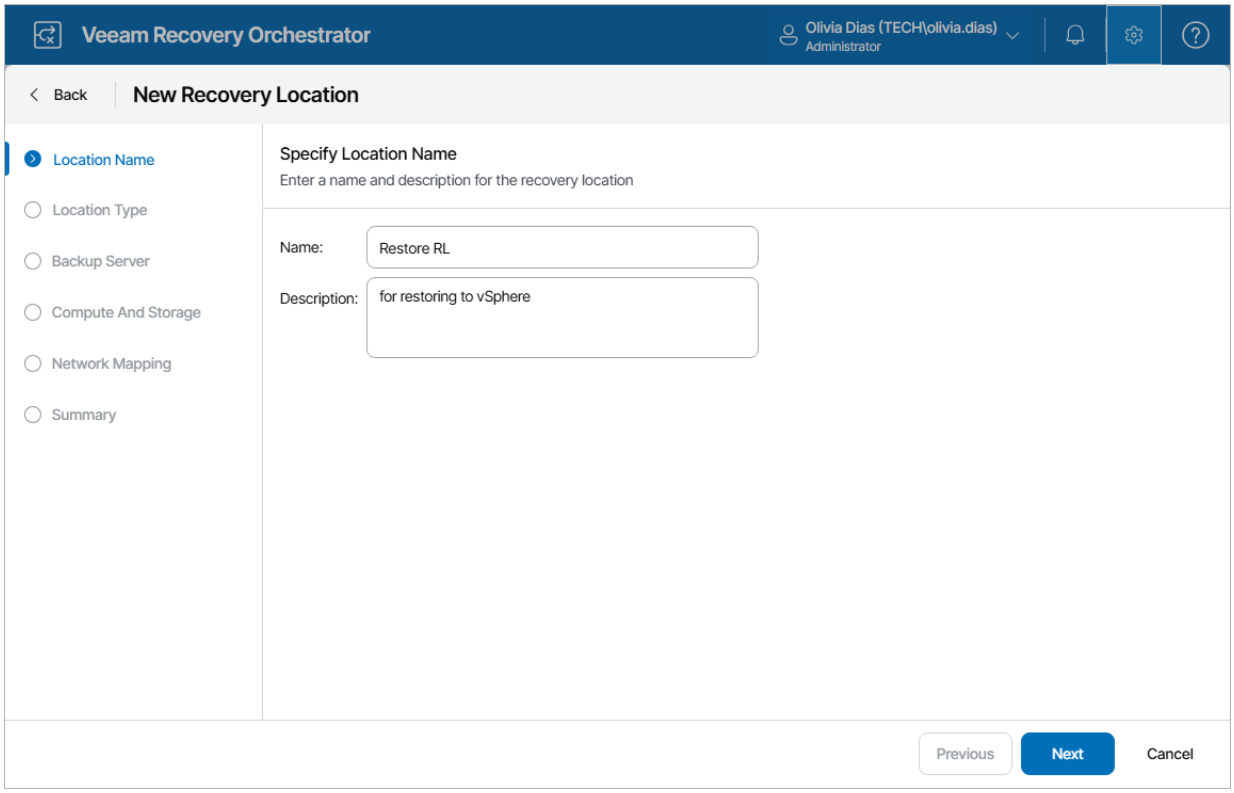

Step 2. Choose Recovery Location Type

At the **Location Type** step of the wizard, select the **VMware vSphere** option.

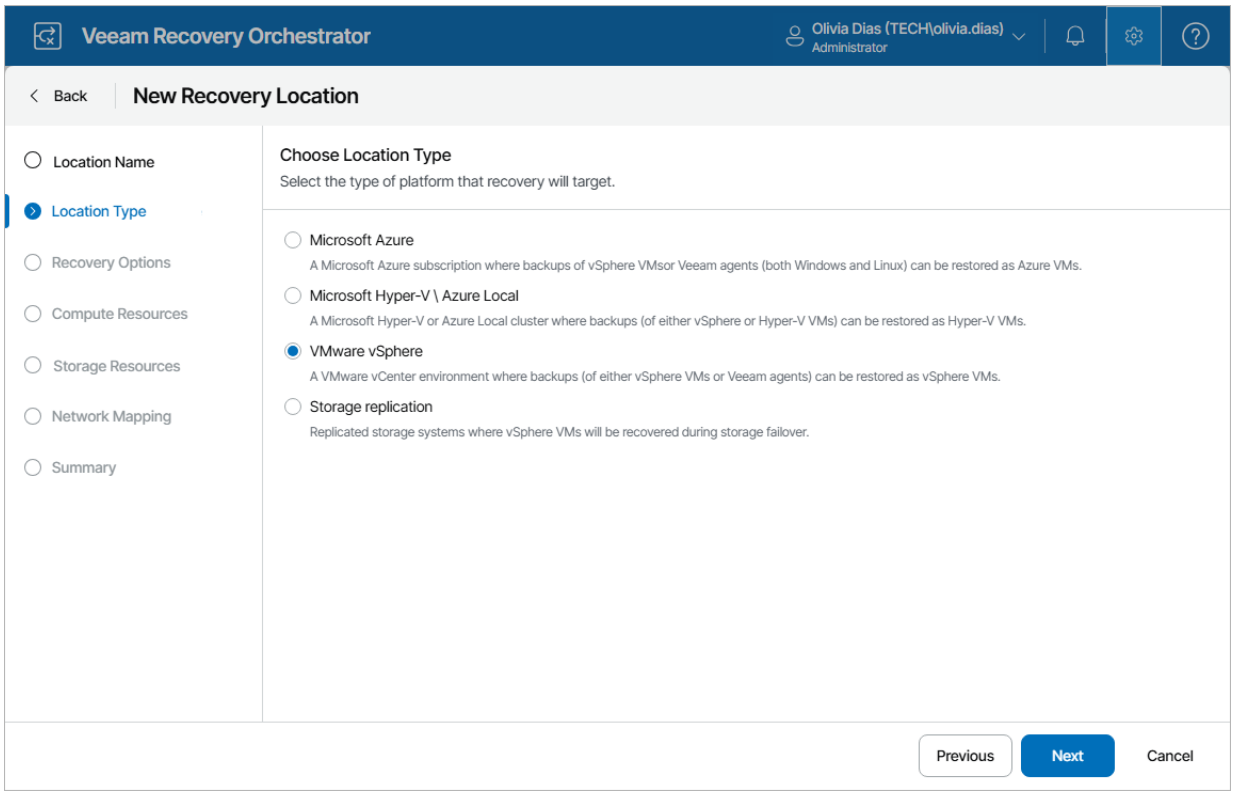

Step 3. Choose Recovery Options

At the **Recovery Options** step of the wizard, choose whether you want Orchestrator to use backup files created by backup jobs, their copies created by backup copy jobs or files that are stored in scale-out backup repositories. Alternatively, you can choose to let the product select the files automatically, based on the most recent restore points. For more information on the way Orchestrator selects backup files and restore points to recover machines when performing restore operations, see the Veeam Recovery Orchestrator User Guide, section How Orchestrator Selects Backup Files.

Additionally, you can configure the following settings:

- Choose whether you want to enable Instant VM Recovery for the location.

  With Instant VM Recovery, all processed machines will be immediately restored in the location by running directly from the backup files. Instant VM Recovery helps improve recovery time objectives (RTO), minimize disruption and downtime of the machines. For more information on the Instant VM Recovery feature, see the Veeam Backup & Replication User Guide, section Instant VM Recovery.

- Choose whether you want Orchestrator to be able to recover VMs to a different location in Veeam Backup & Replication.

  To control data migration in the virtual infrastructure, Veeam Backup & Replication introduces infrastructure locations. A location defines a geographic region where an infrastructure object resides. To learn how to create and assign locations to infrastructure objects in Veeam Backup & Replication, see the Veeam Backup & Replication User Guide, section Locations. To learn how to track geographical locations of production data, their copies and replicas, see the Veeam ONE Reporter User Guide, sections Data Sovereignty Overview and Data Sovereignty Violations.

- Specify the datastore capacity level that must not be breached during the recovery process.

> **NOTE**
>
> Orchestrator currently supports recovering machines protected by Veeam Agent for Microsoft Windows and Veeam Agent for Linux only.

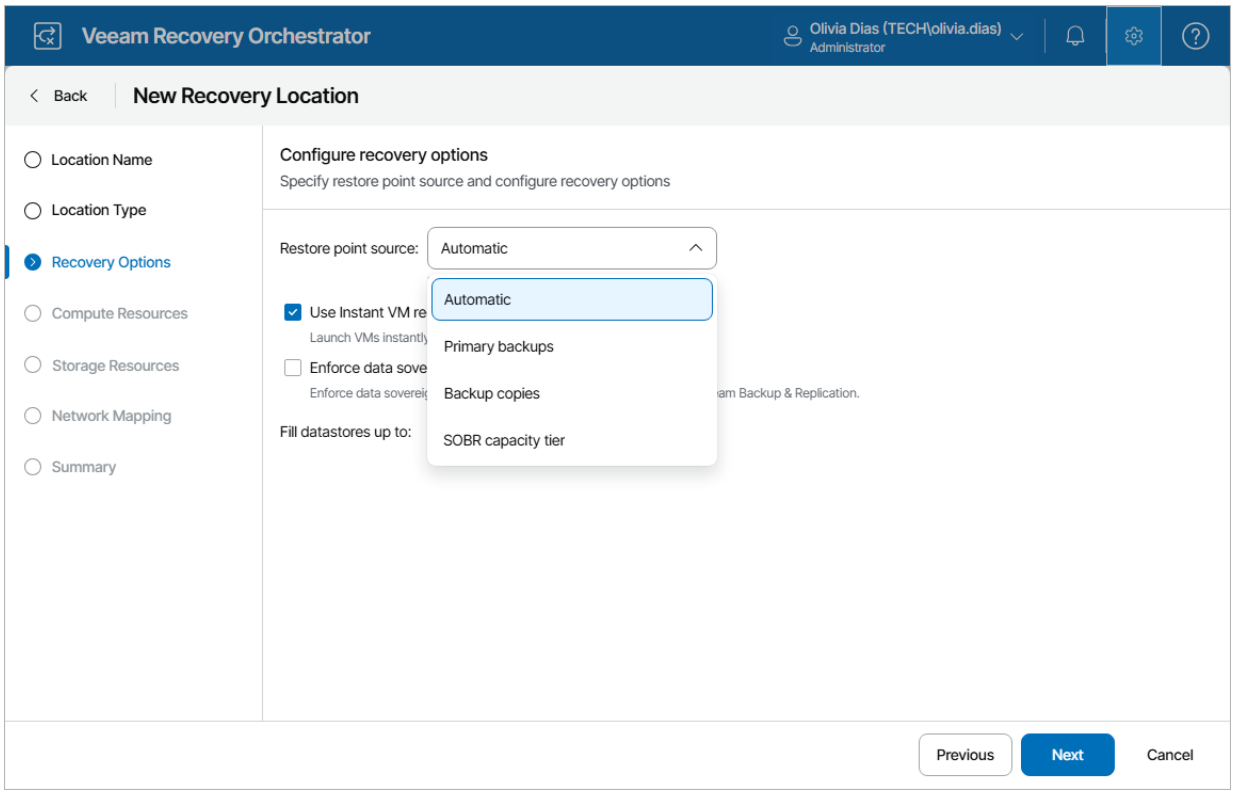
Veeam Recovery Orchestrator
Olivia Dias (TECH\olivia.dias)
Administrator
Back
New Recovery Location
Location Name
Location Type
Recovery Options
Compute Resources
Storage Resources
Network Mapping
Summary
Configure recovery options
Specify restore point source and configure recovery options
Restore point source:
Automatic
Automatic
Primary backups
Backup copies
SOBR capacity tier
Use Instant VM re
Launch VMs instantly
Enforce data sove
Enforce data sovereig
am Backup & Replication.
Fill datastores up to:
Previous
Next
Cancel

Step 4. Specify Compute Resources

At the **Compute Resources** step of the wizard, specify target hosts and clusters where recovered VMs will be registered. To do that, click **Add**, select the required resource groups and click **Save**. To view hosts and clusters included in a resource group, click the group name in the **Category – Tag** list.

For compute resources to be displayed in the **Category – Tag** list, these resources must be categorized into groups in Veeam ONE Client as described in the Veeam Recovery Orchestrator Group Management Guide.

> **IMPORTANT**
>
> If you add a host to a recovery location and then move the host to another datacenter, or if you move a host from one vCenter Server to another, the host will be assigned a new vCenter MoRef ID, Orchestrator will consider the host to be a new infrastructure object, and the configuration of the recovery location will become invalid. As a result, Orchestrator will not be able to use this location for restore.

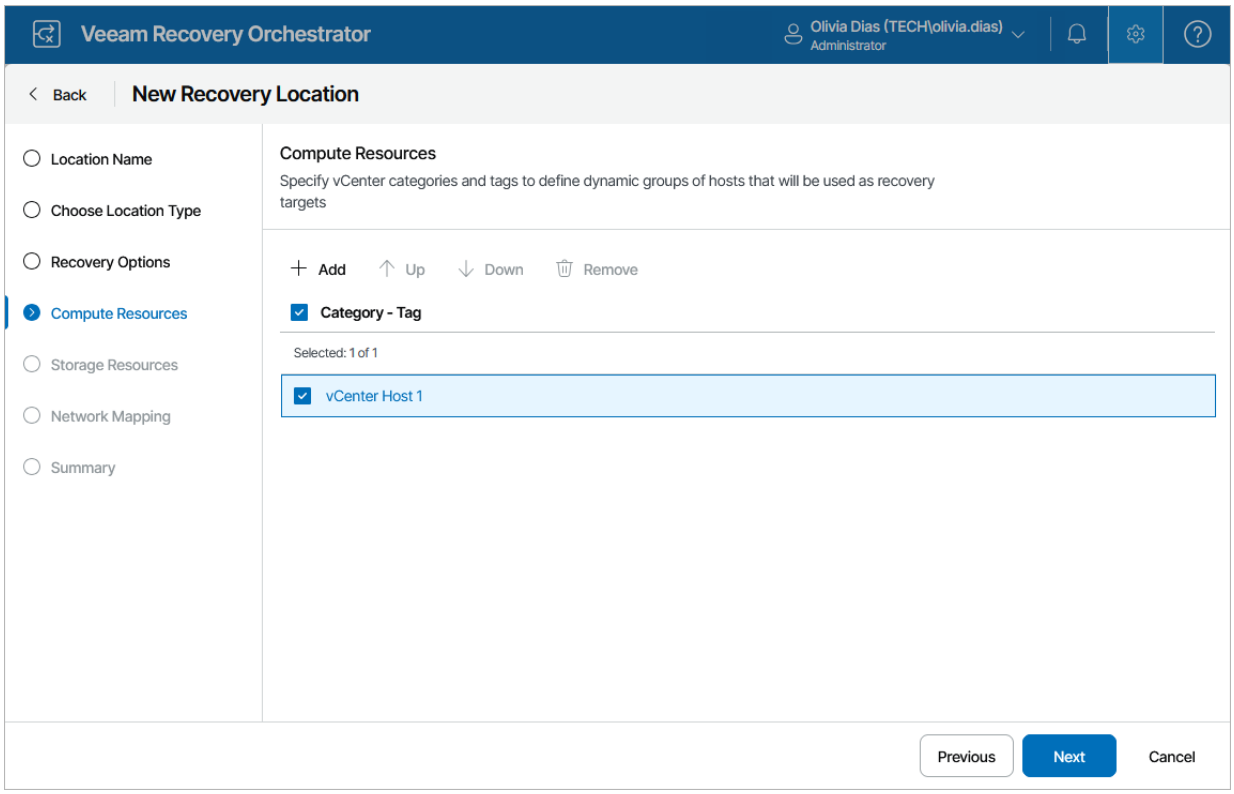

Step 5. Specify Storage Resources

At the **Storage Resources** step of the wizard, specify destination datastores and datastore clusters where recovered VMs will be stored. To do that, click **Add**, select the required resource groups and click **Save**. To view resources included in a resource group, click the group name in the **Category – Tag** list.

For storage resources to be displayed in the **Category – Tag** list, these resources must be categorized into groups in Veeam ONE Client as described in the Veeam Recovery Orchestrator Group Management Guide.

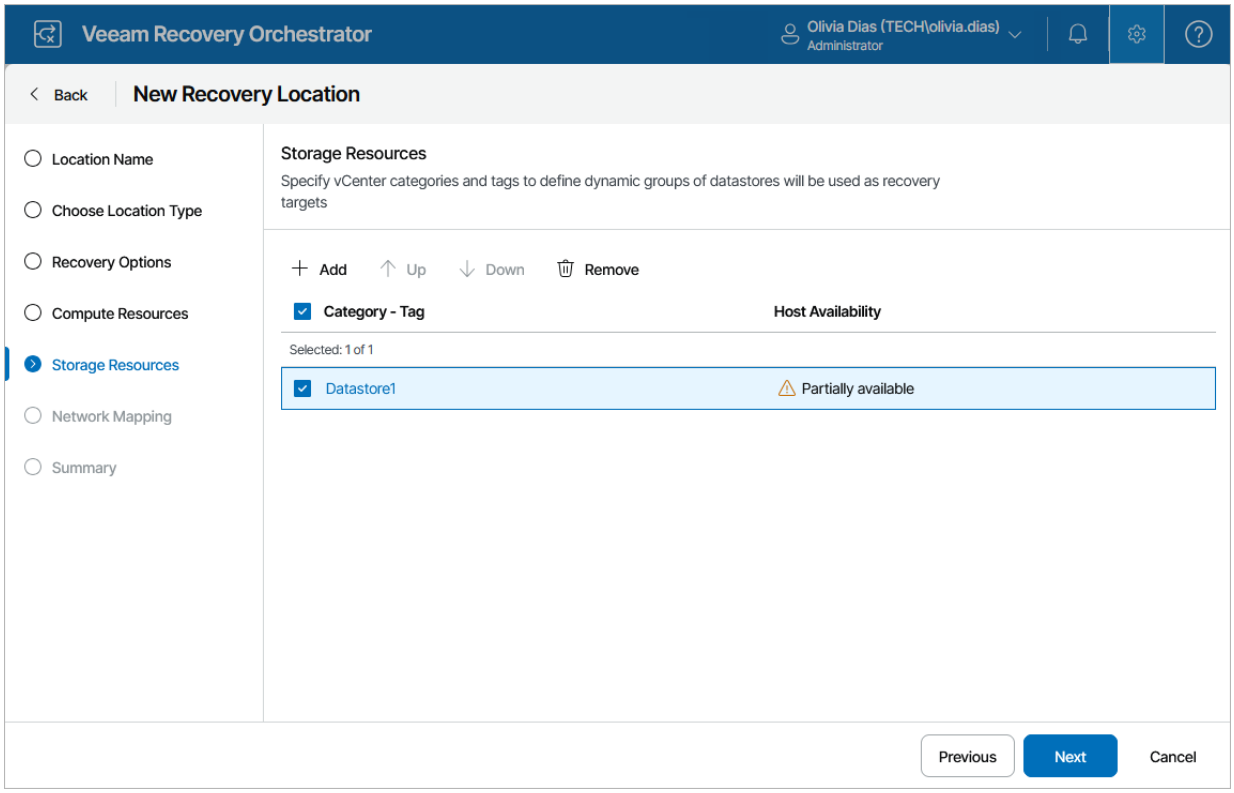

Step 6. Configure Network Mapping

[This step applies only if you want to enable the functionality of network mapping]

When you recover a VM from a vSphere backup, the recovered VM is connected to the same vSphere networks as the source VM; if the same networks are not available in the recovery location, you can create a network mapping table for the location so that the recovered VM is connected to the correct network. However, when you recover a machine from a Veeam agent backup, there are no vSphere networks that can be used — only the IP address of the source agent is known. Therefore, to recover a Veeam agent, you must create at least one network mapping rule that maps an IP address range to a vSphere network so that the recovered VM is connected to the correct network.

To configure network mapping, click **Add Mapping** at the **Network Mapping** step of the wizard and choose whether you want to recover VMs from vSphere or Veeam agent backups. The **Add Network Mapping Rule** window will open.

1. Depending on whether you plan to recover machines from vSphere or Veeam agent backups, do the following in the **Source network** section:
   - To recover VMs from vSphere backups, select a vCenter Server that manages source VMs, a network to which the source VMs are connected, and a datacenter or a cluster where the source VMs reside.

     For a vCenter Server to be displayed in the **vCenter Server** list, it must be connected to Orchestrator as described in section Connecting VMware vSphere Servers.
   - To recover machines from Veeam agent backups, specify a range of IP addresses that contains the IP addresses of the source agent machines. Alternatively, create a separate network mapping rule to map each individual IP address. Note that you must specify at least one network mapping rule.

> **NOTE**
>
> Orchestrator supports IP addresses in the IPv4 format only. If a machine that you want to recover has an IPv6 address only, you must create the *0.0.0.0/0* mapping rule. Otherwise, Orchestrator may halt the recovery process.

2. In the **Target vSphere network** section, select a vCenter Server that will manage recovered VMs, a network to which the recovered VMs will be connected, and a datacenter or a cluster where the target VMs will reside.

   For a vCenter Server to be displayed in the **vCenter Server** list, it must be connected to Orchestrator as described in section Connecting VMware vSphere Servers.

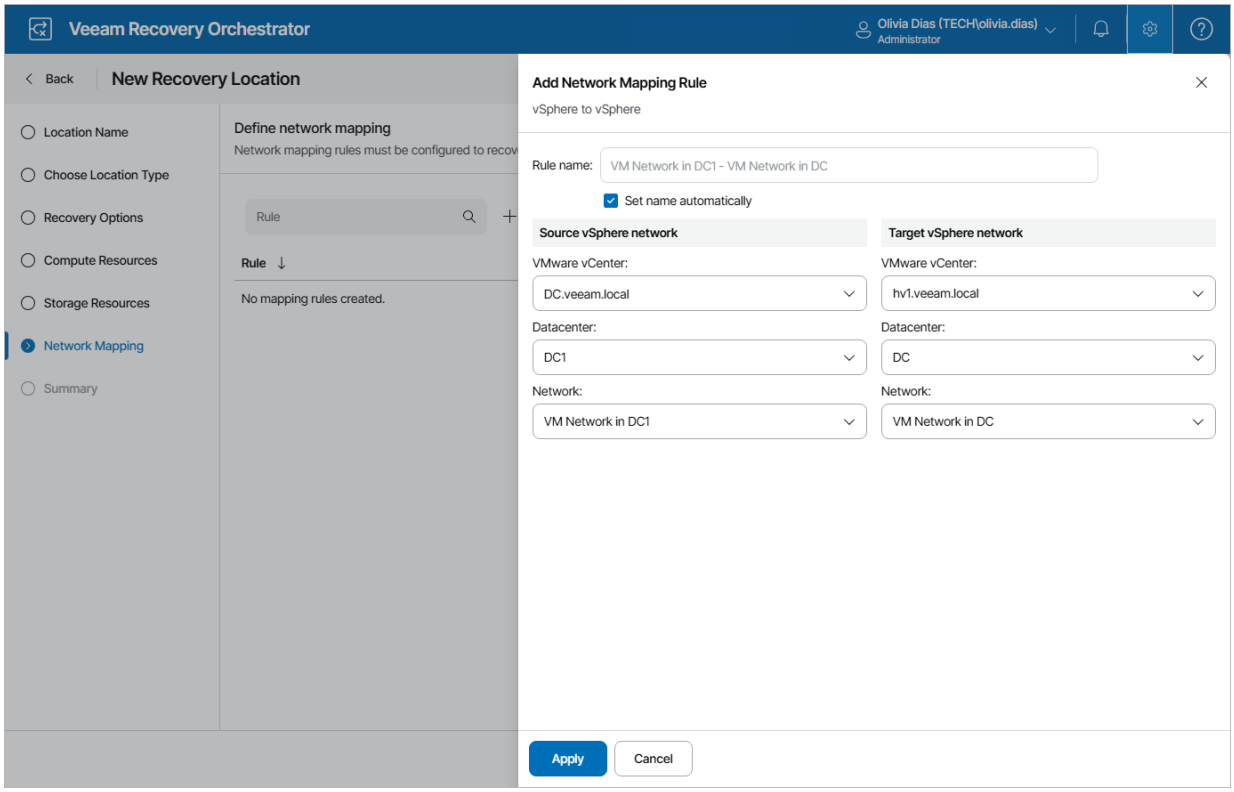

Configuring Re-IP Rules

[This step applies only if you want to enable the functionality of automatic IP address transformation for recovery of Microsoft Windows servers]

If the IP addressing scheme in the source location differs from the target location scheme, you can create re-IP rules for the recovery location, and Orchestrator will automatically reconfigure IP addresses of the recovered VMs. When recovering from a vSphere or agent backup, or failing back to a new location, Orchestrator checks if any of the specified re-IP rules will apply to the recovered VM: if a rule applies, Orchestrator will change the IP address configuration of the recovered VM using the Microsoft Windows registry.

> **IMPORTANT**
>
> To allow Orchestrator to reconfigure IP addresses of a recovered VM, the machine must have VMware Tools installed. This also applies to physical servers protected by Veeam agents if you plan to recover them to the VMware vSphere environment. The readiness check for any plan containing Veeam agents will confirm the presence of VMware Tools.

To configure a re-IP rule, click **Add Mapping** > **Re-IP Rule**. Then, do the following in the **Add Network Mapping Rule** window:

1. In the **Description** field, specify a brief outline for the rule or leave any related comments.
2. In the **Source VM** section, describe an IP numbering scheme adopted in the source location.
3. In the **Target VM** section, describe an IP numbering scheme adopted in the target location.

4. In the **Target network options** section, specify a default gateway that will be used for recovered VMs. If necessary, define the DNS server addresses.
5. Click **Apply**.

> **TIP**
>
> You can use the asterisk character (*) to specify a range of IP addresses.

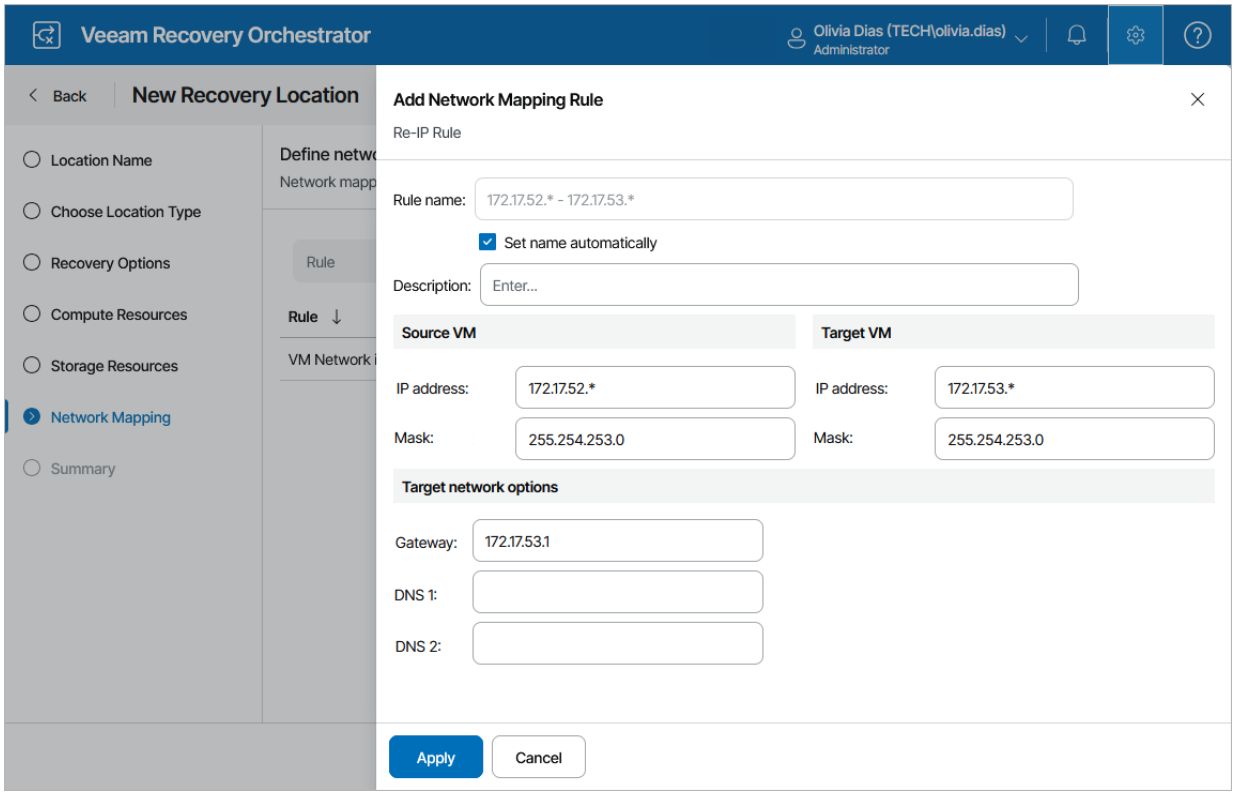

Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.

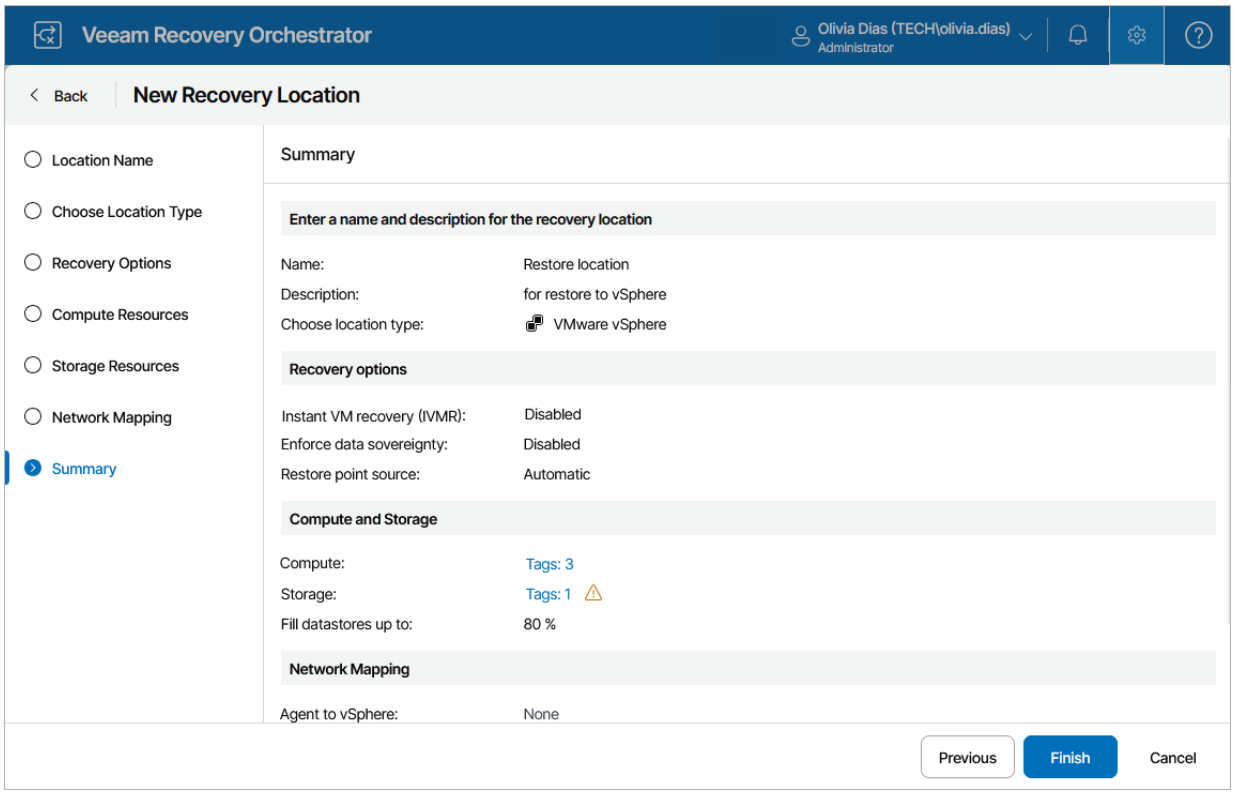

Adding Storage Recovery Locations

To add a storage recovery location:

1. Switch to the **Administration** page.
2. Navigate to **Recovery Locations**.
3. Click **Add**.
4. Complete the **New Recovery Location** wizard:
   a. Specify a recovery location name and description.
   b. Choose a recovery location type.
   c. Choose a storage vendor.
   d. Specify storage and compute resources.
   e. Configure network mapping and re-IP rules.
   f. Finish working with the wizard.

Step 1. Specify Recovery Location Name and Description

At the **Location Name** step of the wizard, use the **Name** and **Description** fields to enter a name for the new location and to provide a description for future reference. The maximum length of the location name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

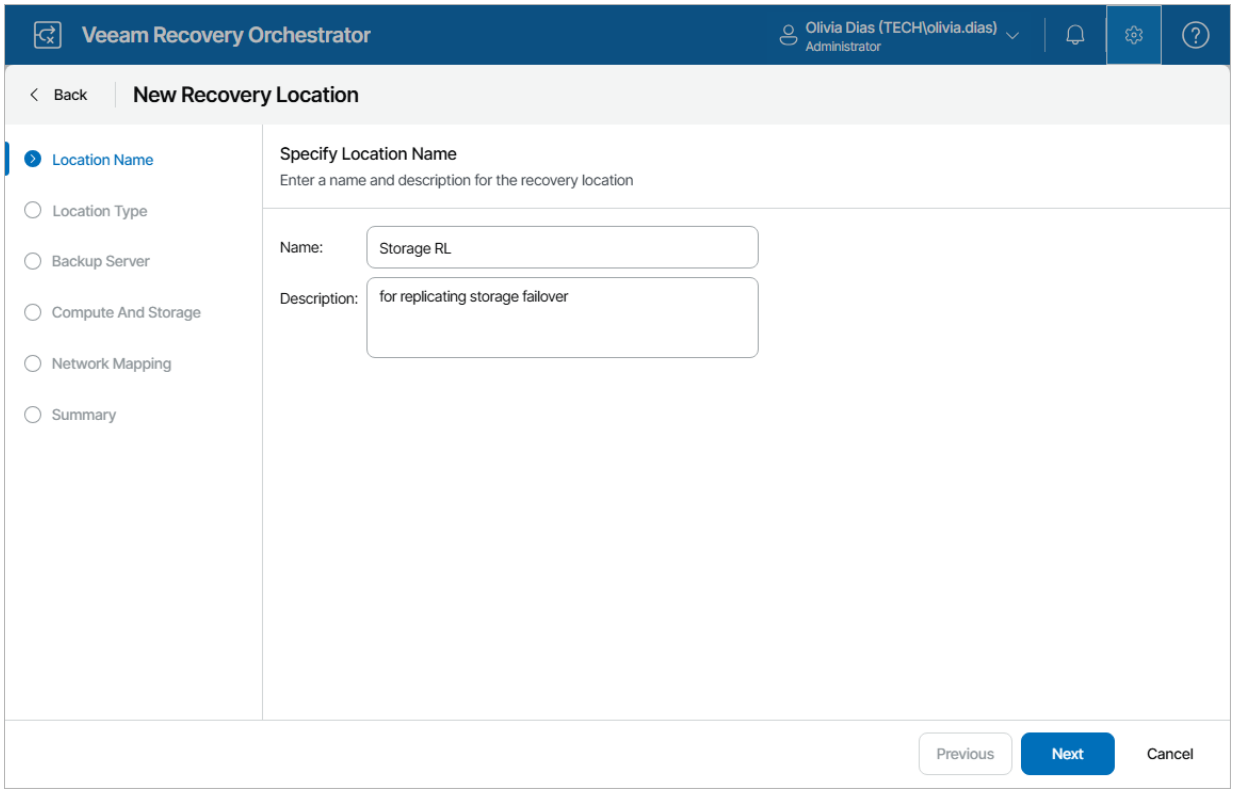

Step 2. Choose Recovery Location Type

At the **Location Type** step of the wizard, select the **Storage replication** option.

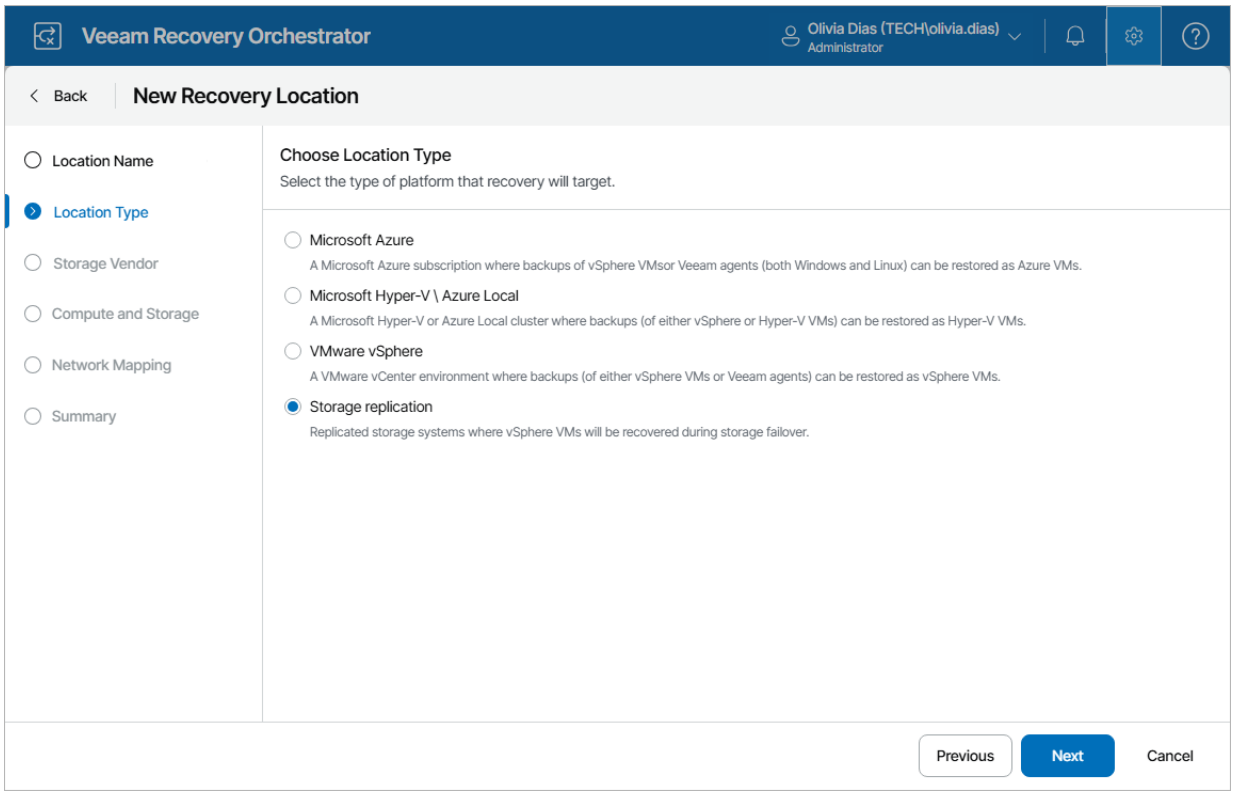

Step 3. Choose Storage Vendor

At the **Storage Vendor** step, choose whether NetApp or HPE storage systems will be used to recover VMs.

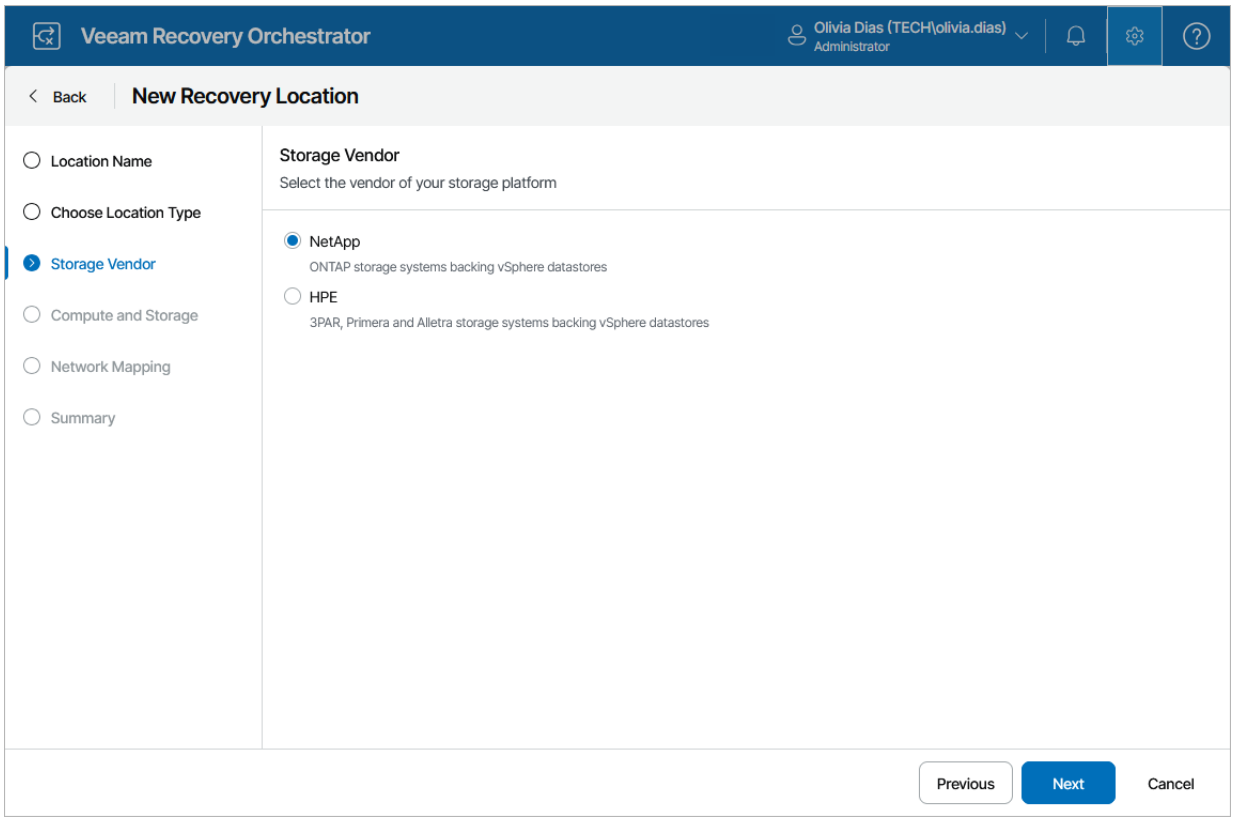

Step 4. Specify Compute and Storage Resources

At the **Compute and Storage** step of the wizard, configure the following settings:

1. Click the link in the **Storage** field and specify target storage systems to be used to recover VMs. To do that, select the required storage systems in the list of available systems and click **Save**.

   To view datastores included in a storage system, click the system name in the **Storage System** list.

2. From the **vCenter Server** drop-down list, select a vCenter Server that will manage recovered VMs.

3. From the **Datacenter** drop-down list, select a datacenter to be used to recover VMs. The recovered VMs will be distributed across resources of this datacenter.

4. Click **Add** to specify target hosts and clusters to which recovered volumes will be mounted. To do that, select resource groups in the list of available groups and click **Save**.

   For compute resources to be displayed in the **Category** list, these resources must be categorized into groups in Veeam ONE Client as described in the Veeam Recovery Orchestrator Group Management Guide and must belong to the datacenter selected at step 3. If a resource group belongs to multiple datacenters at the same time, it will not be displayed in the list.

   To view hosts and clusters included in a resource group, click the group name in the **Category** list.

> **IMPORTANT**
>
> If you add a host to a recovery location and then move the host to another datacenter, or if you move a host from one vCenter Server to another, the host will be assigned a new vCenter MoRef ID, Orchestrator will consider the host to be a new infrastructure object, and the configuration of the recovery location will become invalid. As a result, Orchestrator will not be able to use this location for recovery.

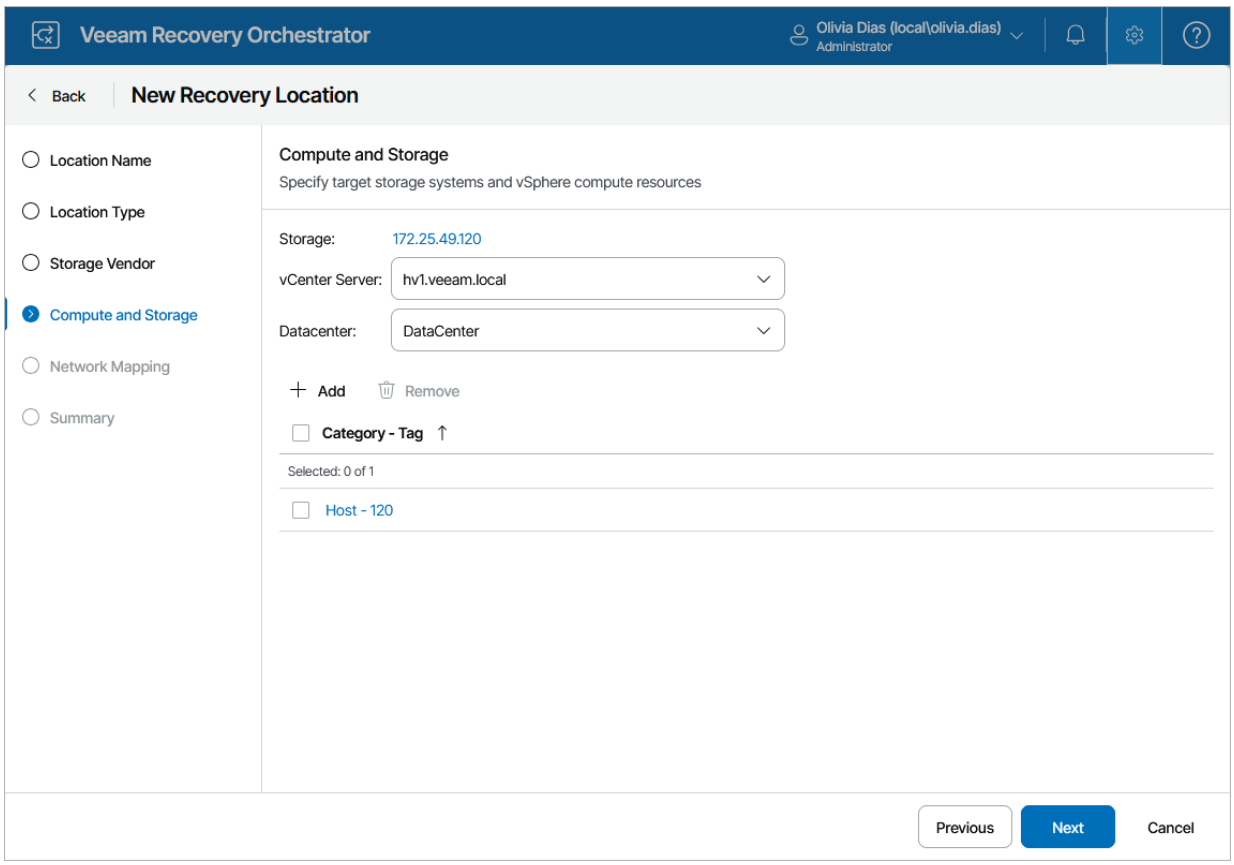

Step 5. Configure Network Mapping

[This step applies only if you want to enable the functionality of network mapping]

By default, a recovered VM is connected to the same networks as the source VM. If the network configuration in the recovery location does not match the production network configuration, you can create a network mapping table for the location so that the recovered VM is connected to the correct network.

At the **Network Mapping** step of the wizard, click **Add Mapping** > **vSphere to vSphere** to configure network mapping. Then, do the following in the **Add Network Mapping Rule** window:

1. In the **Source vSphere network** section, select a vCenter Server that manages source VMs, a network to which the source VMs are connected, and a datacenter where the source VMs reside.

   For a vCenter Server to be displayed in the **VMware vCenter** list, it must be connected to Orchestrator as described in section Connecting VMware vSphere Servers.

2. In the **Target vSphere network** section, select a vCenter Server that will manage recovered VMs, a network to which the recovered VMs will be connected, and a datacenter or a cluster where the target VMs will reside.

   Since you have already specified the target datacenter to be used, the wizard only allows you to change the target network.

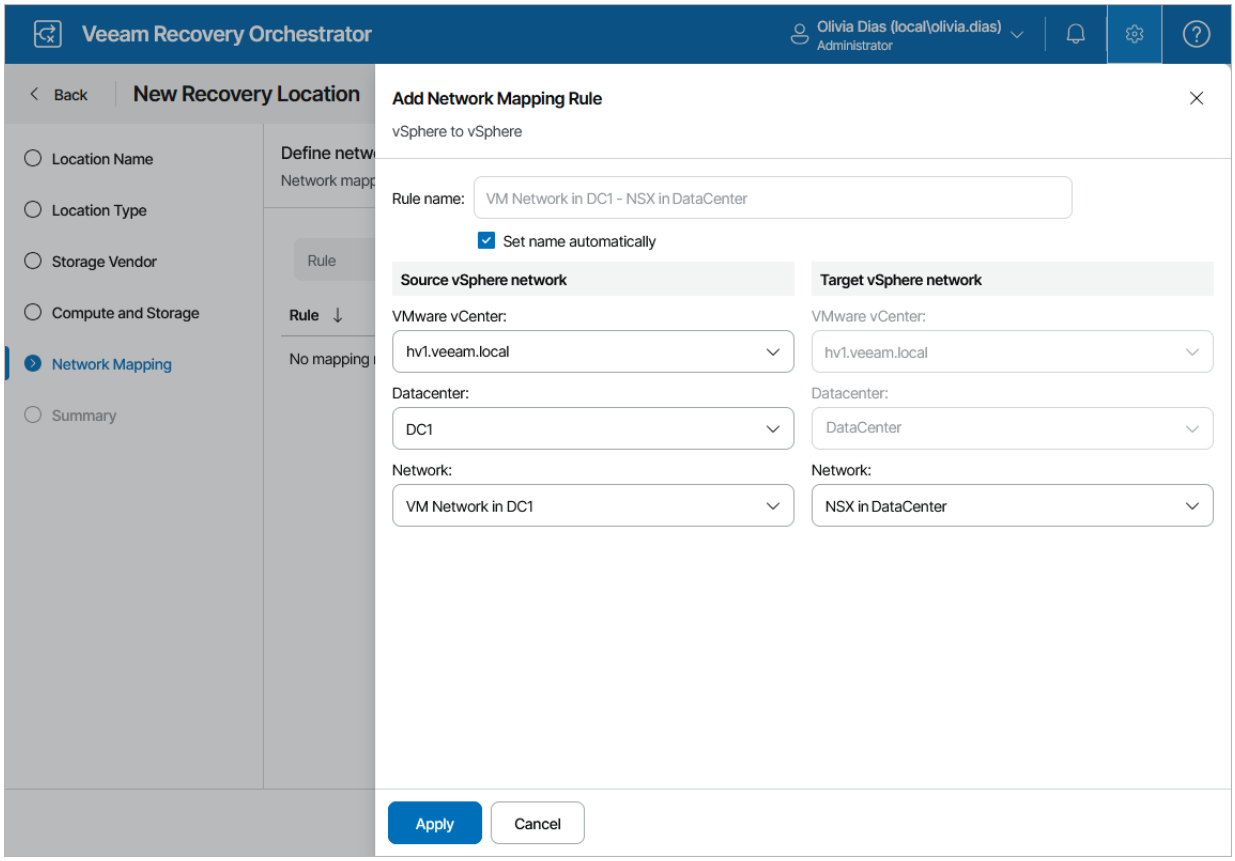

Configuring Re-IP Rules

[This step applies only if you want to enable the functionality of automatic IP address transformation for recovery of Microsoft Windows servers]

If the network configuration in the source location does not match the production network configuration, you can create re-IP rules for the recovery location, and Orchestrator will automatically reconfigure IP addresses of the recovered VMs. During storage failover, Orchestrator checks if any of the re-IP rules will apply to the recovered VM: if a rule applies, Orchestrator will change the IP address configuration of the recovered VM using the Microsoft Windows registry.

> **IMPORTANT**
>
> To allow Orchestrator to reconfigure IP addresses of a recovered VM, the VM must have VMware Tools installed.

To configure a re-IP rule, click **Add Mapping** > **Re-IP Rule**. Then, do the following in the **Add Network Mapping Rule** window:

1. In the **Description** field, specify a brief outline for the rule or leave any related comments.
2. In the **Source VM** section, describe an IP numbering scheme adopted in the source location.
3. In the **Target VM** section, describe an IP numbering scheme adopted in the target location.
4. In the **Target network options** section, specify a default gateway that will be used for recovered VMs.

   If necessary, define the DNS server addresses.
5. Click **Apply**.

> **TIP**
>
> You can use the asterisk character (*) to specify a range of IP addresses.

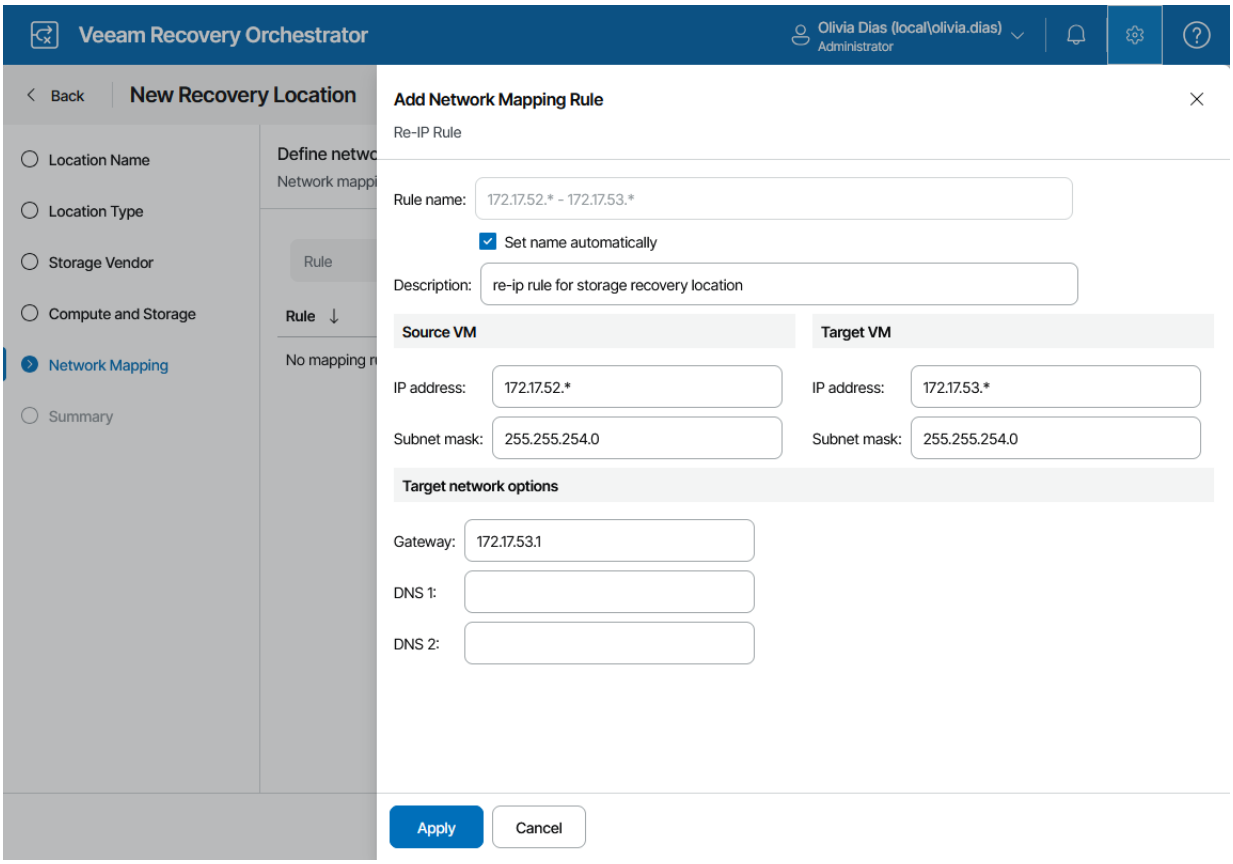

Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.

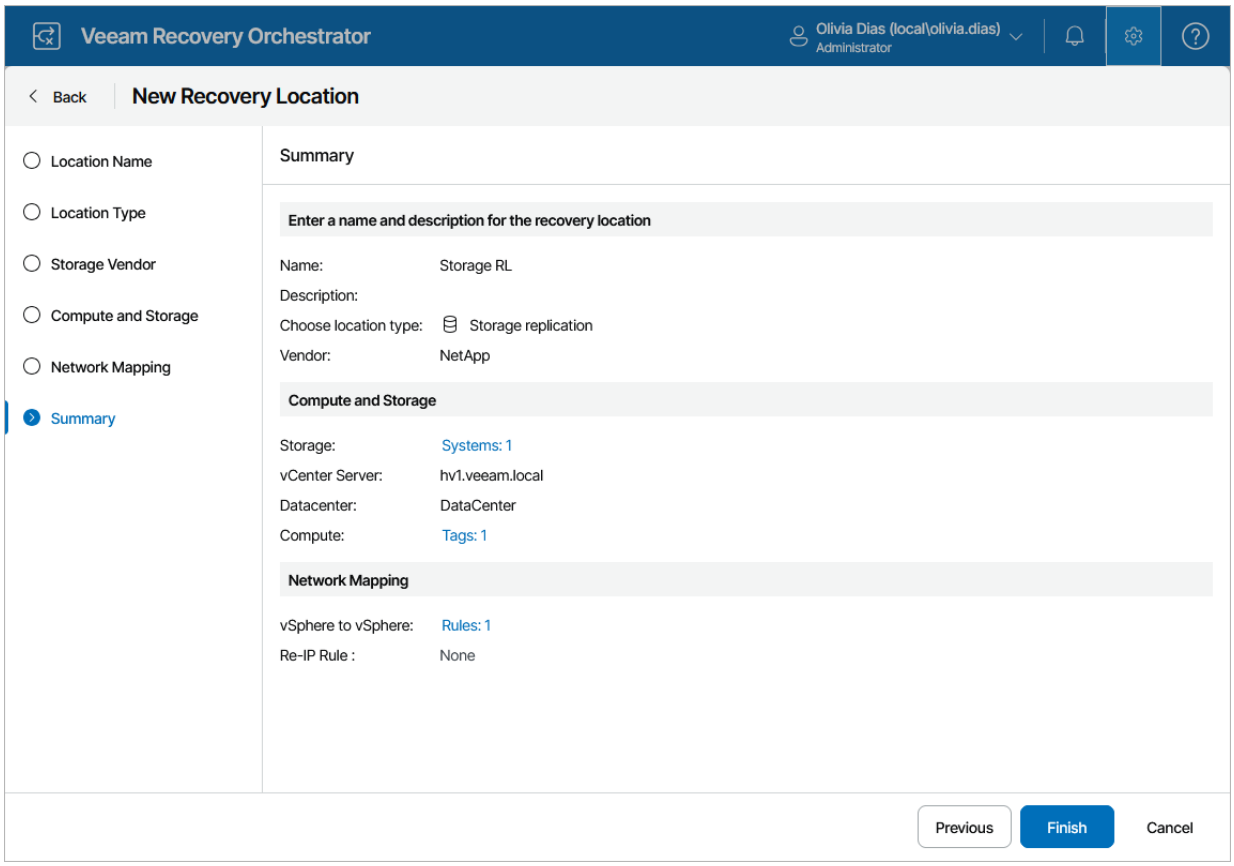

Adding Microsoft Azure Recovery Locations

To add a Microsoft Azure recovery location:

1. Switch to the **Administration** page.
2. Navigate to **Recovery Locations**.
3. Click **Add**.
4. Complete the **New Recovery Location** wizard:
   a. Specify a recovery location name and description.
   b. Choose a recovery location type.
   c. Choose backup servers, specify repositories and configure proxies.
   d. Choose a Microsoft Azure compute account, a region, a resource group and specify a cloud VM configuration.
   e. Configure network mapping.
   f. Finish working with the wizard.

Step 1. Specify Recovery Location Name and Description

At the **Location Name** step of the wizard, use the **Name** and **Description** fields to enter a name for the new location and to provide a description for future reference. The maximum length of the location name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

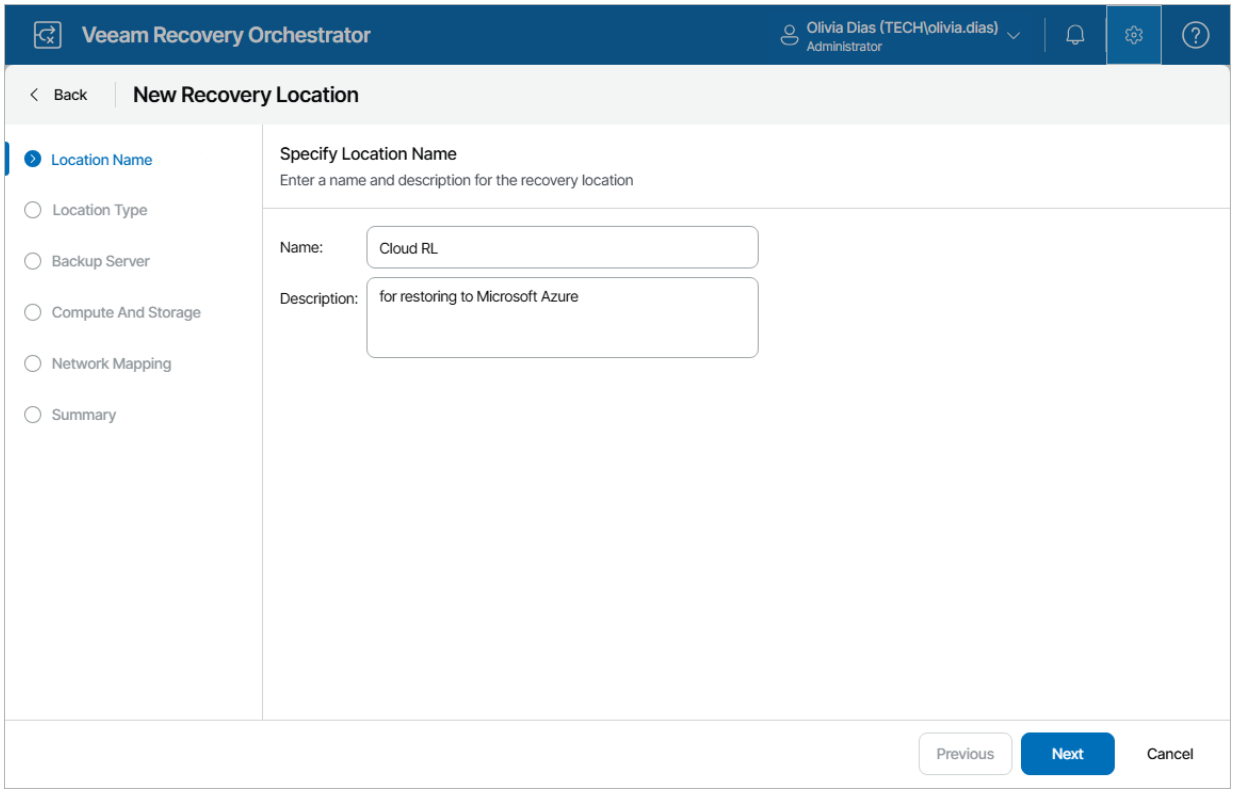

Step 2. Choose Recovery Location Type

At the **Location Type** step of the wizard, select the **Microsoft Azure** option.

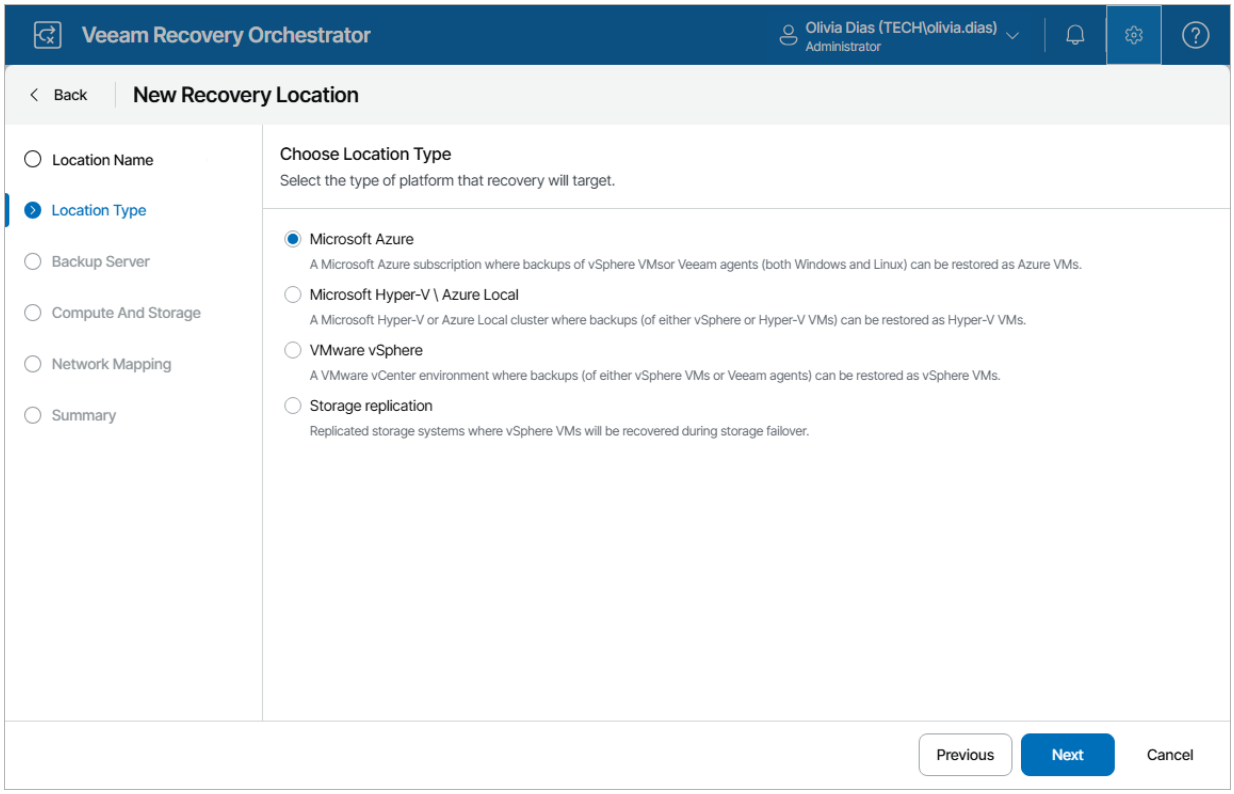

Step 3. Choose Recovery Options

At the **Backup Server** step of the wizard, select a Veeam Backup & Replication server that will manage the process of recovering machines to Microsoft Azure. Note that one recovery location can be associated with one Veeam Backup & Replication server only.

For a Veeam Backup & Replication server to be displayed in the list of available servers, it must be added to Orchestrator as described in section Connecting Veeam Backup & Replication Servers.

> **IMPORTANT**
>
> - To restore workloads to Microsoft Azure, you must add a Microsoft Azure compute account to Veeam Backup & Replication. When you add an Azure compute account, Veeam Backup & Replication imports information about subscriptions and resources associated with this account.
>
> If you do not add a compute account to the configuration of a Veeam Backup & Replication server, you will not be able to choose the server when creating a Microsoft Azure recovery location. For more information on adding Microsoft Azure compute accounts, see the Veeam Backup & Replication User Guide, section Microsoft Azure Compute Accounts.
> - The added account must have the permissions required to access Microsoft Azure resources. For more information, see the Veeam Backup & Replication User Guide, section Creating Custom Role for Azure Account.

Additionally, you can set the **Advanced settings** toggle to *On* and configure the following settings.

Configuring Azure Proxies

Veeam Backup & Replication servers can use Azure restore proxy appliances to speed up the recovery process. For more information on Azure proxy appliances, see the Veeam Backup & Replication User Guide, section Managing Azure Restore Proxy Appliances.

In the **Veeam proxy** field, specify the proxy appliances to be used. To do that, select a proxy appliance in the list of available proxy appliances and click **Add**. If want to select a cloud repository as a target location for storing restore points, choose a cloud proxy appliance to speed up the recovery process.

> **IMPORTANT**
>
> To be able to recover Linux workloads, you must configure a Linux helper appliance when adding a Microsoft Azure compute account to the backup infrastructure.

Configuring Recovery Options

From the **Restore point source** drop-down list, choose whether you want Orchestrator to use backup files created by backup jobs, their copies created by backup copy jobs or files that are stored in scale-out backup repositories. Alternatively, you can choose to let the product select the files automatically, based on the most recent restore points. For more information on the way Orchestrator selects backup files and restore points to recover machines when performing restore operations, see the Veeam Recovery Orchestrator User Guide, section How Orchestrator Selects Backup Files.

Specifying Backup Repositories

Set the **Specify repositories** toggle to *On* to choose backup repositories where the restore points of machines that you plan to recover are stored. To do that, click **Add**, select a repository in the list of available repositories and click **Add**.

For a backup repository to be displayed in the list of the available repositories, it must be added to the backup infrastructure as described in section Adding Backup Repositories. You can use directly attached storage (Windows, Linux), object storage (Veeam Data Cloud Vault, Microsoft Azure, Wasabi) and scale-out backup repositories. Tape and deduplicating storage appliances are not supported.

> **IMPORTANT**
>
> For Veeam Backup & Replication to be able to access the selected repositories using your cloud accounts, the credentials of these accounts must be kept up to date. This means that if you update credentials of a cloud account that is used to connect to a cloud service, you must also update these credentials in the Veeam Backup & Replication console as described in the Veeam Backup & Replication User Guide, section Cloud Credentials Manager.

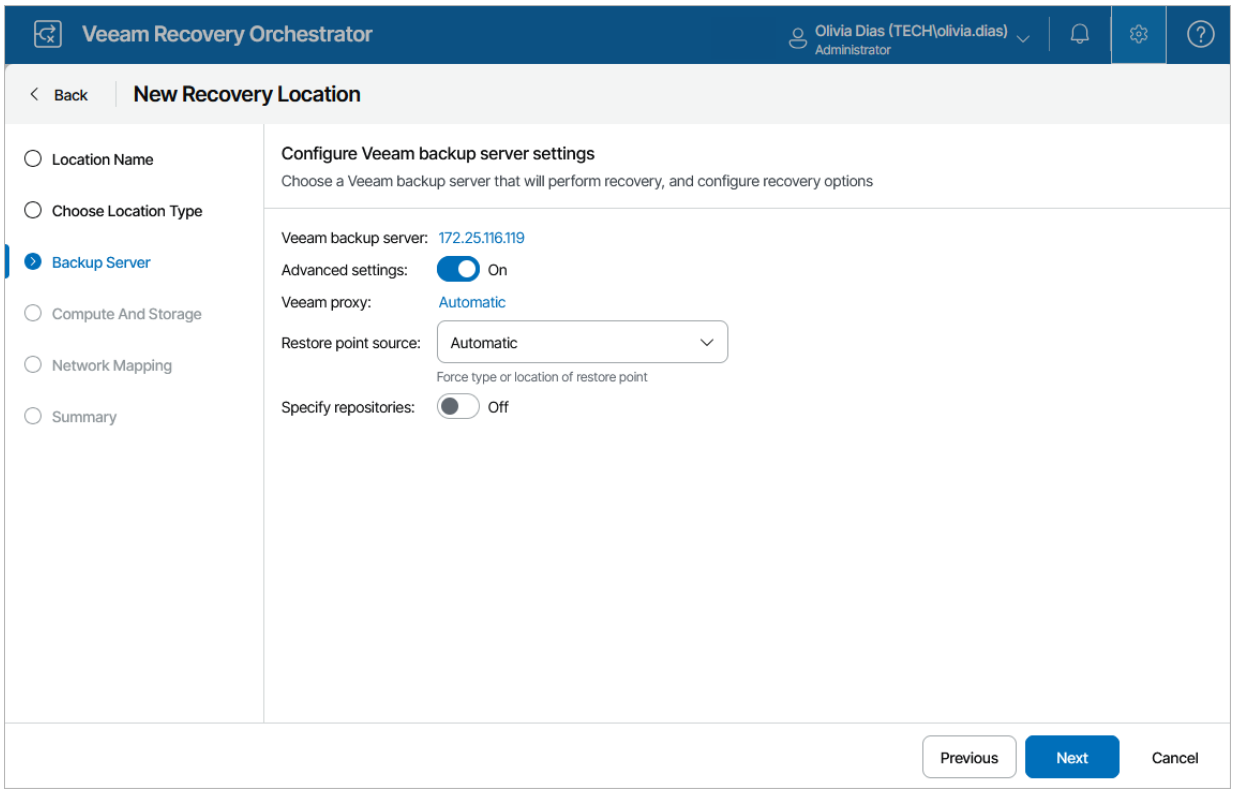

Step 4. Configure Compute and Storage Settings

At the **Compute and Storage** step of the wizard, configure the following settings:

1. Click the link in the **Compute account** field and specify a Microsoft Azure compute account added to Veeam Backup & Replication that you want to use to recover machines.
2. From the **Region** drop-down list, select a Microsoft Azure region in which the recovered VMs will reside.

   For a region to be displayed in the list of available regions, it must belong to the subscription specified at step 2. Note that Early Updates Access Program (EUAP) regions are not supported.

> **NOTE**
>
> For Orchestrator to deploy the recovered VMs in the selected region, you must have sufficient resource quota allocated to your subscription. To learn how to check your quotas, see Microsoft Azure documentation.

3. From the **Resource group** drop-down list, select a resource group to which the recovered VMs will belong.

   For a resource group to be displayed in the Resource Group list, it must be created for the region specified at step 3 in the Microsoft Azure portal, as described in Microsoft Docs.

4. In the **Azure VM Configurations** section, specify a VM configuration (that is, a combination of a VM series and disk type) that Orchestrator will use to create new VMs in Microsoft Azure.

   To help you choose the VM series, the table in the **Choose VM Series** window will provide information on the maximums for the number of vCPU cores, system RAM and attached disks for each available VM size. For the full description of Microsoft Azure VM sizes, see Microsoft Docs.

The specified VM series will be used as the basis for machines recovered in Microsoft Azure as new VMs. The created VMs will be customized to best match the CPU and memory configuration of the source machines. If you want different machines to be recovered using different settings, you can add up to 2 more VM configurations. If you create another VM configuration, you must also modify the parameters of the **Restore to Azure** step to use it, as described in section Configuring Plan Steps.

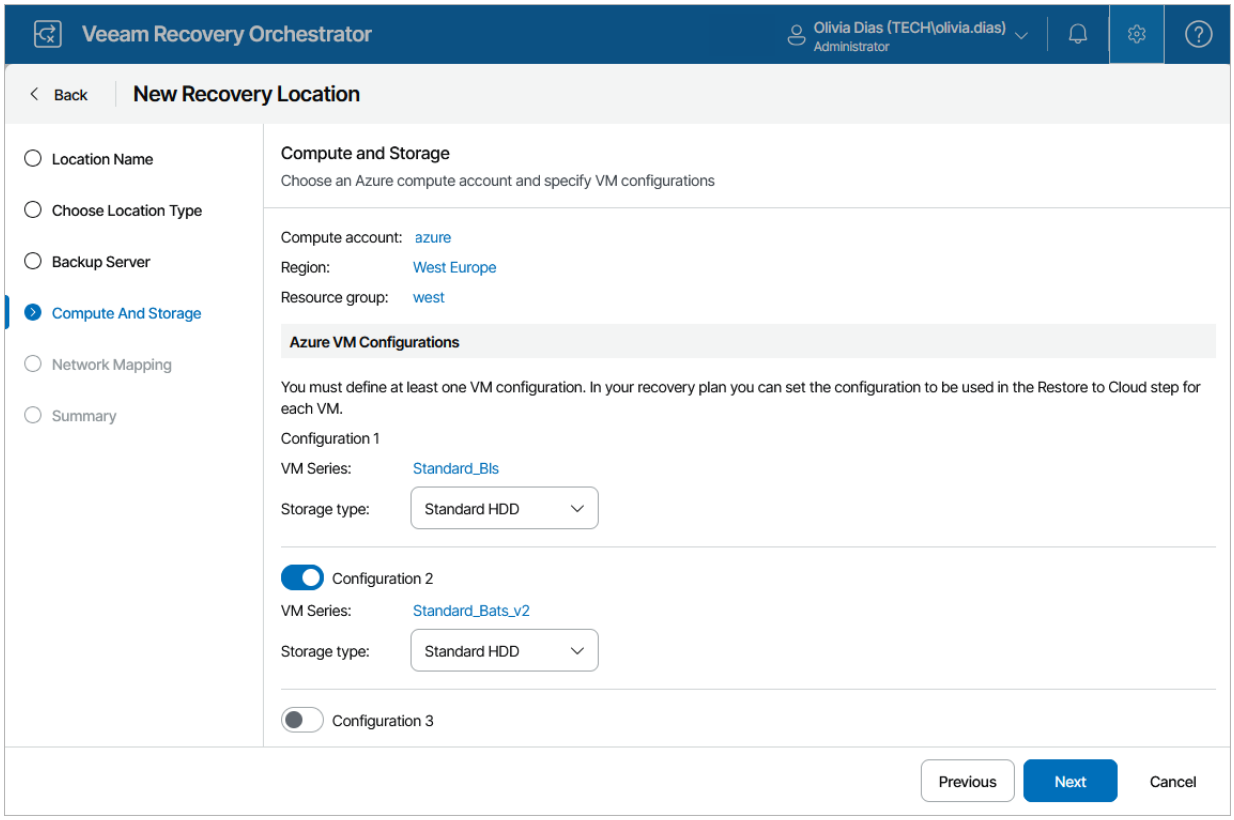

Step 5. Configure Network Mapping

When you recover a machine from a Veeam agent backup or a VM from a vSphere backup to a cloud environment, Orchestrator is not able to connect the recovered VM to the same networks as the source machine — that is why you must create at least one network mapping rule for the location so that the recovered VM is connected to the correct network.

To configure network mapping, click **Add** at the **Network Mapping** step of the wizard. Then, do the following in the **Add Network Mapping Rule** window:

1. Depending on whether you plan to recover machines from vSphere or agent backups, do the following in the **Source network** section:
   - To recover machines from Veeam agent backups, specify a range of IP addresses that contains the IP addresses of the source agent machines. Alternatively, create a separate network mapping rule to map each individual IP address.
   - To recover VMs from vSphere backups, select a vCenter Server that manages source VMs, a network to which the source VMs are connected, and a datacenter or a cluster where the source VMs reside.

     For a vCenter Server to be displayed in the **vCenter Server** list, it must be connected to Orchestrator as described in section Connecting VMware vSphere Servers.

> **NOTE**
>
> Orchestrator supports IP addresses in the IPv4 format only. If a machine that you want to recover has a single IPv6 address, you must create a mapping rule *0.0.0.0/0* (this will map all networks). Otherwise, Orchestrator may halt the recovery process.

2. In the **Target network** section, select a virtual network and a subnet to which you want to connect the recovered VMs. For a virtual network to be displayed in the **Cloud network** list, it must be created in the Microsoft Azure portal for the region selected at step 4, as described in Microsoft Docs. For a subnet to be displayed in the **Subnets** list, it must be created in the Microsoft Azure portal for the specified virtual network, as described in Microsoft Docs.

   You can also specify a security group (virtual firewall) that will be associated with the recovered VMs. For a network security group to be displayed in the **Security Groups** list, it must be created and associated to the necessary subnet in the Microsoft Azure portal as described in Microsoft Docs.

> **IMPORTANT**
>
> - If the IP address ranges for different rules overlap, Orchestrator will use the mapping in the rule with the narrowest range.
> - Even if a source machine had multiple network adapters, the recovered VM will have only one network adapter (vNic).

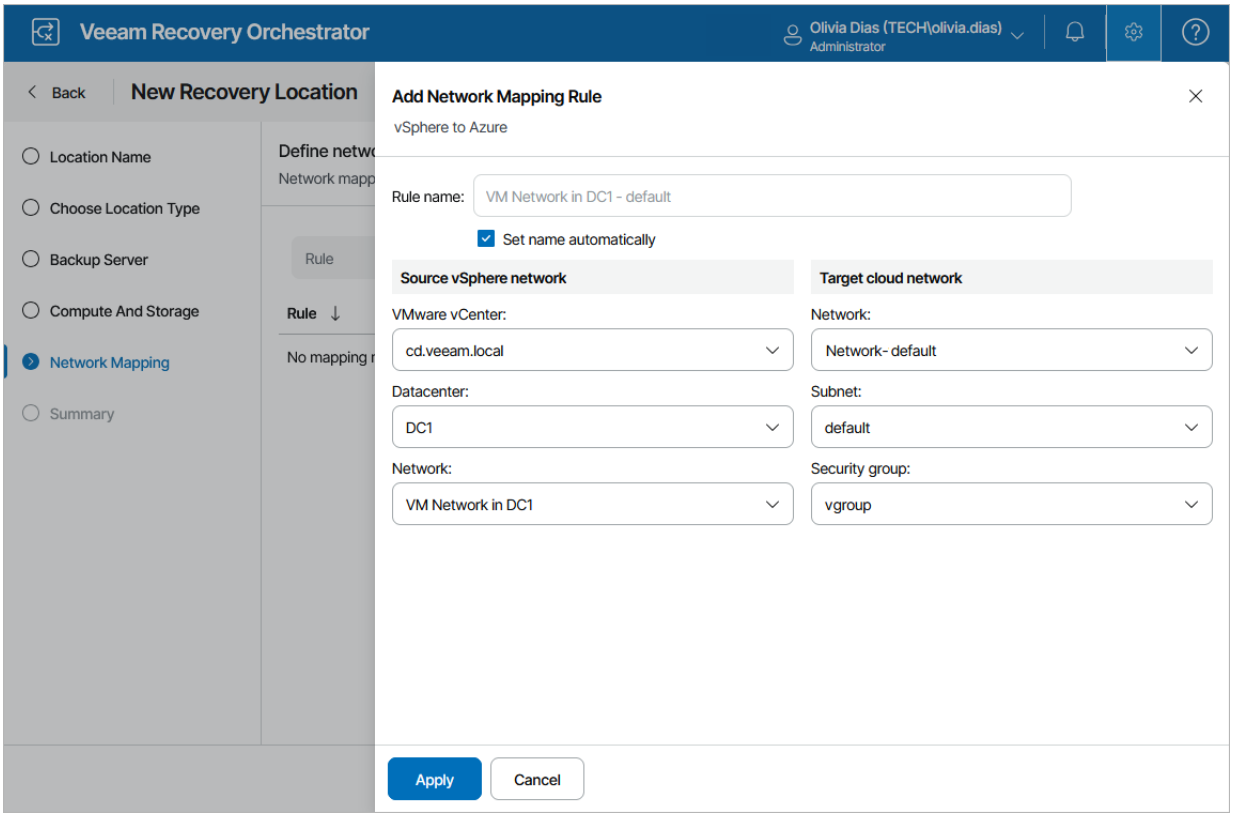

Configuring Quarantine Network

You can scan machine disks for possible malware before restoring them to the production environment. During malware scan, Orchestrator iterates through the number of restore points specified while running the plan one by one to detect a restore point with no malware. By default, if all restore points of a machine are infected, Orchestrator halts the restore. However, you can instruct Orchestrator to restore the infected machine to a quarantine network.

At the **Network Mapping** step of the wizard, click **Add Mapping** > **Quarantine Network** and specify a network and a subnet to which you want to connect the infected machines. For a virtual network to be displayed in the **Target cloud network** list, it must be created in the Microsoft Azure portal for the region selected at step 4, as described in Microsoft Docs. For a subnet to be displayed in the **Subnet** list, it must be created in the Microsoft Azure portal for the specified virtual network, as described in Microsoft Docs.

You can also specify a security group (virtual firewall) that will be associated with the recovered VMs. For a network security group to be displayed in the **Security group** list, it must be created and associated to the necessary subnet in the Microsoft Azure portal as described in Microsoft Docs.

> **NOTE**
>
> Orchestrator allows you to perform malware scan only for those machines whose restore points are stored in on-premises or scale-out repositories. If a restore point is stored in an object storage repository, Orchestrator will be unable to perform the scan. For more information, see the Veeam Recovery Orchestrator User Guide, section Malware Scan.

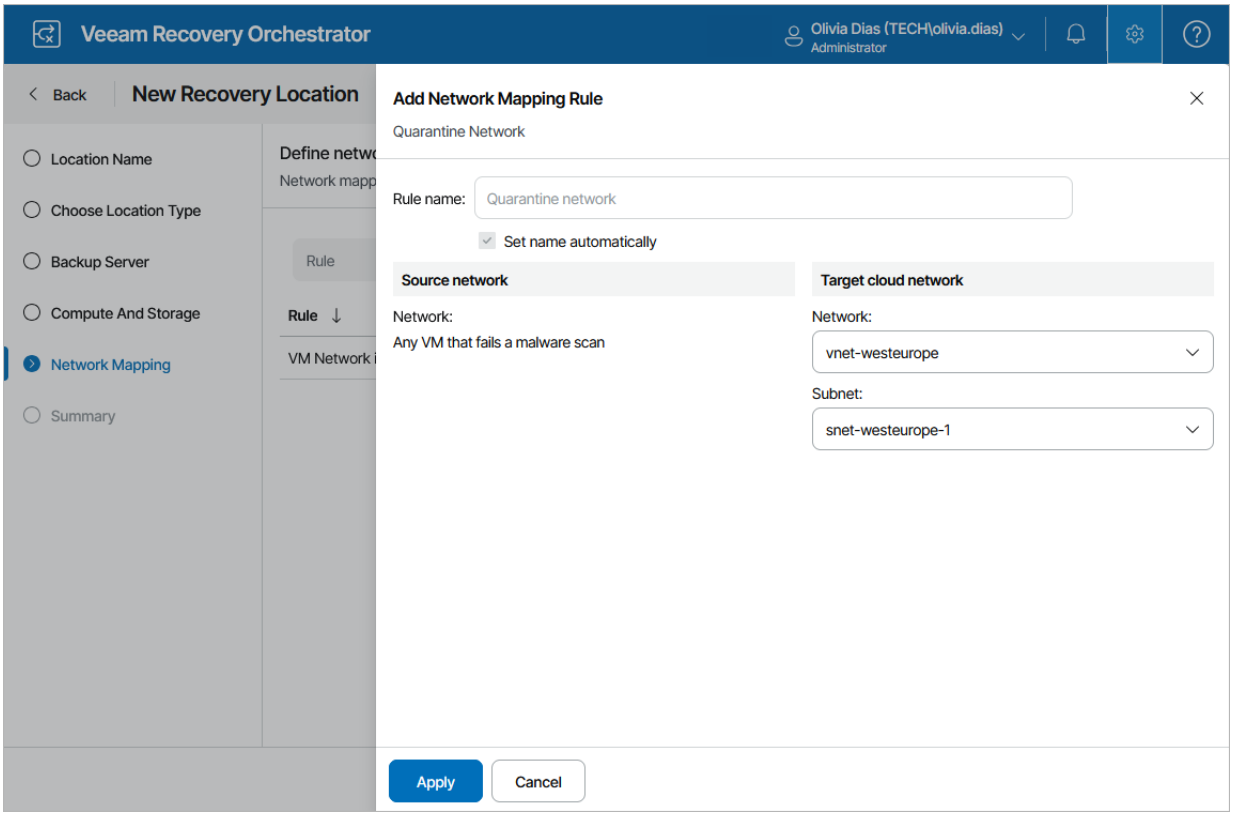
Veeam Recovery Orchestrator
Olivia Dias (TECH\olivia.dias)
Administrator
Back
New Recovery Location
Location Name
Choose Location Type
Backup Server
Compute And Storage
Network Mapping
Summary
Define netwo
Network mapp
Rule
Rule
VM Network i
Add Network Mapping Rule
Quarantine Network
Rule name:
Quarantine network
Set name automatically
Source network
Network:
Any VM that fails a malware scan
Target cloud network
Network:
vnet-westeurope
Subnet:
snet-westeurope-1
Apply
Cancel

Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.

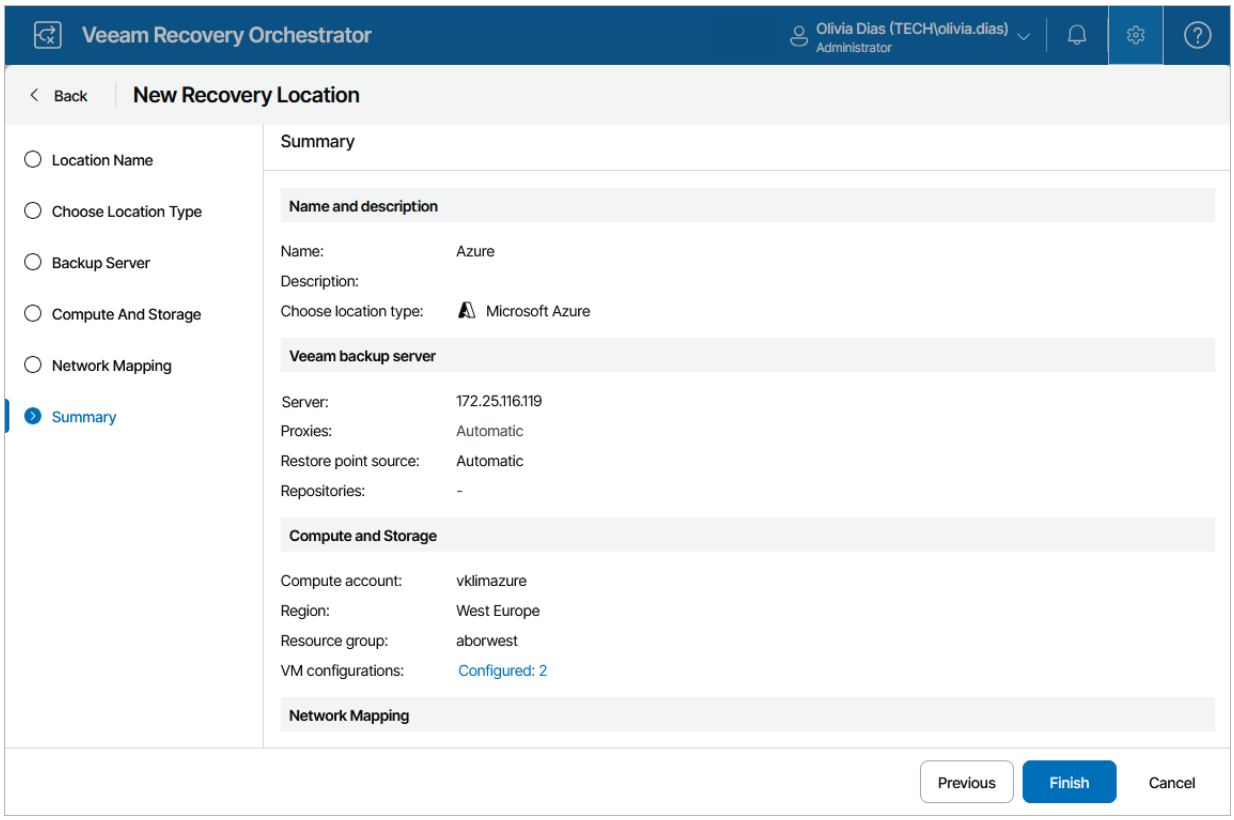

Adding Microsoft Hyper-V Recovery Locations

To add a Microsoft Hyper-V or Azure Local recovery location:

1. Switch to the **Administration** page.
2. Navigate to **Recovery Locations**.
3. Click **Add**.
4. Complete the **New Recovery Location** wizard:
   a. Specify a recovery location name and description.
   b. Choose a recovery location type.
   c. Specify the connection type of a target cluster.
   d. Choose recovery options.
   e. Choose a target SCVMM server and cluster.
   f. Configure network mapping.
   g. Finish working with the wizard.

IMPORTANT

Before you start adding a Microsoft Hyper-V recovery location, consider the following limitations:

- CSV disks that you plan to use to recover VMs must have the *Online* status.
- All hosts that belong to the cluster where the recovered VMs will reside must be online and available.

Keep in mind that after you bring a Hyper-V object online or perform any infrastructure configuration changes, the changes may not appear in the Orchestrator UI immediately — the data synchronization process between Orchestrator and Microsoft Hyper-V may take up to 2 hours to complete.

Step 1. Specify Recovery Location Name and Description

At the **Location Name** step of the wizard, use the **Name** and **Description** fields to enter a name for the new location and to provide a description for future reference. The maximum length of the location name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

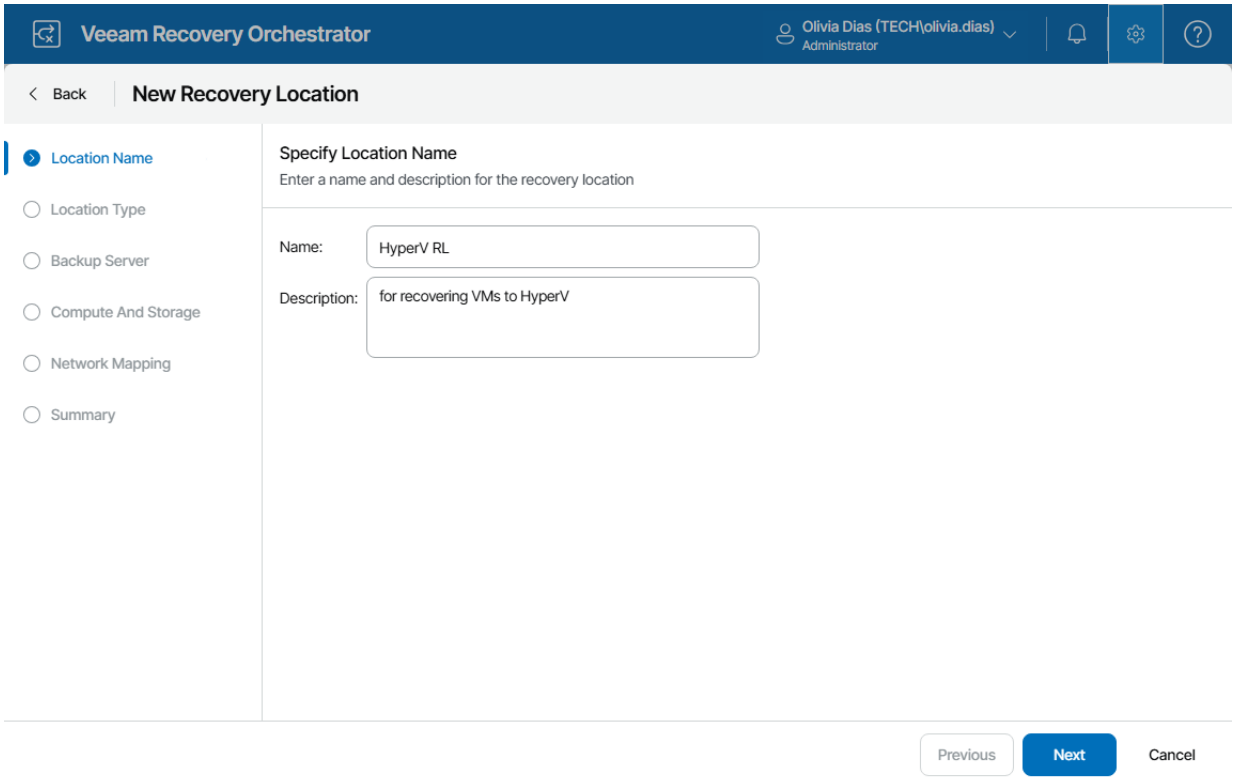

Step 2. Choose Recovery Location Type

At the **Location Type** step of the wizard, select the **Microsoft Hyper-V \ Azure Local** option.

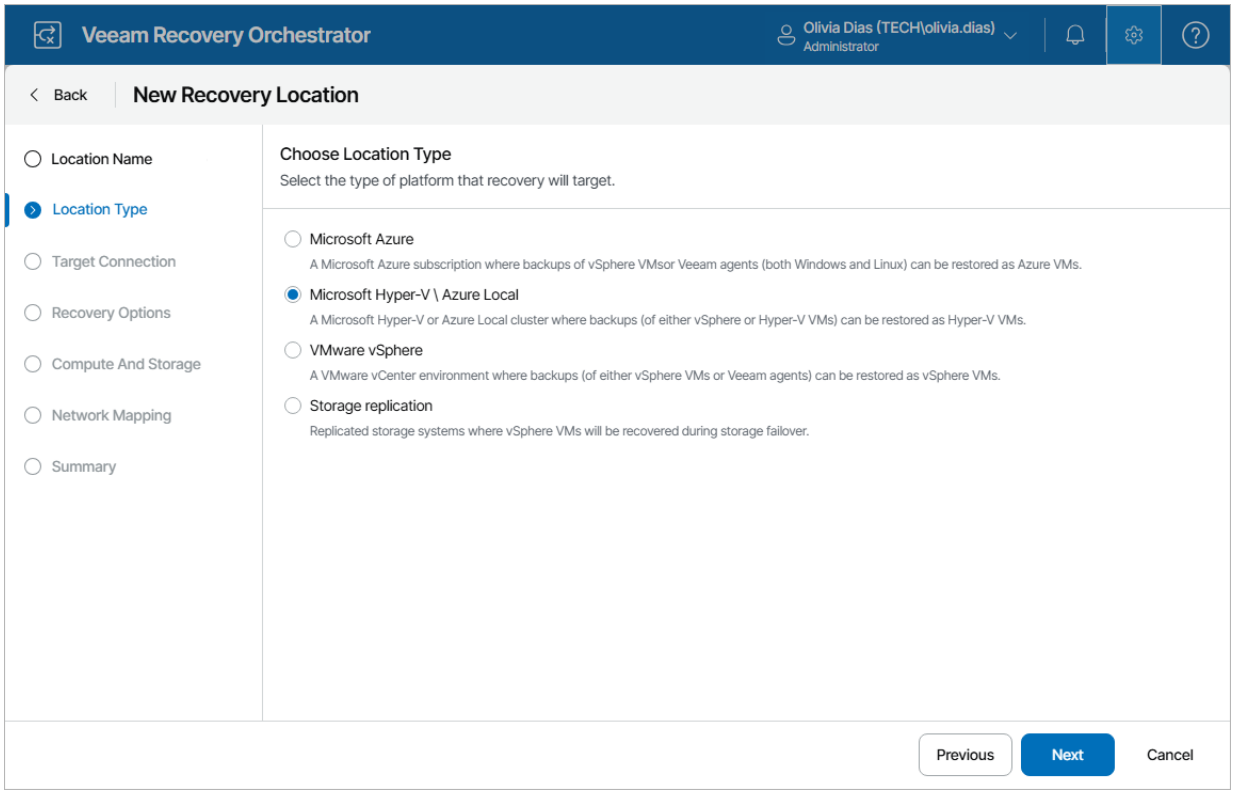

Step 3. Specify Connection Type

At the **Target Connection** step of the wizard, choose whether you want to use an SCVMM connection or a direct connection to a standalone Microsoft Hyper-V or Azure Local cluster to restore machines to this recovery location.

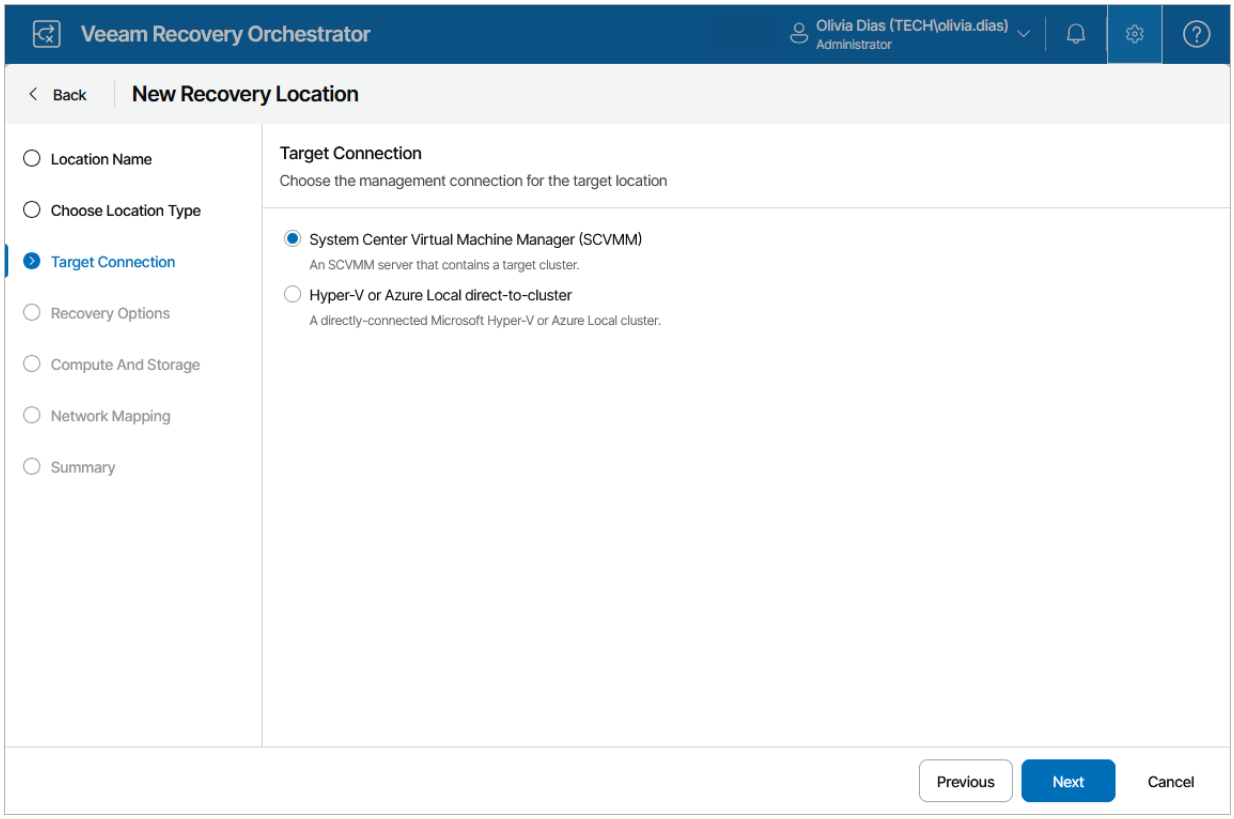

Step 4. Choose Recovery Options

At the **Recovery Options** step of the wizard, choose whether you want Orchestrator to use backup files created by backup jobs, their copies created by backup copy jobs or files that are stored in scale-out backup repositories. Alternatively, you can choose to let the product select the files automatically, based on the most recent restore points. For more information on the way Orchestrator selects backup files and restore points to recover machines when performing restore operations, see the Veeam Recovery Orchestrator User Guide, section How Orchestrator Selects Backup Files.

Additionally, you can configure the following settings:

- Choose whether you want to enable Instant VM Recovery for the location.

  With Instant VM Recovery, all processed machines will be immediately restored in the location by running directly from the backup files. Instant VM Recovery helps improve recovery time objectives (RTO), minimize disruption and downtime of the machines. For more information on the Instant VM Recovery feature, see the Veeam Backup & Replication User Guide, section Instant VM Recovery.

- Choose whether you want Orchestrator to be able to recover VMs to a different location in Veeam Backup & Replication.

  To control data migration in the virtual infrastructure, Veeam Backup & Replication introduces infrastructure locations. A location defines a geographic region where an infrastructure object resides. To learn how to create and assign locations to infrastructure objects in Veeam Backup & Replication, see the Veeam Backup & Replication User Guide, section Locations. To learn how to track geographical locations of production data, their copies and replicas, see the Veeam ONE Reporter User Guide, sections Data Sovereignty Overview and Data Sovereignty Violations.

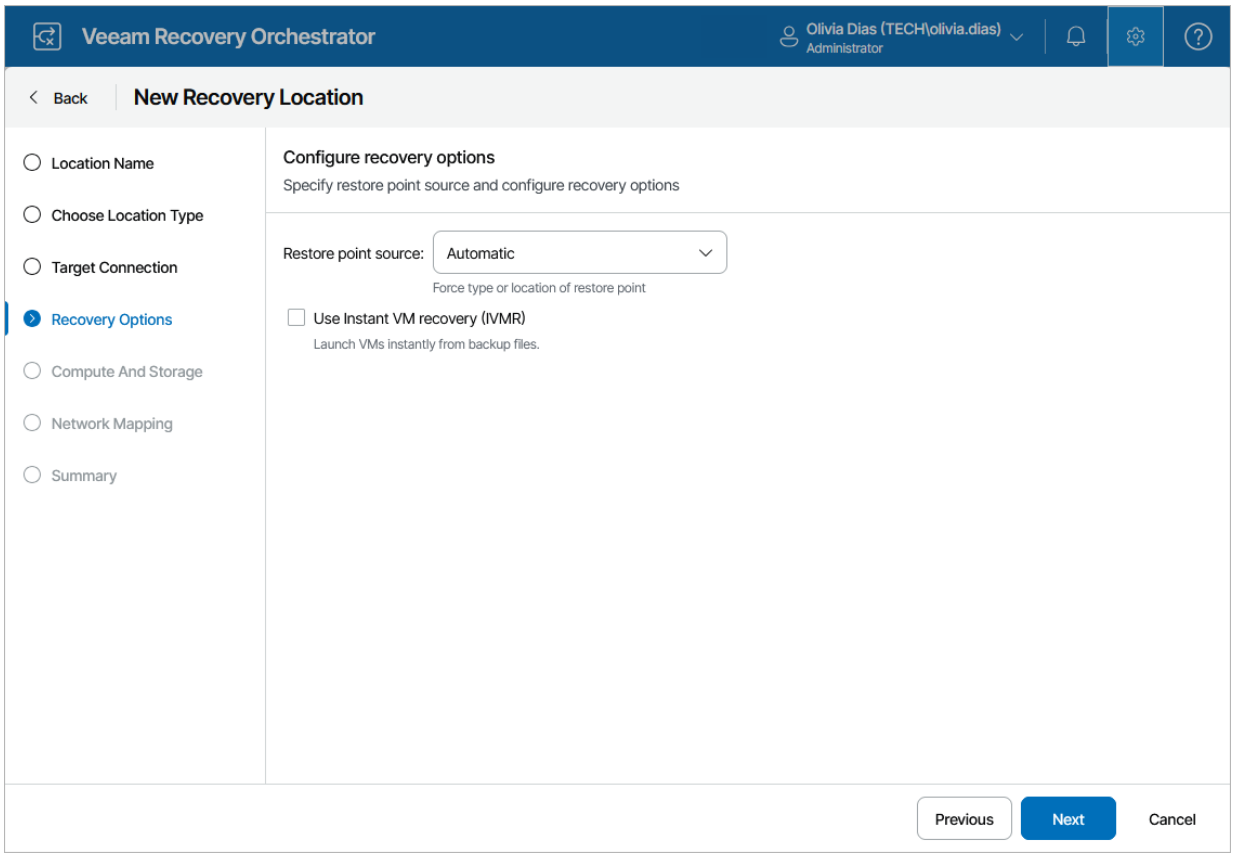

Step 5. Specify SCVMM Server and Cluster

At the **Compute and Storage** step of the wizard, do the following:

1. [Applies only if you have selected the System Center Virtual Machine Manager (SCVMM) option at the Target Connection step of the wizard] Select an SCVMM server that will manage the process of recovering machines to a Microsoft Hyper-V environment. Note that one recovery location can be associated with one SCVMM server only.

   For an SCVMM server to be displayed in the list of available servers, the SCVMM console must be installed on the machine that runs Orchestrator, and the server must be connected to Orchestrator as described in section Connecting Microsoft Hyper-V Servers.

2. Select a Microsoft Hyper-V or Azure Local cluster where the recovered VMs will reside.

   For a cluster to be displayed in the list of available clusters, it must be either managed by the SCVMM server selected at step 1, or added to Orchestrator as a direct connection as described in section Connecting Microsoft Hyper-V Servers.

3. Click **Add** to specify destination CSV disks where recovered VMs will be stored.

   For storage resources to be displayed in the **Storage** list, these resources must belong to the cluster selected at step 2.

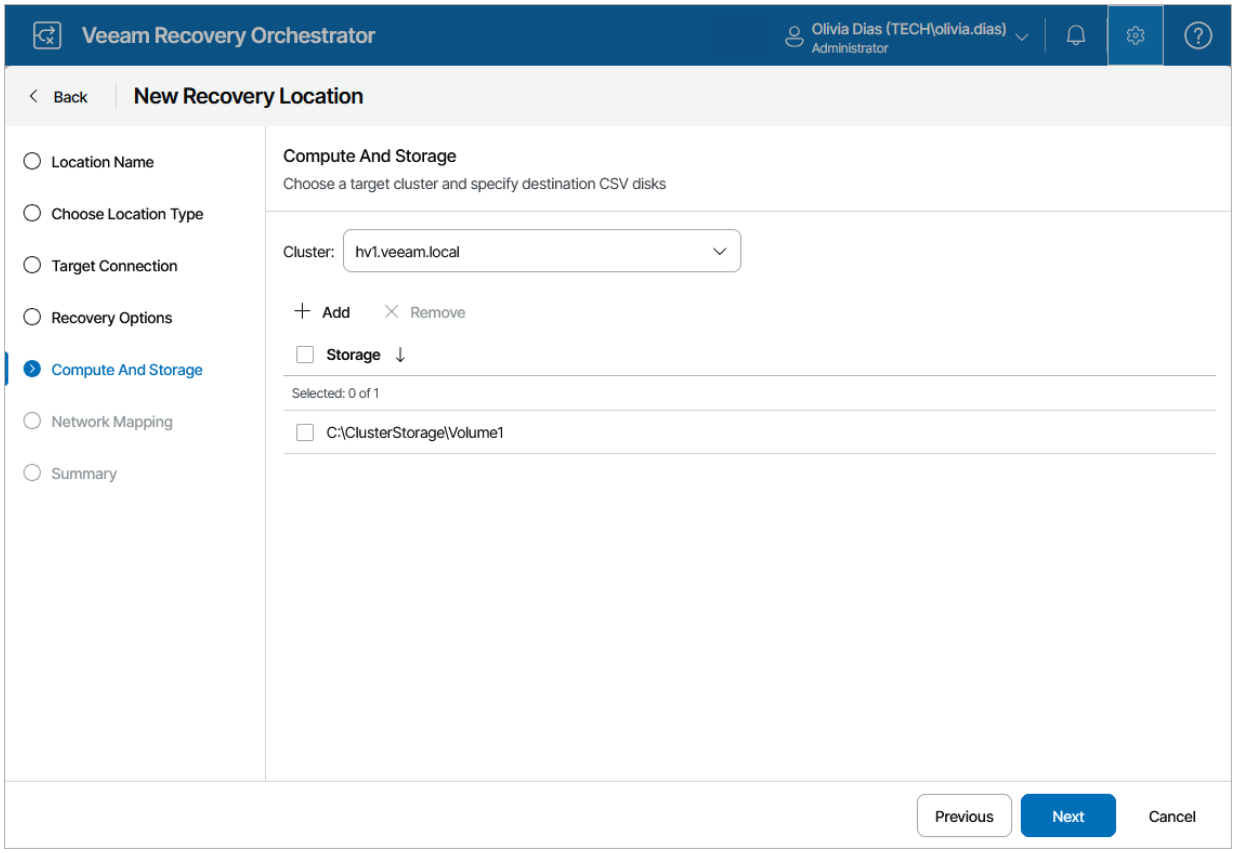

Step 6. Configure Network Mapping

When you recover a VM from a Hyper-V backup, the recovered VM is connected to the same Hyper-V networks as the source VM; if the same networks are not available in the recovery location, you can create a network mapping table for the location so that the recovered VM is connected to the correct network. However, when you recover a VM from a vSphere backup, there are no Hyper-V networks that can be used. Therefore, to recover a vSphere VM, you must create at least one network mapping rule so that the recovered VM is connected to the correct network.

To configure network mapping, at the **Network Mapping** step of the wizard, click **Add** > **VMware Mapping** to recover VMs from vSphere backups or **Add** > **Hyper-V Mapping** to recover VMs from Hyper-V backups. Then, do the following in the **Add Network Mapping Rule** window:

1. In the **Source network** section, select a vCenter Server or an SCVMM server that manages source VMs, a network to which the source VMs are connected, and a datacenter or a cluster where the source VMs reside.

   For a vCenter Server to be displayed in the **vCenter server** list, it must be connected to Orchestrator as described in section Connecting VMware vSphere Servers. For an SCVMM server to be displayed in the **SCVMM server** list, it must be connected to Orchestrator as described in section Connecting Microsoft Hyper-V Servers.

2. In the **Target Hyper-V network** section, select a network to which the recovered VMs will be connected.

   For a network to be displayed in the **Network** list, it must be configured on all the hosts of the cluster selected at step 5 of the wizard. For more information on how to configure network for Hyper-V hosts, see Microsoft Docs.

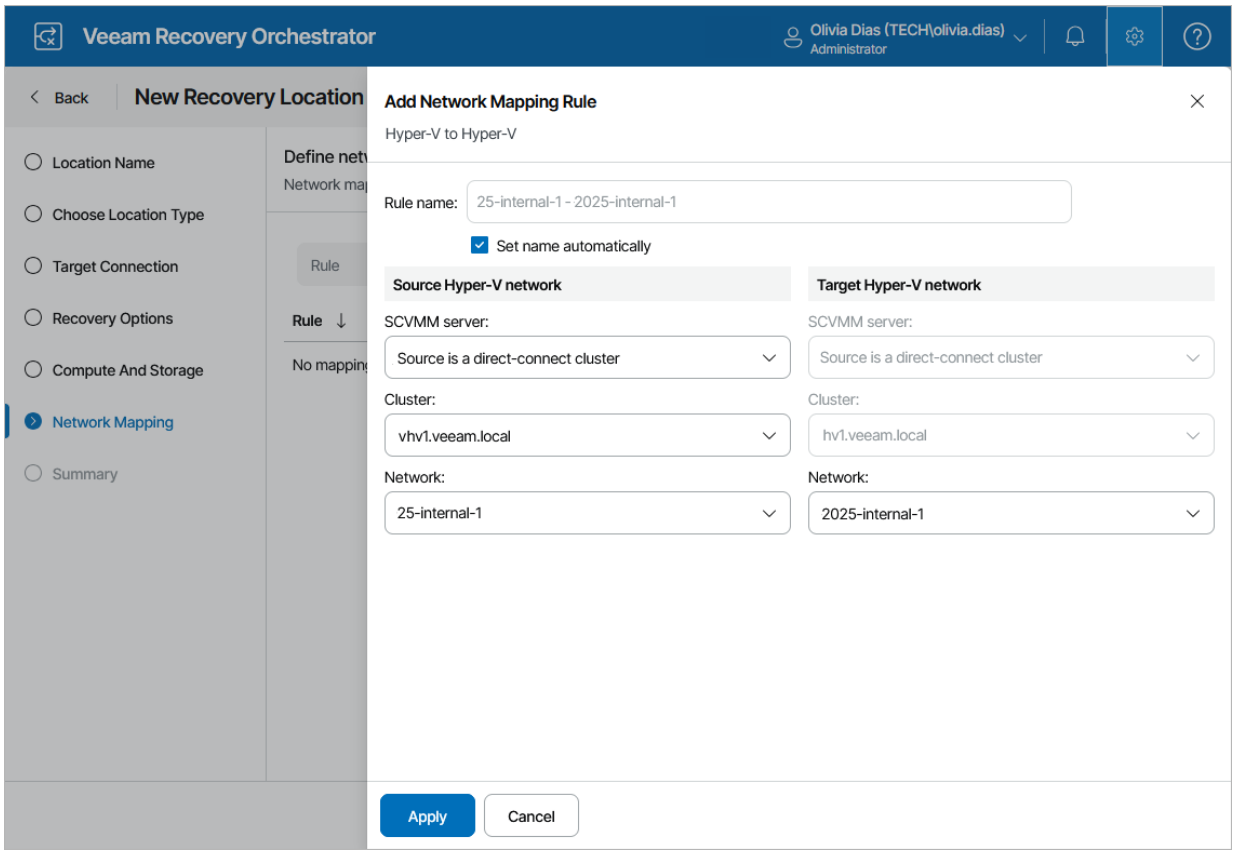

Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.

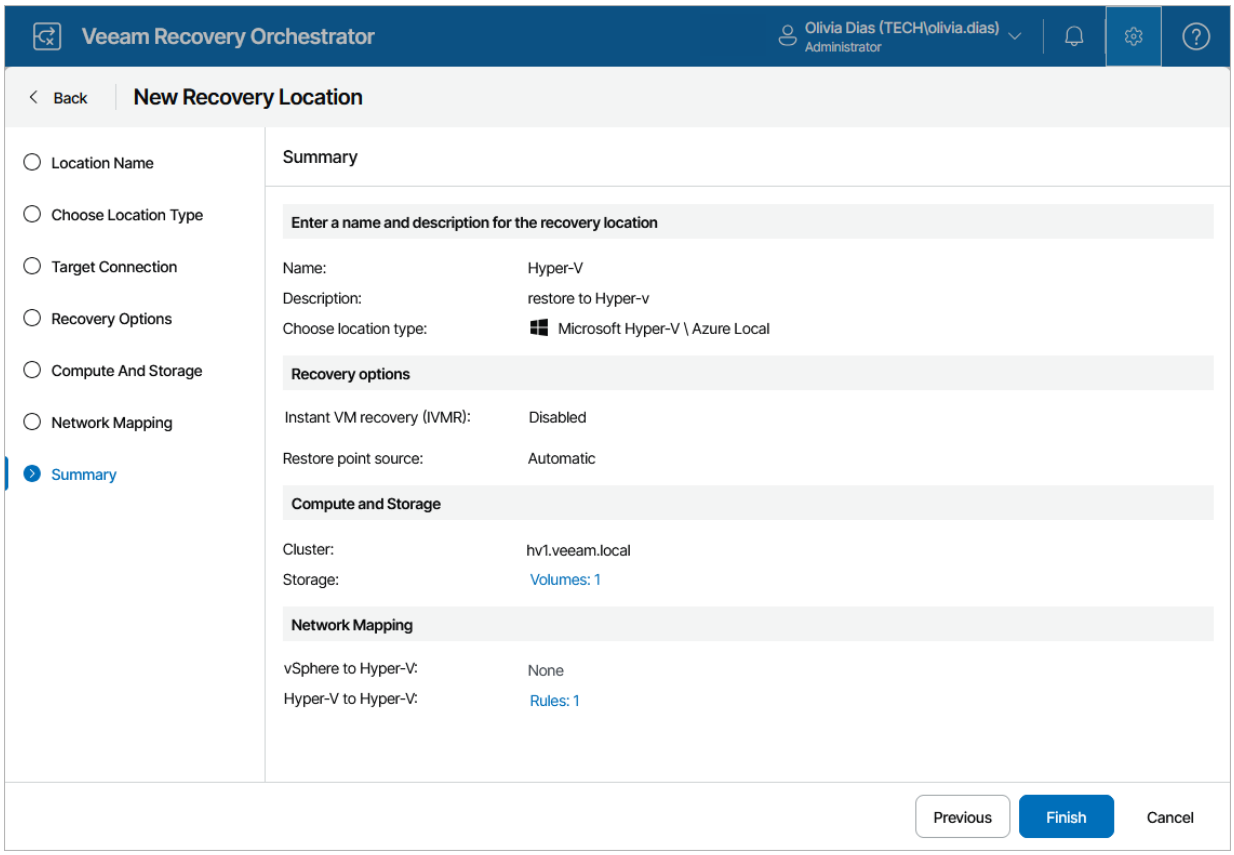

Editing Recovery Locations

If you want to change settings specified while adding a recovery location, the Orchestrator UI allows you to customize the location.

Editing VMware vSphere Recovery Locations

For each VMware vSphere recovery location, you can modify settings configured while creating the location:

1. Switch to the **Administration** page.
2. Navigate to **Recovery Locations**.
3. Select the location and click **Edit**.
4. Complete the **Edit Recovery Location** wizard:
    a. To change the name and description of the location, follow the instructions provided in section Adding VMware vSphere Recovery Locations (step 1).
    b. To change the specified recovery options, follow the instructions provided in section Adding VMware vSphere Recovery Locations (step 3).
    c. To change the specified compute resources, follow the instructions provided in section Adding VMware vSphere Recovery Locations (step 4).
    d. To change the specified storage resources, follow the instructions provided in section Adding VMware vSphere Recovery Locations (step 5).
    e. To configure network mapping and re-IP rules, follow the instructions provided in section Adding VMware vSphere Recovery Locations (step 6).
    f. At the **Summary** step of the wizard, review configuration information and click **Finish** to confirm the changes.

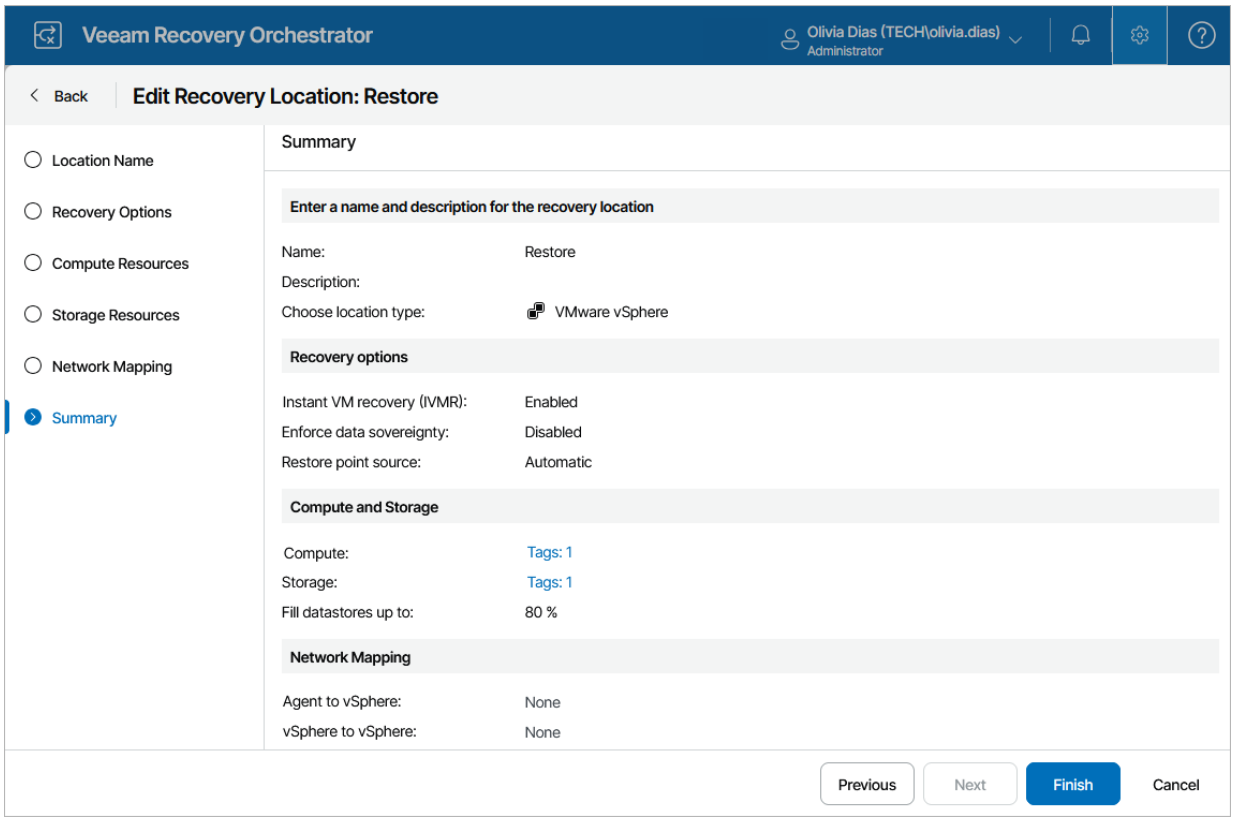

Editing Microsoft Hyper-V Recovery Locations

For each Microsoft Hyper-V recovery location, you can modify settings configured while creating the location:

1. Switch to the **Administration** page.
2. Navigate to **Recovery Locations**.
3. Select the location and click **Edit**.
4. Complete the **Edit Recovery Location** wizard:
   a. To change the name and description of the location, follow the instructions provided in section Adding Microsoft Hyper-V Recovery Locations (step 1).
   b. To change the connection of the selected cluster, follow the instructions provided in section Adding Microsoft Hyper-V Recovery Locations (step 3).
   c. To change the specified recovery options, follow the instructions provided in section Adding Microsoft Hyper-V Recovery Locations (step 4).
   d. To change the specified SCVMM server, cluster and target CSV disks, follow the instructions provided in section Adding Microsoft Hyper-V Recovery Locations (step 5).
   e. To configure network mapping, follow the instructions provided in section Adding Microsoft Hyper-V Recovery Locations (step 6).
   f. At the **Summary** step of the wizard, review configuration information and click **Finish** to confirm the changes.

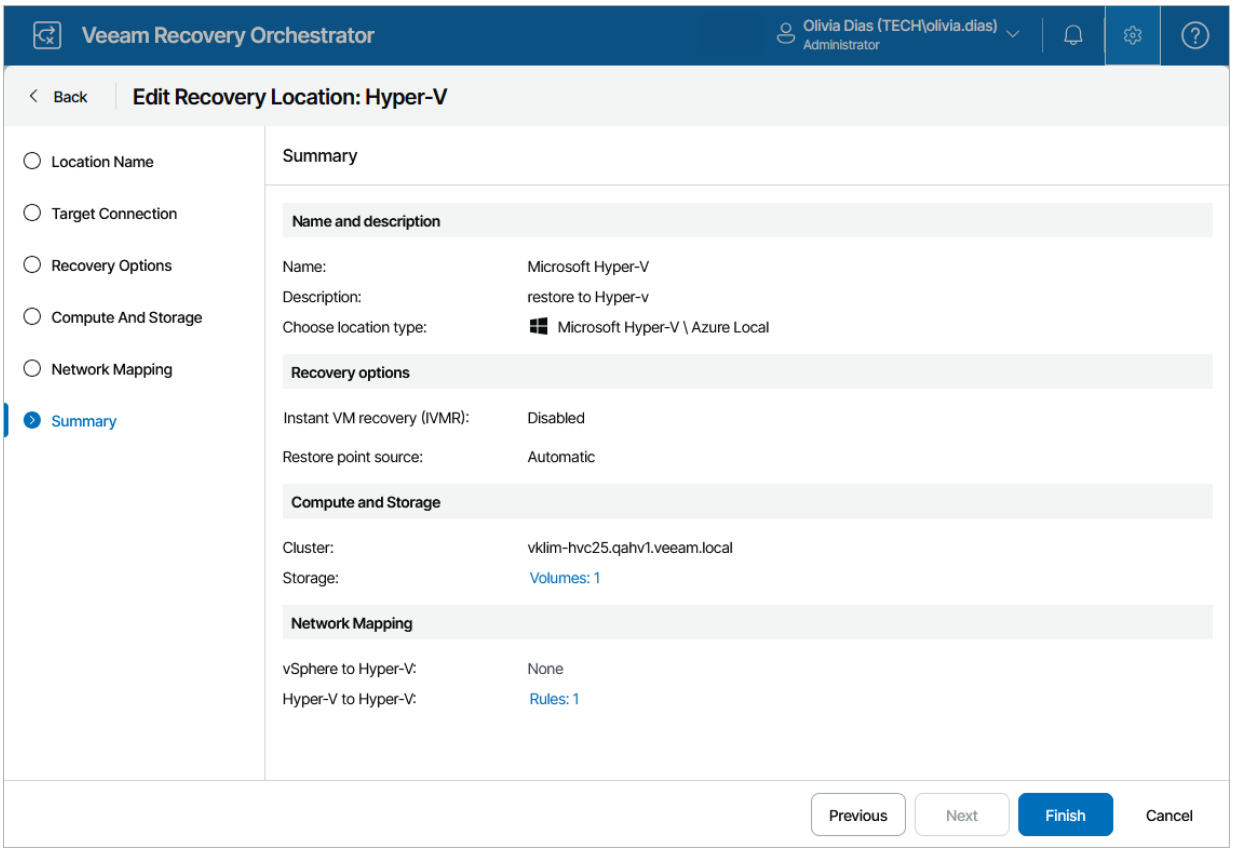

Editing Microsoft Azure Recovery Locations

For each Microsoft Azure recovery location, you can modify settings configured while creating the location:

1. Switch to the **Administration** page.
2. Navigate to **Recovery Locations**.
3. Select the location and click **Edit**.
4. Complete the **Edit Recovery Location** wizard:
   a. To change the name and description of the location, follow the instructions provided in section Adding Microsoft Azure Recovery Location (step 1).
   b. To modify the specified recovery options and change the Veeam Backup & Replication server that will manage the process of recovering machines to Microsoft Azure, follow the instructions provided in section Adding Microsoft Azure Recovery Locations (step 3).
   c. To configure the specified Microsoft Azure settings, follow the instructions provided in section Adding Microsoft Azure Recovery Locations (step 4).
   d. To configure network mapping and modify the specified quarantine network, follow the instructions provided in section Adding Microsoft Azure Recovery Locations (step 5).

e. At the **Summary** step of the wizard, review configuration information and click **Finish** to confirm the changes.

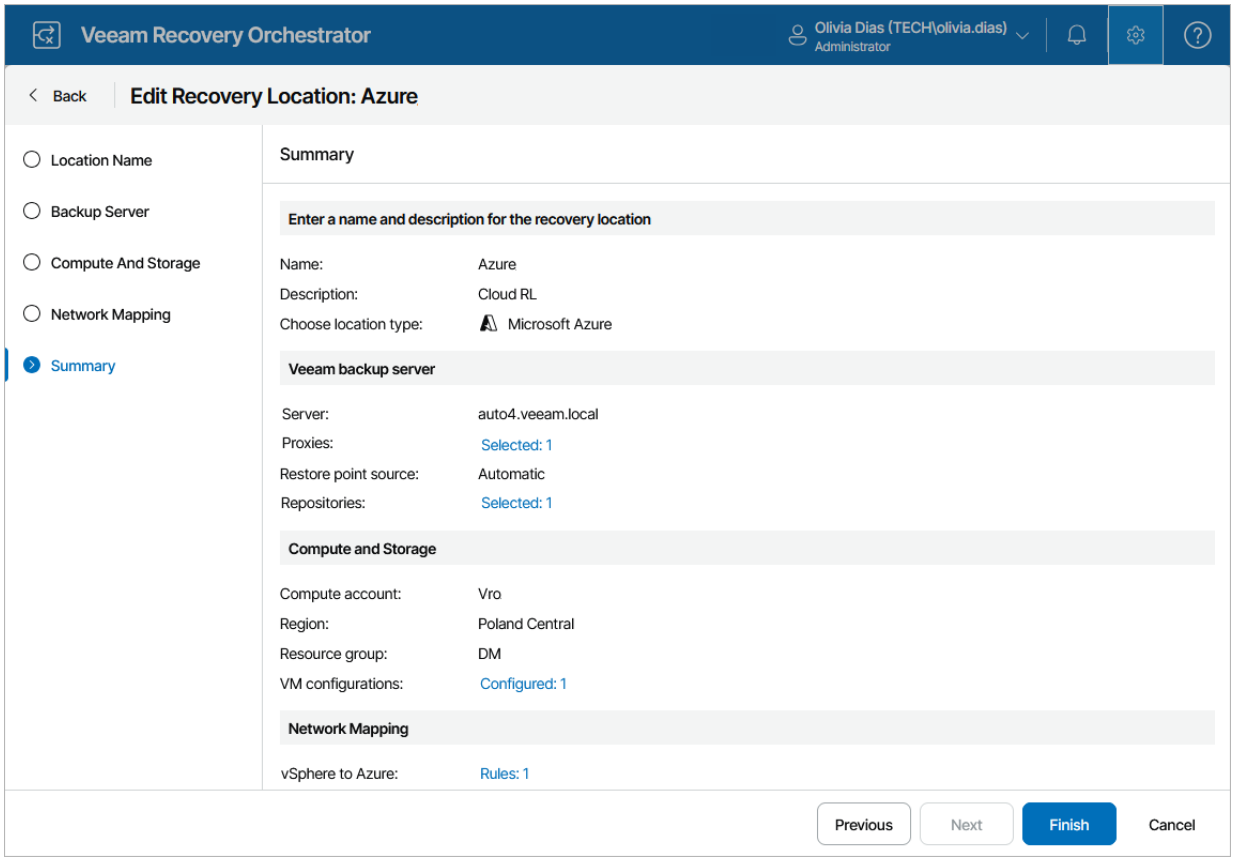

Editing Storage Recovery Locations

For each storage recovery location, you can modify settings configured while creating the location:

1. Switch to the **Administration** page.
2. Navigate to **Recovery Locations**.
3. Select the location and click **Edit**.
4. Complete the **Edit Recovery Location** wizard:
   a. To change the name and description of the location, follow the instructions provided in section Adding Storage Recovery Locations (step 1).
   b. To change the specified compute and storage resources, follow the instructions provided in section Adding Storage Recovery Locations (step 4).
   c. To configure network mapping and re-IP rules, follow the instructions provided in section Adding Storage Recovery Locations (step 5).
   d. At the **Summary** step of the wizard, review configuration information and click **Finish** to confirm the changes.

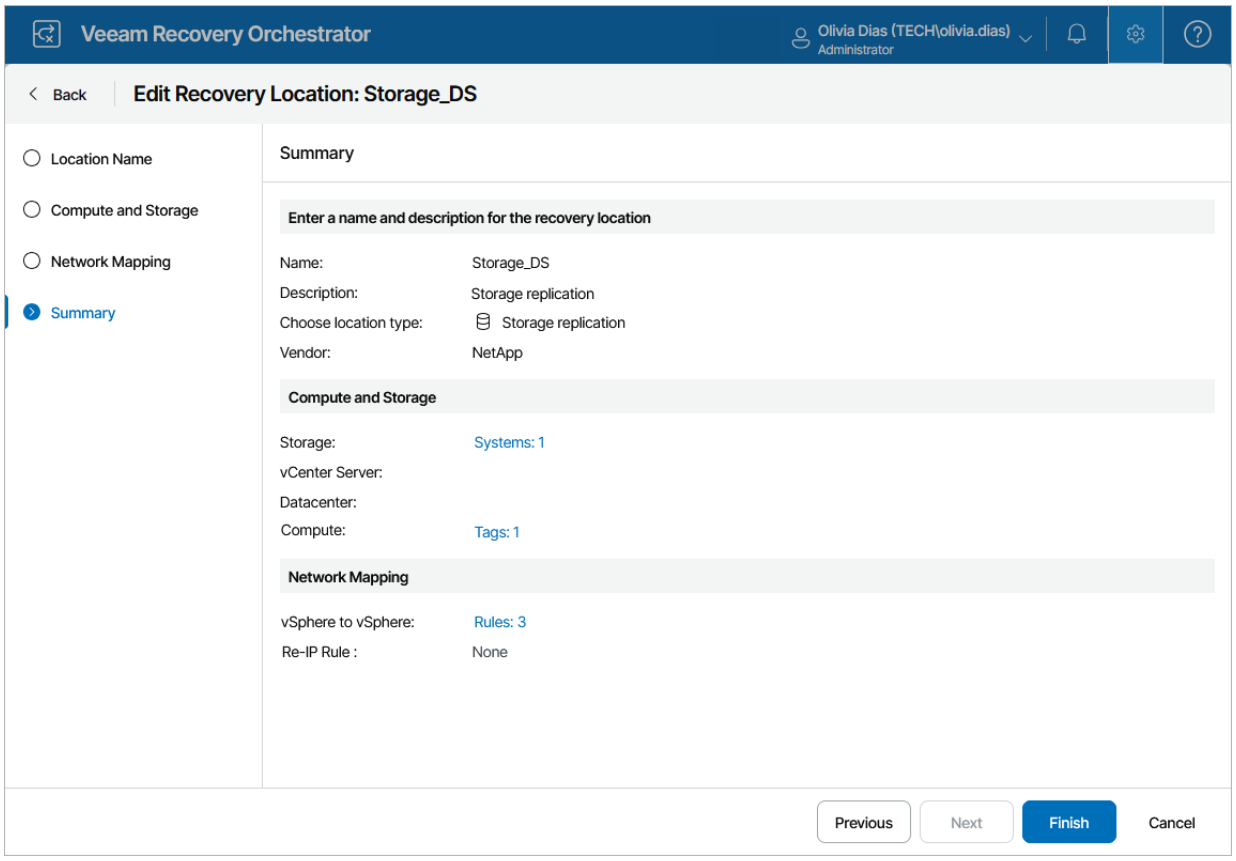

Editing Original VM Location

Since all resource groups included in the *Original VM Location* are empty by default and depend on processed machines, you cannot customize its resource settings. However, you can still customize some general settings for the location:

1. Switch to the **Administration** page.
2. Navigate to **Recovery Locations**.
3. In the list of recovery locations, select *Original VM Location*, and click **Edit**.

4. Complete the **Edit Recovery Location** wizard:

   a. To change the configured recovery options, follow the instructions provided in section Adding VMware vSphere Recovery Locations (step 3).

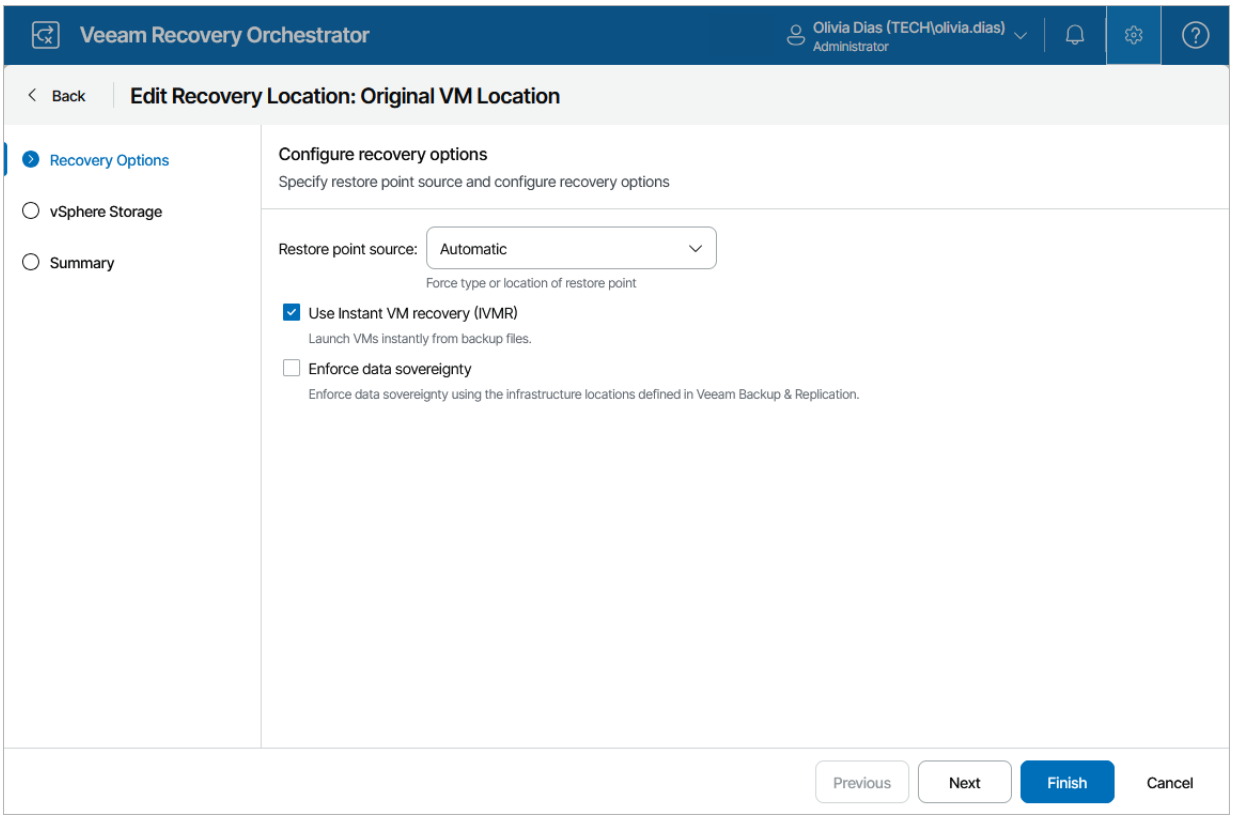

b. To change the datastore capacity level that must not be breached during the recovery process, follow the instructions provided in section Adding VMware vSphere Recovery Locations (step 3).

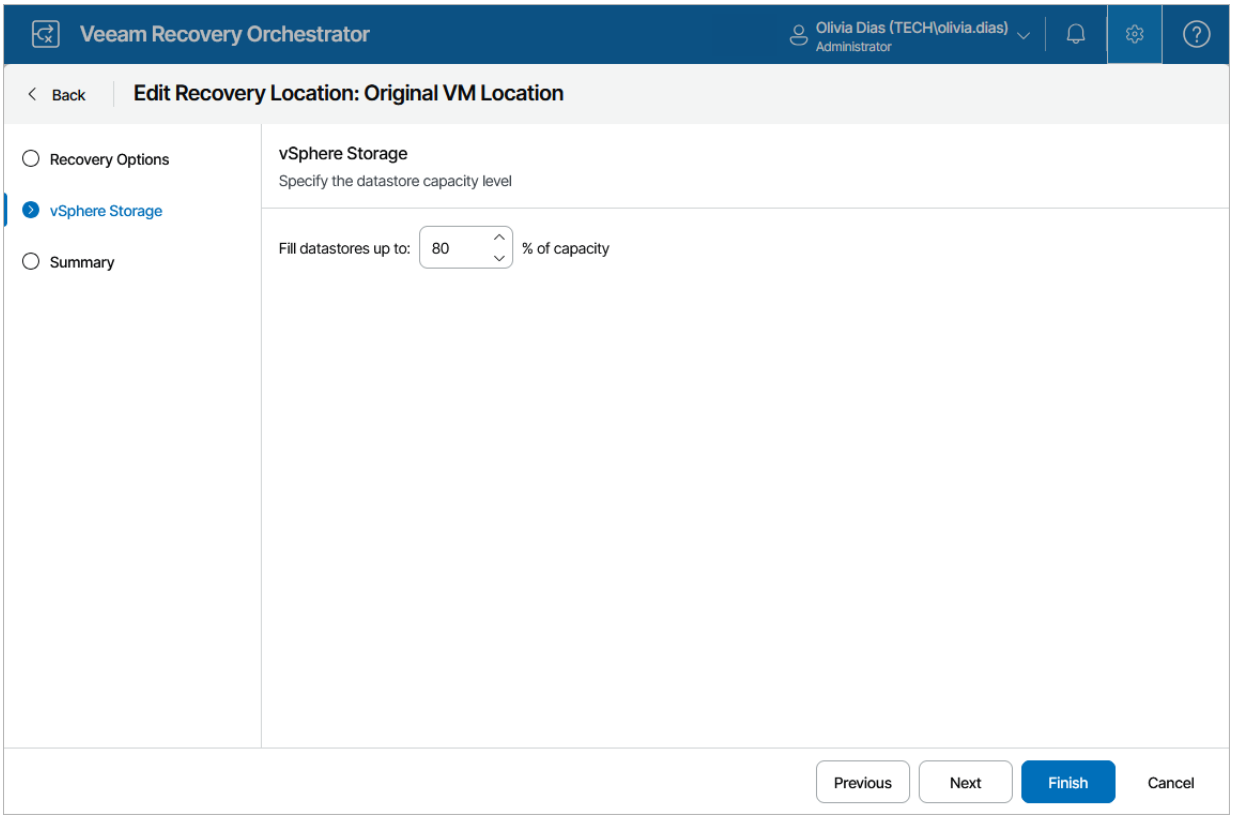

c. At the **Summary** step of the wizard, review configuration information and click **Finish** to confirm the changes.

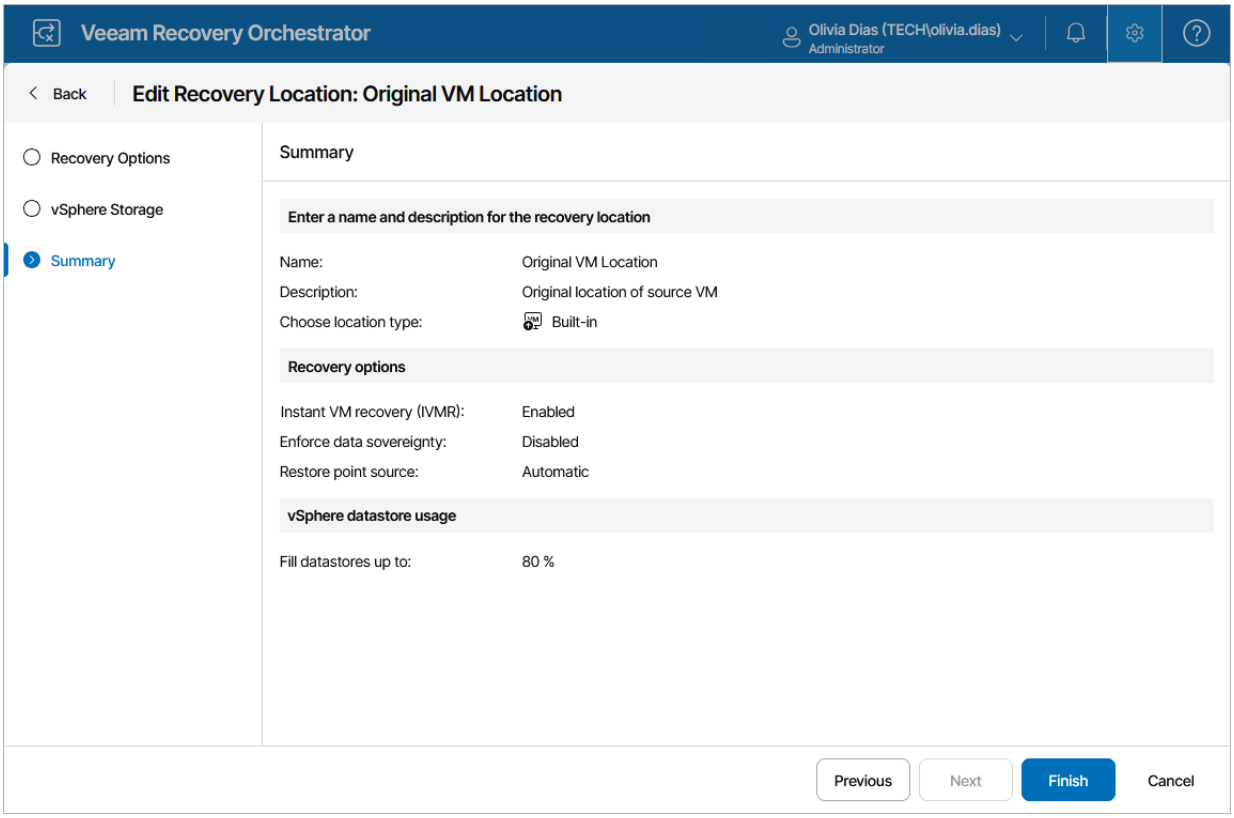

Cloning Recovery Locations

You can also create a new recovery location by cloning a location that already exists. The new location will have the same configuration as the existing location, which means that all items configured for the existing location will be applied to the new location.

To clone a recovery location:

1. Select an existing location that you want to use as a template for the new scope and click **Clone**.
2. In the **Clone Recovery Location** window:
   a. Use the **Name** and **Description** fields to enter a name for the new location and to provide a description for future reference.

      The maximum length of the scope name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

   b. Click **Clone** to save the location.

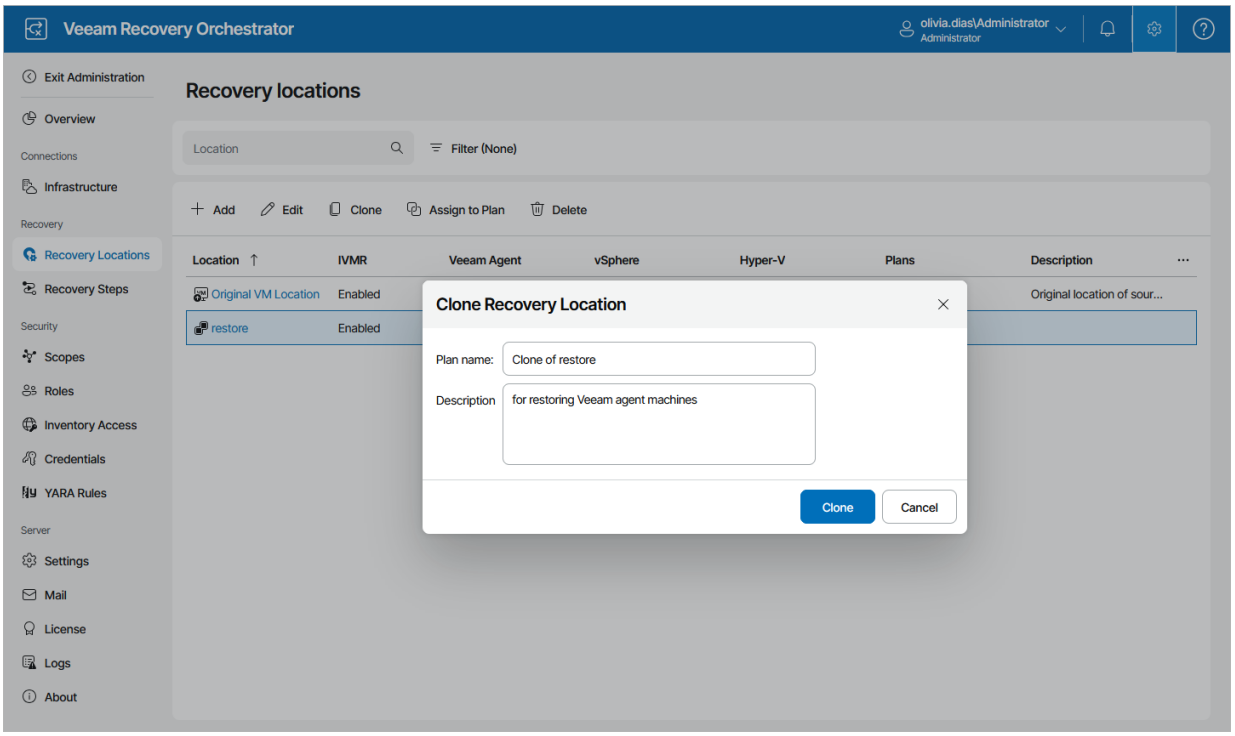

Configuring Plan Steps

Plan steps are the sequence of actions taken by the Orchestrator server during plan execution. For each machine in an inventory group included in a recovery plan, 2 types of steps can be performed:

- Default steps provided by Veeam
- Custom steps added by end user

For detailed description of all available recovery plan steps, see Appendix A. Recovery Plan Steps.

> **NOTE**
>
> You cannot delete built-in steps provided by Veeam. This option is available for custom steps only.

Configuring Default Parameter Settings

To modify default parameter settings for a plan step:

1. Switch to the **Administration** page.
2. Navigate to **Recovery Steps**.
3. In the **Step** column, select the step and click **Edit**.
4. In the **Parameters** section of the **Step Editor** window, set the desired parameter values and click **Save**.

> **NOTE**
>
> The parameter settings will be changed only for plans that are NOT in the *IN-USE* mode. For the list of modes that a recovery plan can acquire, see Replica Plans, CDP Replica Plans, Restore Plans, Storage Plans and Cloud Plans.

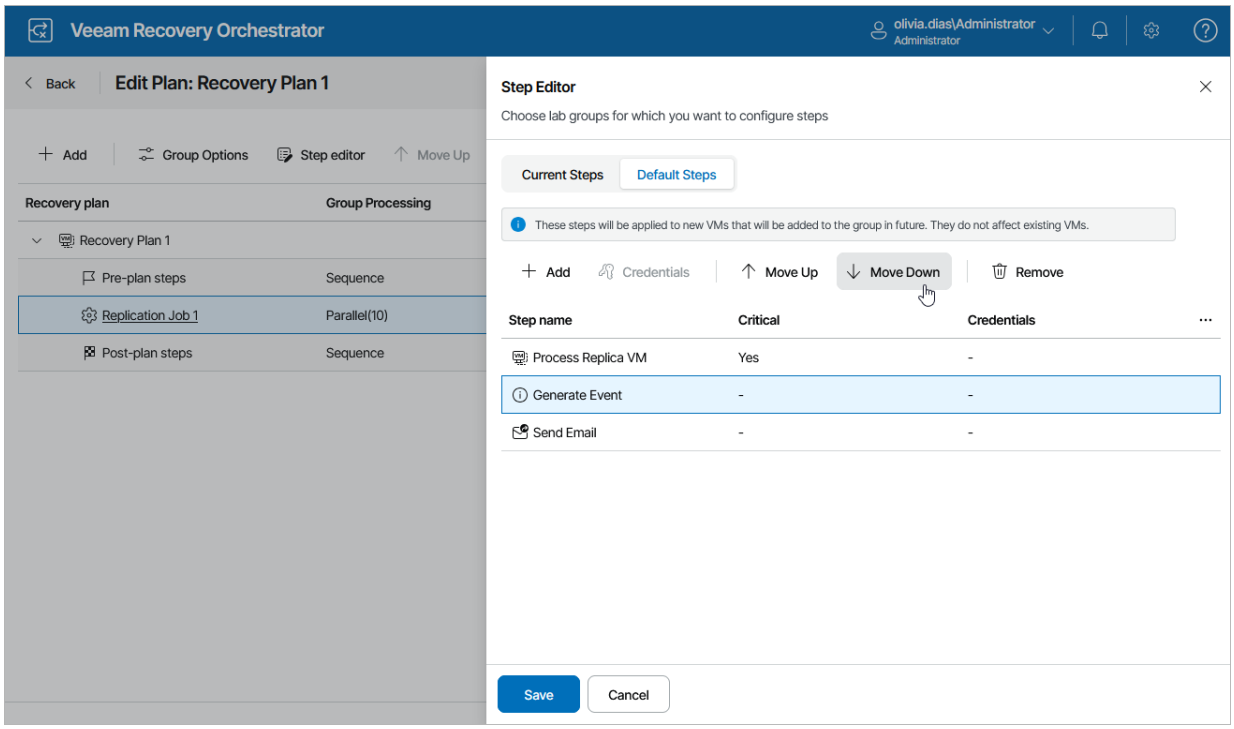

Managing Credentials

Orchestrator Administrators can choose which credentials will be visible to Plan Authors and available for use in recovery plans. These credentials can be used in plans, for example, to run verification scripts inside the guest OS, or in other checks.

The Orchestrator UI allows you to edit the existing credentials, and add any domain and non-domain accounts.

> **IMPORTANT**
>
> Starting from Orchestrator version 13, credentials from the Orchestrator UI are not synchronized into the Veeam Backup & Replication console. This means that if you change a password in the Orchestrator UI, you must also update this password manually in the Veeam Backup & Replication console as described in the Veeam Backup & Replication User Guide, section Managing Credentials.

Adding Credentials

If you want to manually add credentials under which recovery plan steps will be launched:

1. Switch to the **Administration** page.
2. Navigate to **Credentials**.
3. Click **Add**.
4. In the **Add Credentials** window, click **Browse**.
5. In the **Browse for Account** window:
    a. In the **Location** field, select whether the account that you want to add belongs to a domain or to a local OS workgroup.
    b. In the **Account** field, enter the account name.
    c. Select the account and click **Add**.
6. In the **Add Credentials** window, enter a password for the account that you want to add, provide a description for future reference, and click **Save**.

> **TIP**
>
> You can also add any credentials of your choice, even those that do not exist yet. To do that, in the **Add Credentials** window, use the **Account** and **Password** fields to enter an account name and a password for the account, and click **Save**.

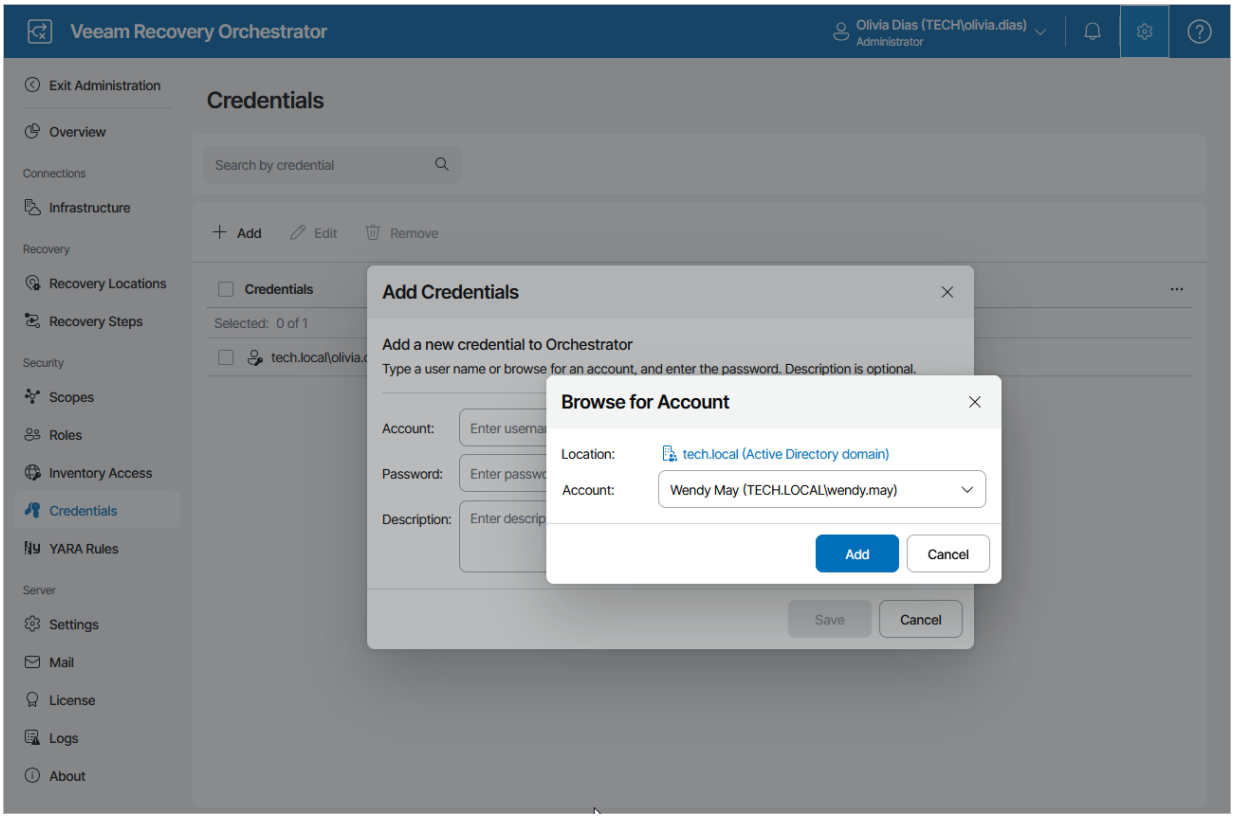

By default, all credentials are not added to newly created scopes; only the *Default Scope* has all credentials added. To edit the list of credentials available for a scope and to add the new credentials, follow the instructions provided in section Managing Inventory Items.

Changing Passwords

To change a password for an account in the Orchestrator UI, do the following:

1. Switch to the **Administration** page.
2. Navigate to **Credentials**.
3. Select a credential that you want to modify and click **Edit**.
4. In the **Edit Credentials** window:
   a. In the **Password** field, enter a new password.
   b. Click **Save**.

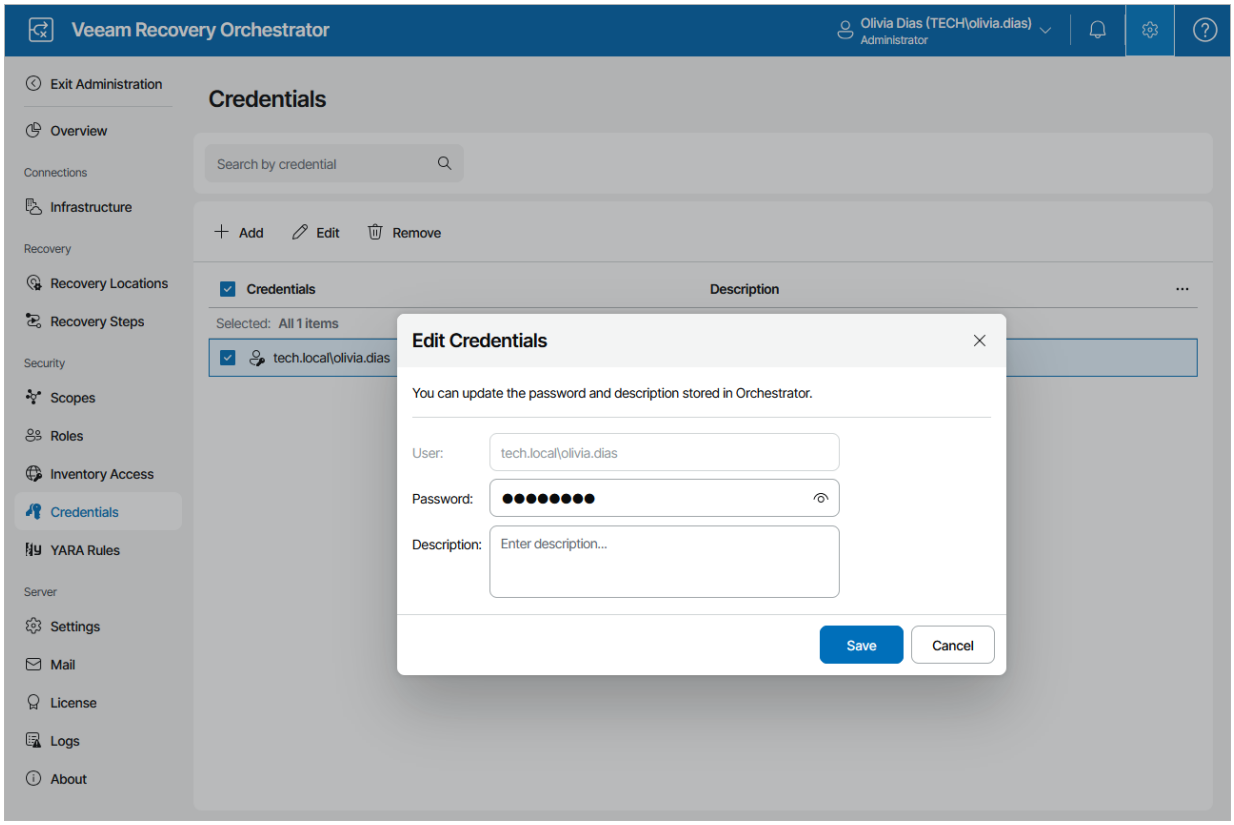

Editing Template Jobs

When you create a replica or restore plan, you have an option to reprotect machines included in the plan as soon as the recovery process completes. Orchestrator will automatically create a new replication or backup job to reprotect the recovered VMs as part of the plan execution process. Keep in mind that you cannot reprotect VMs recovered to a Microsoft Hyper-V environment.

To accomplish this, configure a template job on the Veeam Backup & Replication server that protects the required machines and is connected to your Orchestrator server. For Orchestrator to discover the job as a template, you must create a standard backup or replication job and include the text *[VRO Template]* in the job description.

> **NOTE**
>
> When creating a new replication or backup job, Orchestrator will copy all settings configured for the template job — except for the guest processing settings. If you want to enable application-aware processing for machines included in the plan, edit settings of the newly created job as described in the Veeam Backup & Replication User Guide, sections Creating Backup Jobs and Creating Replication Jobs.

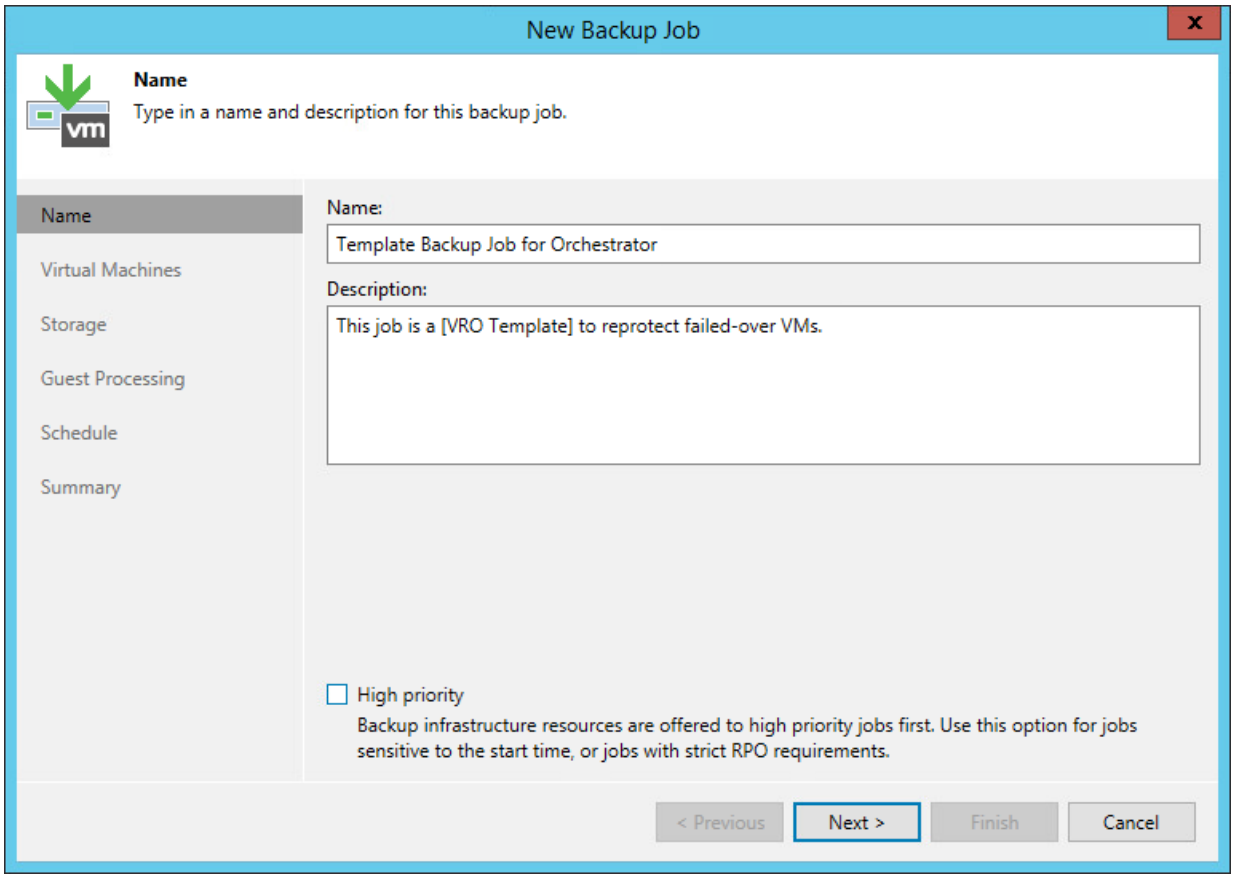

After you create a template job on the Veeam Backup & Replication server, Orchestrator will collect this data and display it in the Orchestrator UI. Note that the data synchronization process between Orchestrator and the Veeam Backup & Replication server may take several minutes to complete.

> **NOTE**
>
> The Orchestrator template job can be configured per-inventory group in a plan. To do that, select the **Reprotect group** check box when editing the plan. With this option selected, a new job will be created for the inventory group. For more information, see Configuring Group Settings.

Managing YARA Rules

YARA is a tool that allows you to create malware detection patterns (rules) that can be customized to detect attacks and security threats specific to your environment. To perform YARA scan for a plan, you must import a YARA rule file to Orchestrator and then add this rule to the list of inventory items for a scope. When you run a plan with YARA scan enabled, Veeam Backup & Replication first mounts machine disks from backups to the mount server and then initiates a scan session.

You can use YARA scan to do the following:

- To find the most recent clean restore point. In this case the backup server iterates through the number of restore points specified while running the plan one by one to detect a restore point with no malware. If no clean restore point is found, the scan session completes with an error.
- To analyze the content of restore points for specific information (for example, sensitive data). If sensitive data is found, the scan session completes with an error.

If the scan session completes with an error, Orchestrator performs either of the following actions:

- When running a cloud plan, Orchestrator halts the plan or restores the machine to a quarantine network depending on the configured restore point settings.
- When running a restore plan, Orchestrator either halts the plan or restores the machine to the selected recovery location without connecting it to any network, depending on the configured restore point settings.

Adding YARA Rule Files

To add a YARA file, do the following:

1. Switch to the **Administration** page.
2. Navigate to **YARA Rules**.
3. Click **Add**.
4. Complete the **Add YARA Rule File** wizard:
   a. Click **Browse** to locate a YARA rule file.
   b. Provide a nickname and description for the file. The maximum length of the rule name is 128 characters; the following characters are not supported: * : / \ ? " < > | .
   c. Click **Save**.

> **NOTE**
>
> You can add only one YARA rule file at a time.

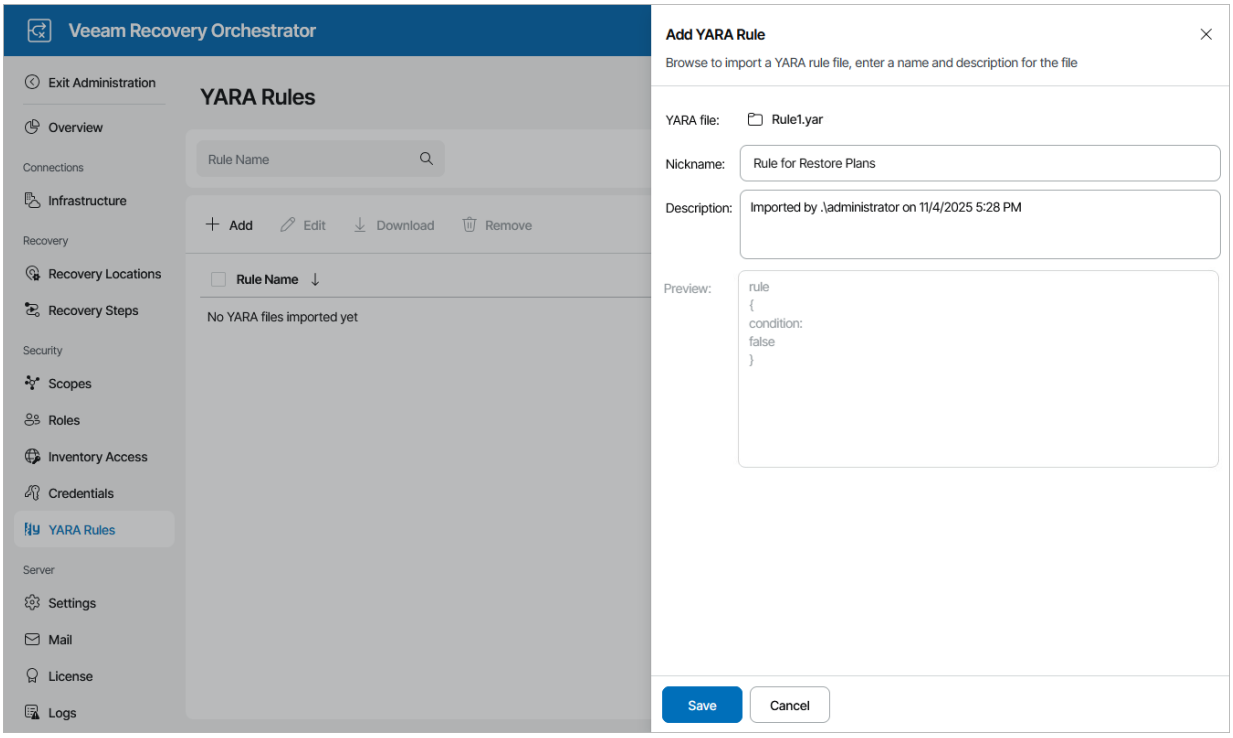

By default, all YARA rules are not added to newly created scopes; only the *Default Scope* has all rules added. To edit the list of rules available for a scope and to add the new rules, follow the instructions provided in section Managing Inventory Items.

Managing Inventory Items

Unless an item is added to the list of inventory items for a scope, it will not be available for use in the scope. By default, all items are not added to newly created scopes; only the *Default Scope* has all items added.

> **TIP**
>
> If an item is added to more than one scope, it will appear more than once in the list of inventory items — once for each scope. To view the relevant scope memberships, use the search fields displayed on the item tabs.

To modify the list of items available for a scope:

1. Switch to the **Administration** page.
2. Navigate to **Inventory Access**.
3. Switch to the necessary tab.
4. Select an item that you want to add to the scope:
   a. Click **Add Scope**.
   b. In the **Add to Scope** window, select the scope from the drop-down list, and click **Apply**.

> **TIP**
>
> You can simultaneously modify the list of inventory items available for multiple scopes. To do that, select check boxes next to the required items and click **Add Scope** or **Edit Scope**. After you select a scope from the drop-down list in the **Add to Scope** or **Edit Scope** window, the changes will be applied to all the selected item at the same time.

After you add an item to a scope, Plan Authors will be able to use this item when building recovery plans for the scope, particularly:

- Use template jobs to protect machines included in replica and restore plans.
- Use credentials when configuring the **Windows Credentials** and **SQL Credentials** parameters for plan steps.
- Use recovery locations to recover machines included in restore and cloud plans, and also when running failback.
- Use DataLabs when performing on-demand and scheduled testing of recovery plans. For more information, see Testing Recovery Plans.

For more information on creating and editing recovery plans, see Working with Replica Plans, Working with CDP Replica Plans, Working with Restore Plans and Working with Cloud Plans.

> **NOTE**
>
> By design, the list of available recovery locations will always display VMware vSphere, Microsoft Hyper-V and Microsoft Azure recovery locations only. For storage recovery locations, there is no need to allow access — Orchestrator automatically identifies the locations to be used when running storage plans. For more information on the way Orchestrator analyzes storage recovery locations, see the Veeam Recovery Orchestrator User Guide, section Storage Failover.

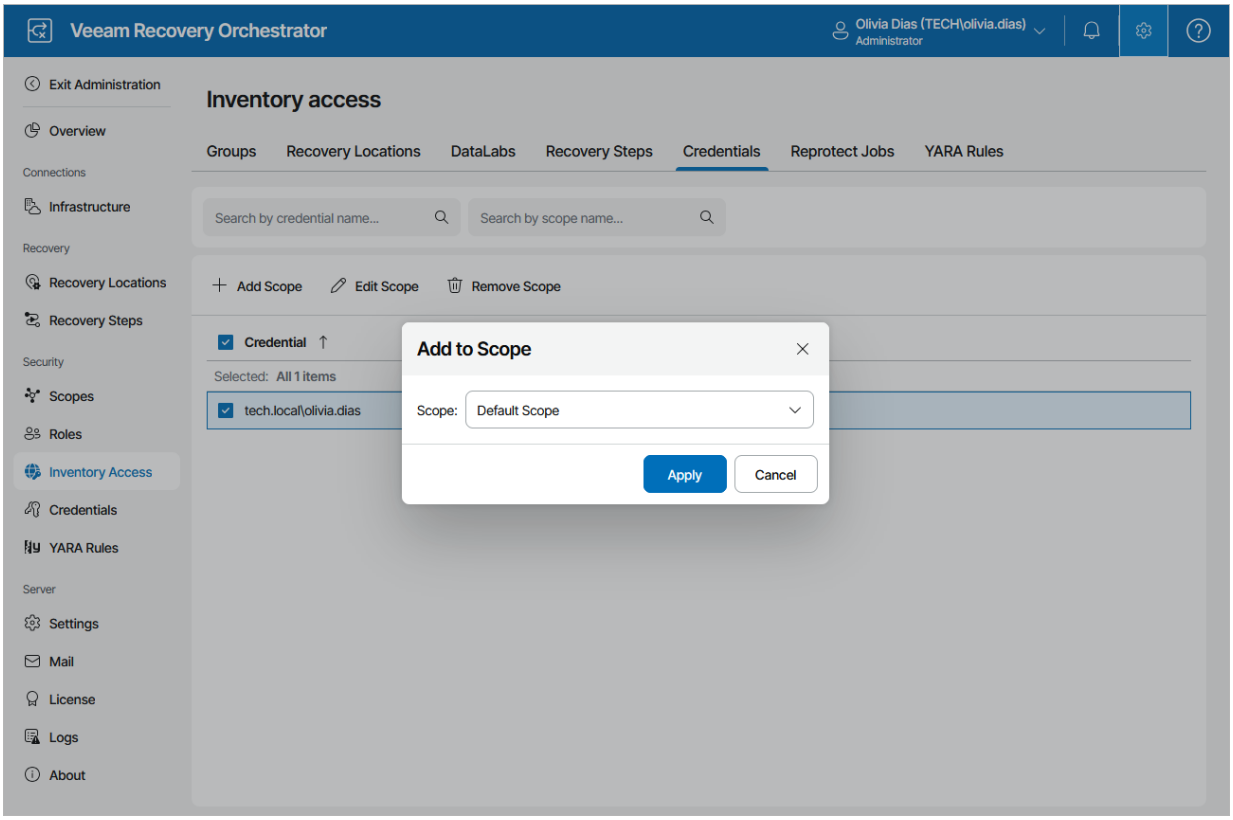
Veeam Recovery Orchestrator
Olivia Dias (TECH\olivia.dias)
Administrator
Exit Administration
Overview
Connections
Infrastructure
Recovery
Recovery Locations
Recovery Steps
Security
Scopes
Roles
Inventory Access
Credentials
YARA Rules
Server
Settings
Mail
License
Logs
About
Inventory access
Groups
Recovery Locations
DataLabs
Recovery Steps
Credentials
Reprotect Jobs
YARA Rules
Search by credential name...
Search by scope name...
Add Scope
Edit Scope
Remove Scope
Credential
Selected: All 1 items
tech.local\olivia.dias
Add to Scope
Scope:
Default Scope
Apply
Cancel

Working with Recovery Plans

Orchestrator offers 5 types of recovery plans:

- Replica plans
- CDP replica plans
- Restore plans
- Storage plans
- Cloud plans

Recovery plans are based on inventory groups. For more information on inventory groups, see the Veeam Recovery Orchestrator User Guide, section Inventory Groups.

> **TIP**
>
> To see the full list of discovered inventory groups, navigate to **Scope Inventory**. You can select any inventory group, click **Add to Plan**, and then click **New Plan** to create a recovery plan that will have the selected group preselected in the list of groups to recover.

Recovery plans can be scheduled and chained to execute in sequence, and Orchestrator will automatically produce and update detailed documentation. Execution of recovery plans is simplified to allow simultaneous management of multiple plans that contain hundreds of machines.

Working with Replica Plans

The type of a recovery plan you create depends on whether you intend to use Orchestrator to switch to VM replicas, to restore machines from backups or backup copies, or to serve data from a destination (NetApp) or secondary (HPE) volume in case a disaster strikes.

If you want to recover vSphere VMs protected by Veeam replication jobs, create a replica plan.

Creating Replica Plans

To create a replica plan:

1. Navigate to **Recovery Plans**.
2. Click **Add New Plan**.
3. Complete the **New Recovery Plan** wizard:
    a. Choose a type of the plan.
    b. Choose the necessary replica type.
    c. Choose a scope for the plan.
    d. Specify a plan name and description.
    e. Specify the target RTO and RPO.
    f. Select a template for plan reports.
    g. Finish working with the wizard.

Step 1. Choose Plan Type

At the **Plan Type** step of the wizard, select the **Replica** option.

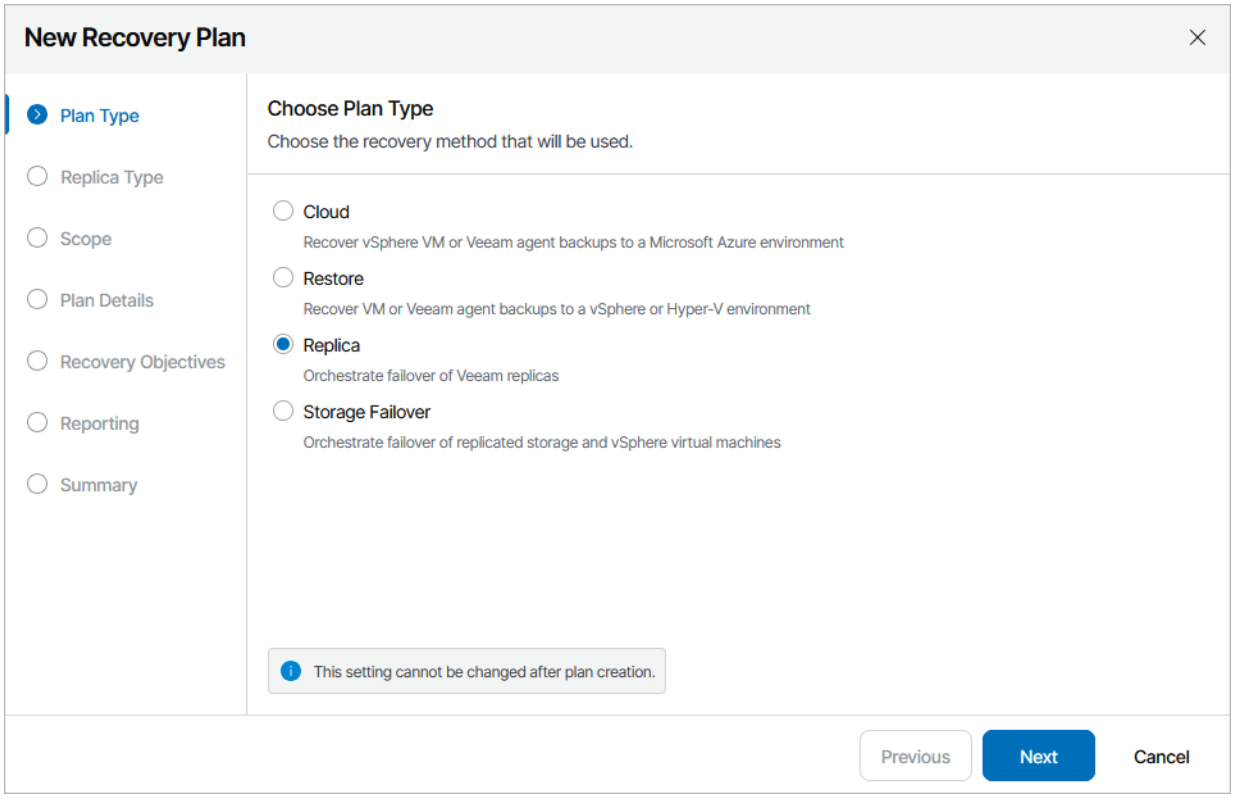

Step 2. Choose Replica Type

At the **Replica Type** step of the wizard, select the **Standard Replica** option.

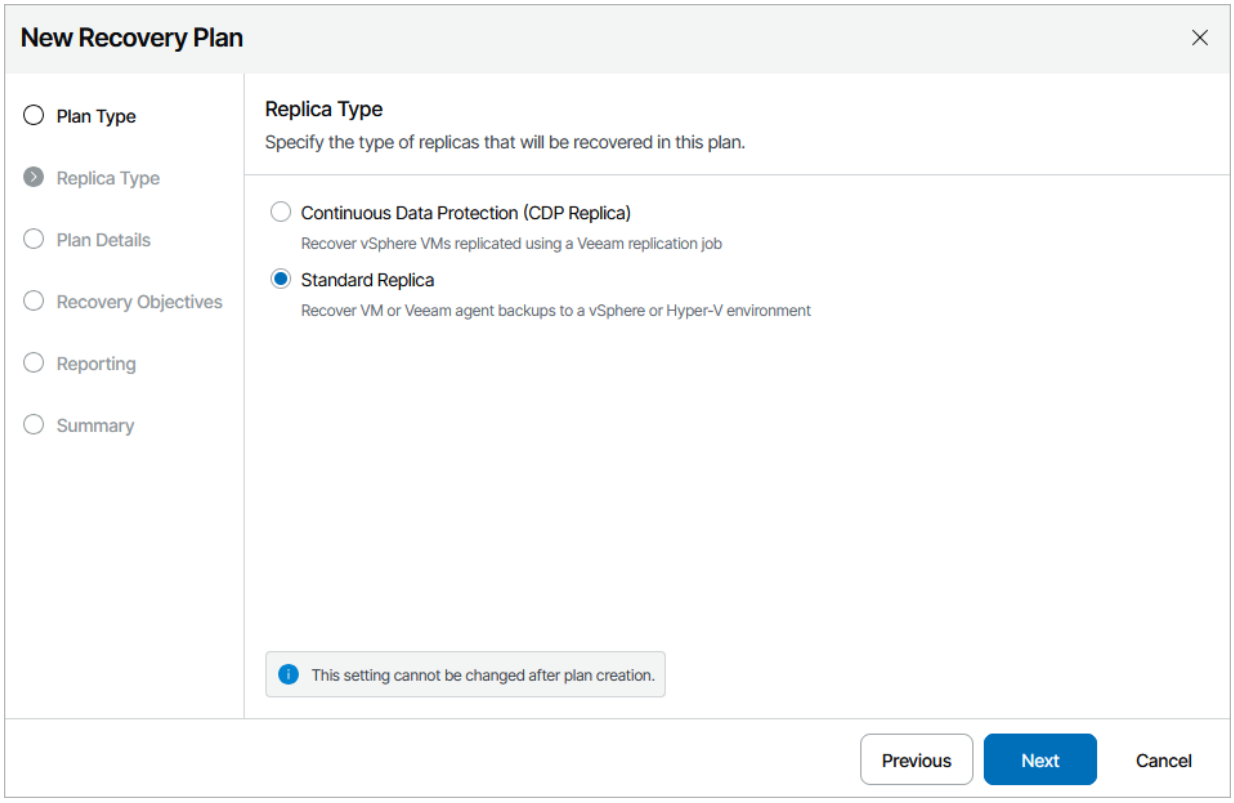

Step 3. Choose Plan Scope

At the **Scope** step of the wizard, select a scope for which you want to create the plan.

For a scope to be displayed in the list, it must be created and customized as described in section Managing Scopes.

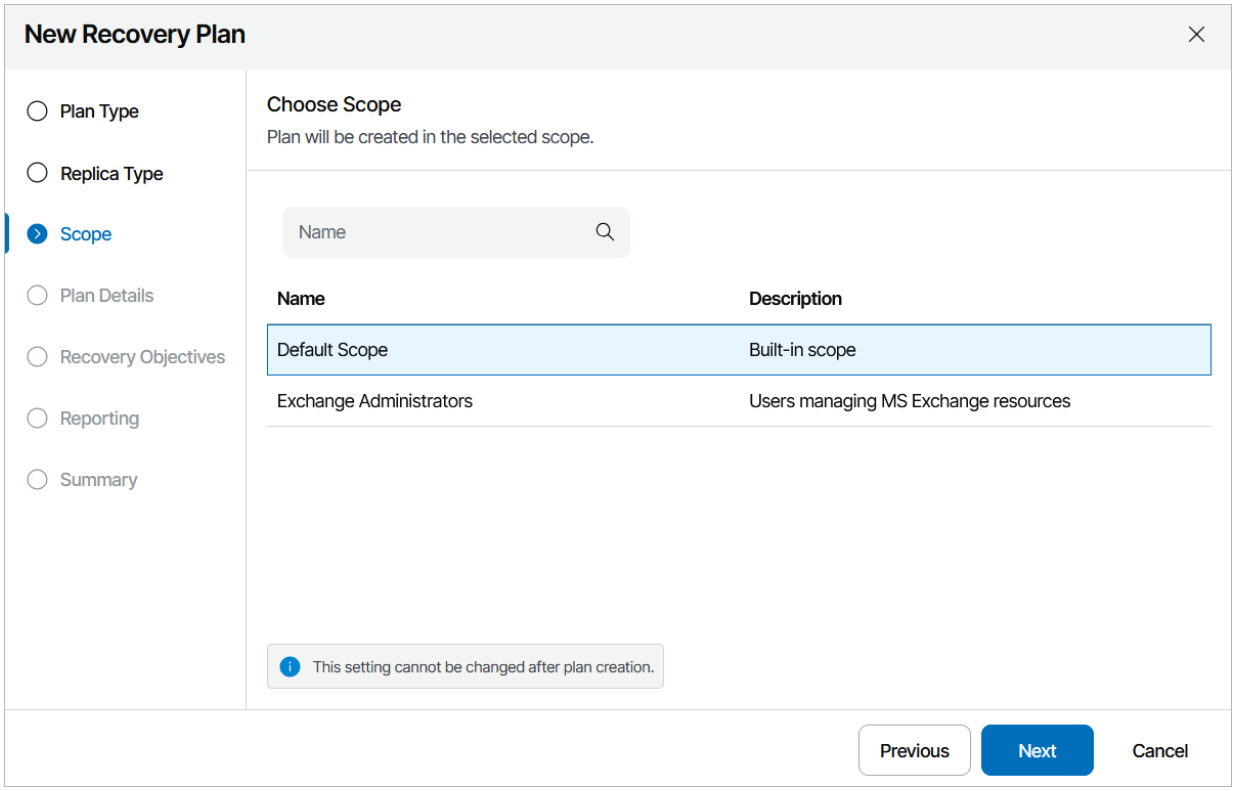

Step 4. Specify Plan Name and Description

At the **Details** step of the wizard, use the **Plan Name** and **Description** fields to enter a name for the new plan and to provide a description for future reference. The maximum length of the plan name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

You can also provide a contact name, email and telephone number of a person responsible for the plan.

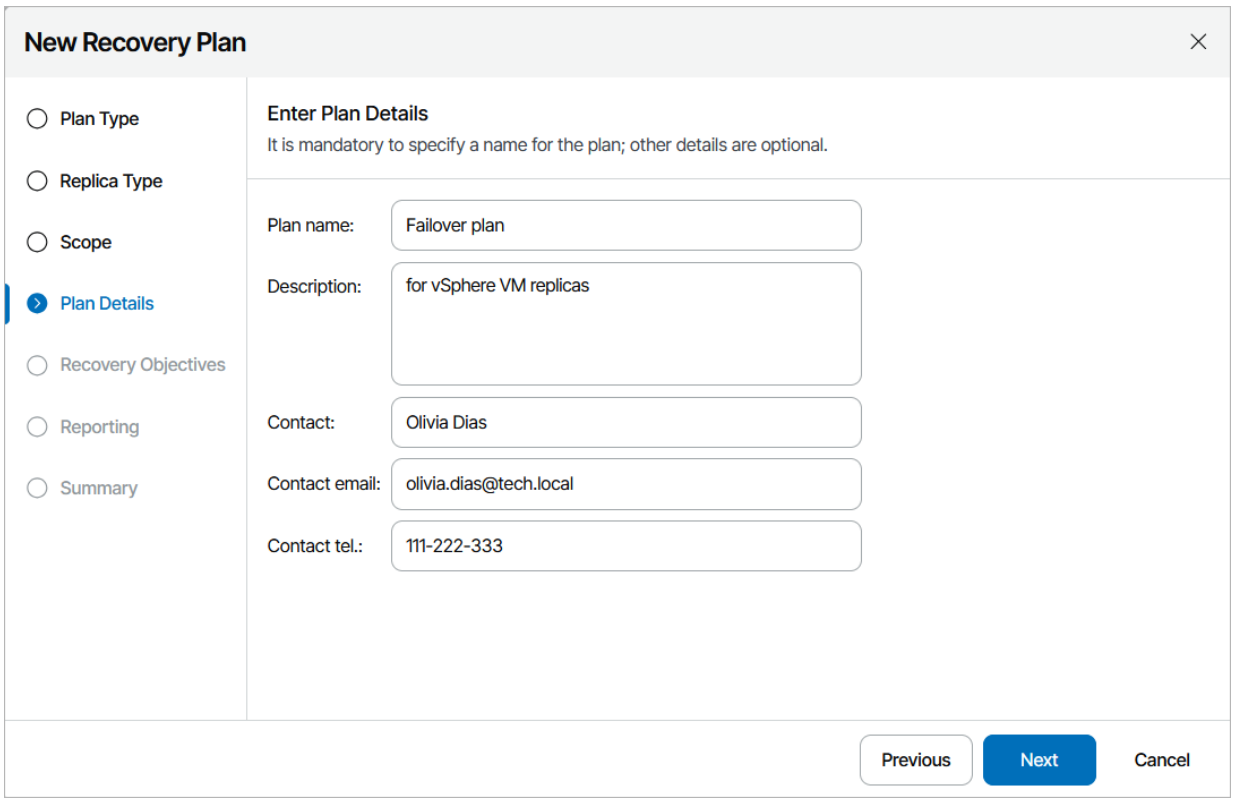

Step 5. Specify Target RTO and RPO

At the **Recovery Objectives** step of the wizard, define your Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for the plan:

- The **RPO** defines the maximum acceptable period of data loss.
- The **RTO** represents the amount of time it should take to recover from an incident.

RTO and RPO performance will be recorded in the Plan Readiness Check, Plan Execution, Plan Audit and DataLab Test reports, and you will be able to track the achieved RTO and RPO objectives for each plan on the Home Page Dashboard.

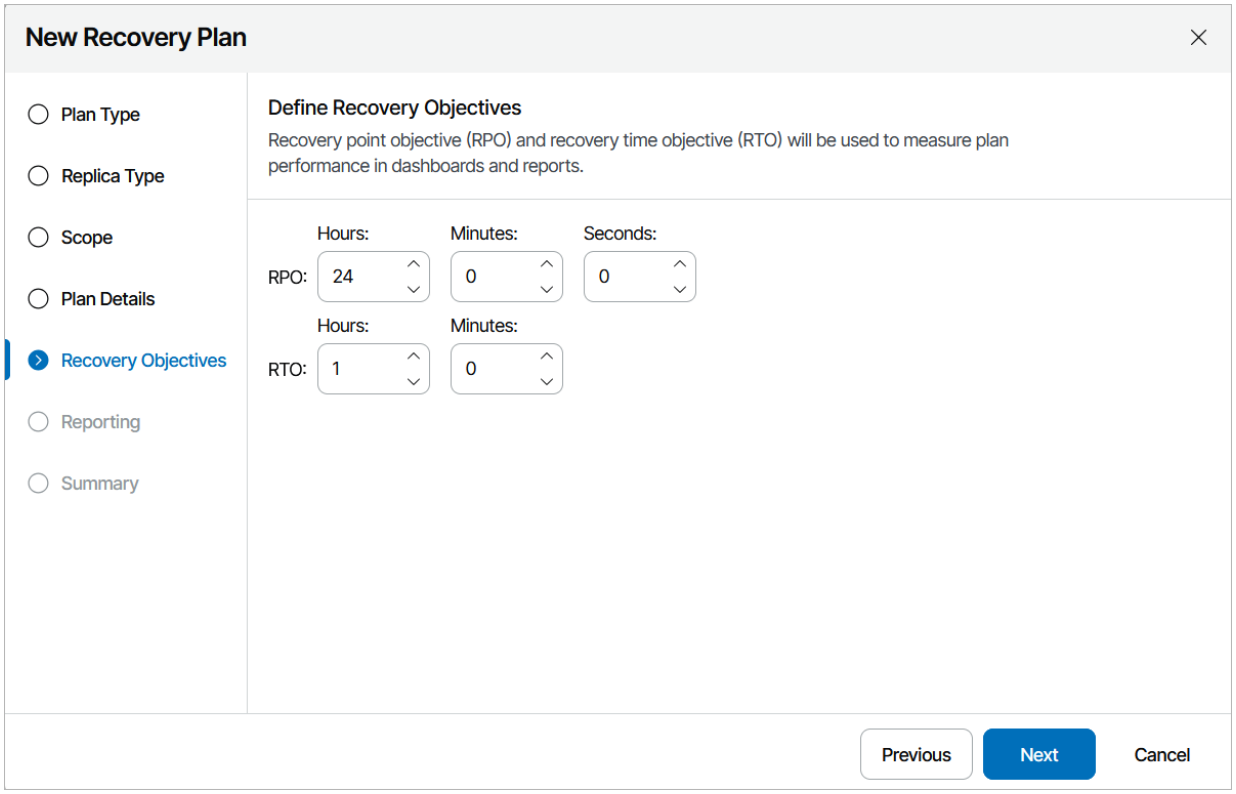

Step 6. Select Report Template

At the **Reporting** step of the wizard, select a document template that will be used as the cover page for all Orchestrator reports.

For a custom document template to be displayed in the list, it must be created and customized as described in section Managing Templates.

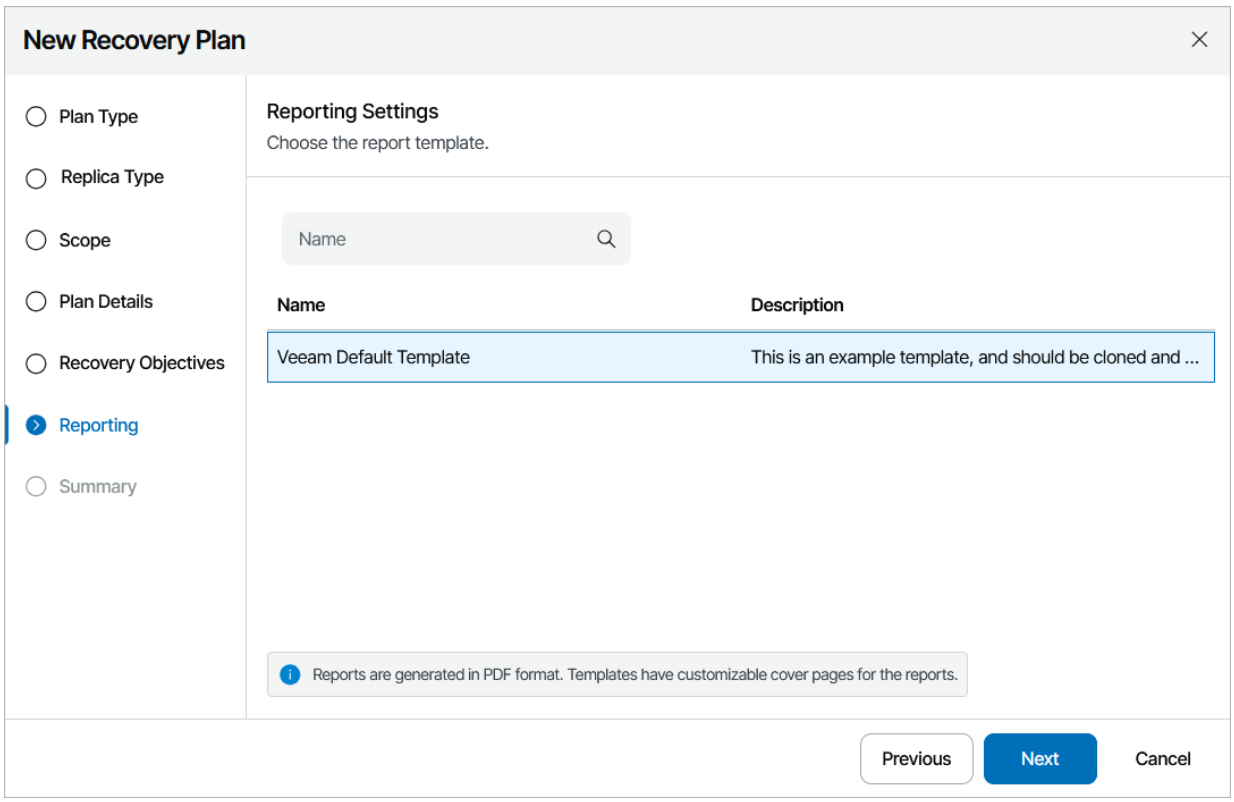

Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.

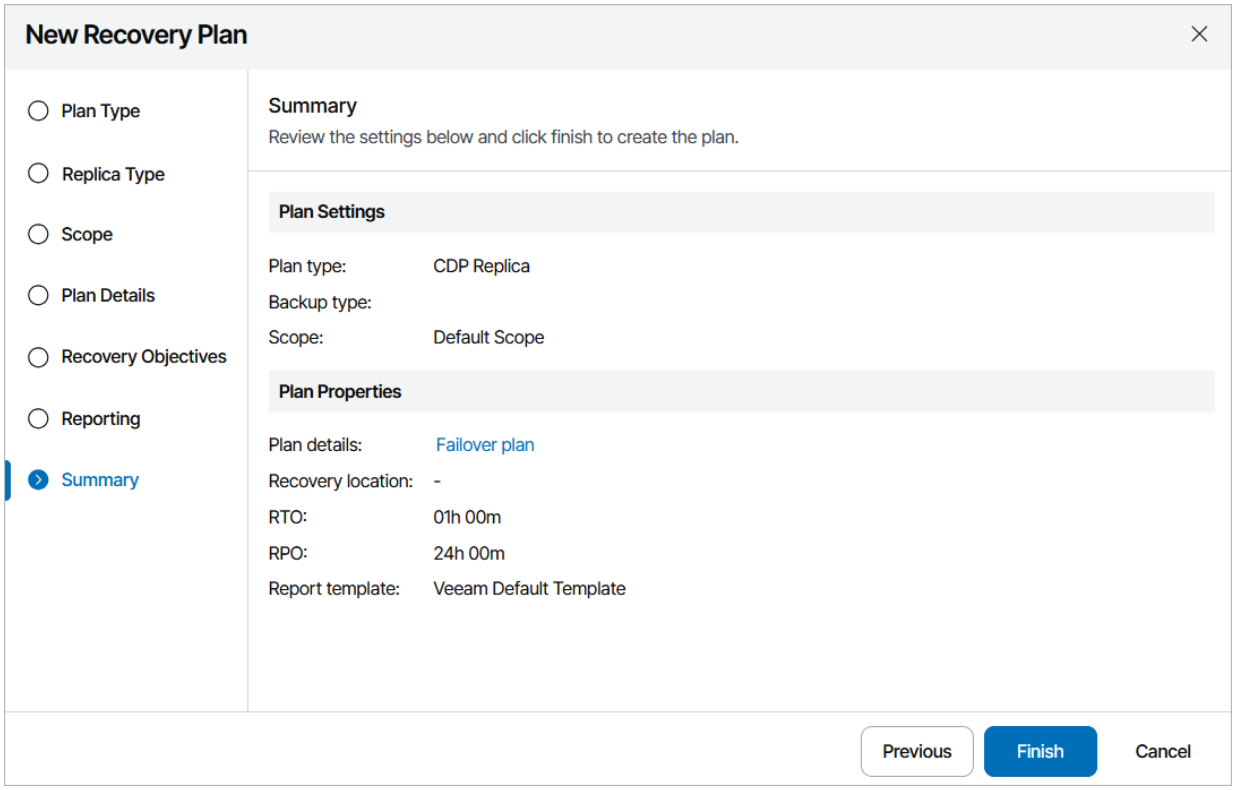

Editing Replica Plans

If you want to specify granular settings not provided in the New Recovery Plan wizard, the Orchestrator UI allows you to customize replica plans and configure the settings for groups, recovered VMs, plan steps and step parameters.

The procedures to edit replica, CDP replica, restore, storage and cloud plans are almost identical. For more information, see Editing Recovery Plans.

Testing Replica Plans

You can start on-demand plan testing and configure test scheduling for any replica plan. There is almost no difference between the procedures performed for replica, CDP replica, restore and storage plans. For more information, see Testing Recovery Plans.

Scanning Replica Plans

You can start on-demand plan scanning for malware and configure scan scheduling for any replica plan. There is almost no difference between the procedures performed for replica, CDP replica, restore and cloud plans. For more information, see Scanning Recovery Plans.

Running and Scheduling Replica Plans

To run a replica plan, it must be *ENABLED*. To enable a plan:

1. Navigate to **Recovery Plans**.
2. Select the plan.
3. From the **Manage** menu, select **Properties**.

   OR-

   Right-click the plan name and select **Manage** > **Properties**.
4. Set the **Availability** toggle to *Enabled*.
5. Click **Save**.

If you do not enable a plan before you run it, the Run Plan wizard will force you to do that as soon as you try running the plan.

> **NOTES**
>
> 1. An Orchestrator Administrator or Plan Operator can force-enable a plan in the **Run Plan** wizard. However, a Plan Operator will not be able to run a disabled replica plan.
>
>   For more information on roles that can be assigned to users and user groups working with the Orchestrator UI, see Managing User Accounts.
> 2. For security purposes, all 'real-world' actions associated with replica plans (such as failover and failback) require password confirmation.

Scheduling Failover

You can schedule a time for a replica plan to execute. Only the failover process can be scheduled — all other operations (failback, undo failover and so on) must be performed manually in the Orchestrator UI.

To schedule a replica plan:

1. Navigate to **Recovery Plans**.
2. Select the plan. From the **Manage** menu, select **Schedule**.

   -OR-

   Right-click the plan name and select **Manage** > **Schedule**.
3. In the **Scheduled Tasks** window, do the following:

   a. Set the **Schedule plan execution** toggle to *Enabled*.

   b. Click the **Configure schedule** link and choose whether you want to run the plan on schedule or after any other plan:

   - If you want to run the plan at a specific time, click the **Schedule** icon in the **Run on** field, set the desired date and time, and click **Apply**.
   - If you want to run the plan after another plan, select the **Schedule after plan** check box and click **Choose plan**. Then, in the **Select Plan** window, select the necessary plan and click **Apply**.

     For a plan to be displayed in the list of available plans, it must be *ENABLED* as described in section Running and Scheduling Replica Plans.

c. Set the **Malware actions** toggle to *Enabled* if you want to check restore points created for machines included in the plan for malware flags.

   For more information on how Orchestrator performs malware scan, see the Veeam Recovery Orchestrator User Guide, section Overview.

d. Review the configuration information and click **Save**.

> **TIP**
>
> You can also scan a recovery plan for possible malware without scheduling the plan execution. To do that, follow the instructions provided in section Scanning Recovery Plans.

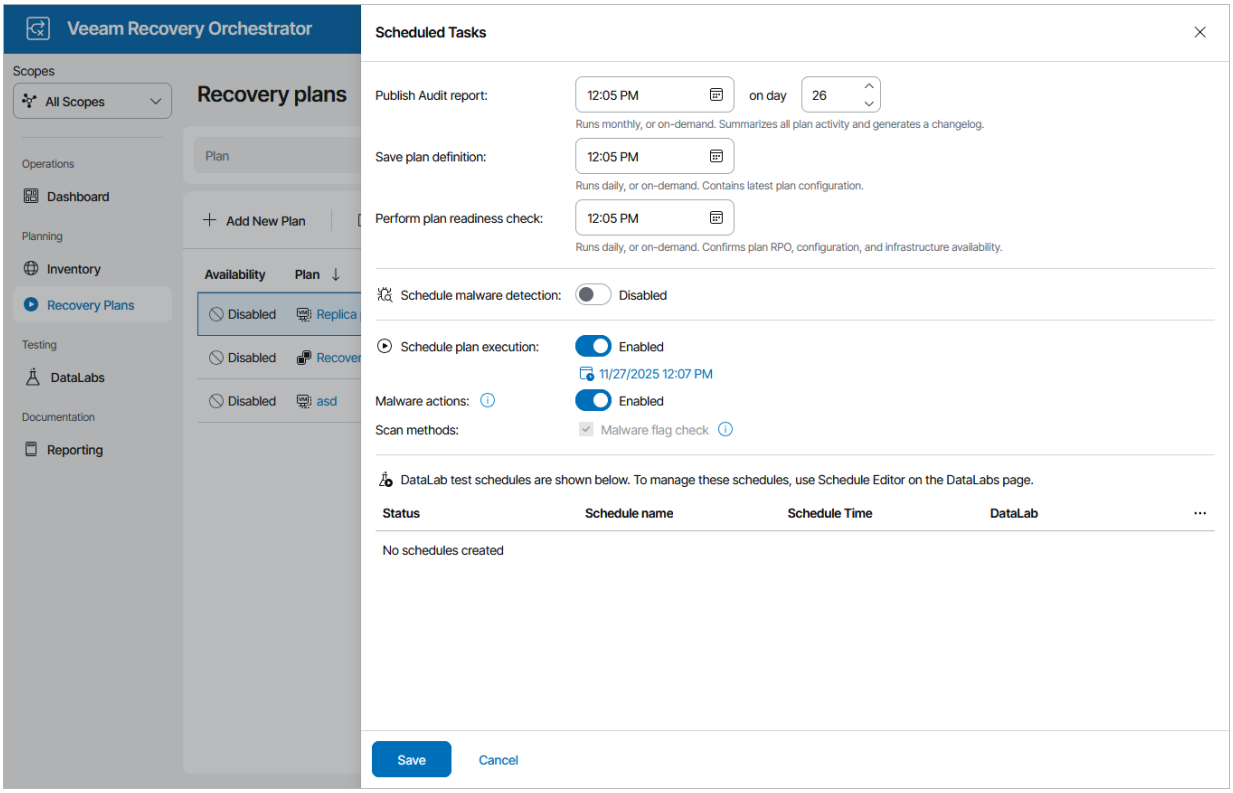

> **TIP**
>
> You can disable a configured schedule if you no longer need it. To do that, set the **Schedule plan execution** toggle to *Disabled* in the **Scheduled Tasks** window.

Running Failover

The **Run** action causes VMs in a plan to fail over to their replicas. For more information on the replica failover process, see the Veeam Backup & Replication User Guide, section Failover.

To run a replica plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Run**.

3. In the **Run Plan** window, do the following:

   a. For security purposes, retype your password and click **Next**.

      You must also select the **Force-enable the plan** check box if you have not enabled the plan yet.

   b. In the **Timestamp** field, choose a restore point that will be used to recover VM replicas.

   c. In the **Malware actions** field, choose whether you want to check restore points created for machines included in the plan for malware flags.

      For more information on how Orchestrator performs malware scan, see the Veeam Recovery Orchestrator User Guide, section Overview.

   d. Review configuration information and click **Run**.

> **TIP**
>
> You can also scan a recovery plan for possible malware without running the plan. To do that, follow the instructions provided in section Scanning Recovery Plans.

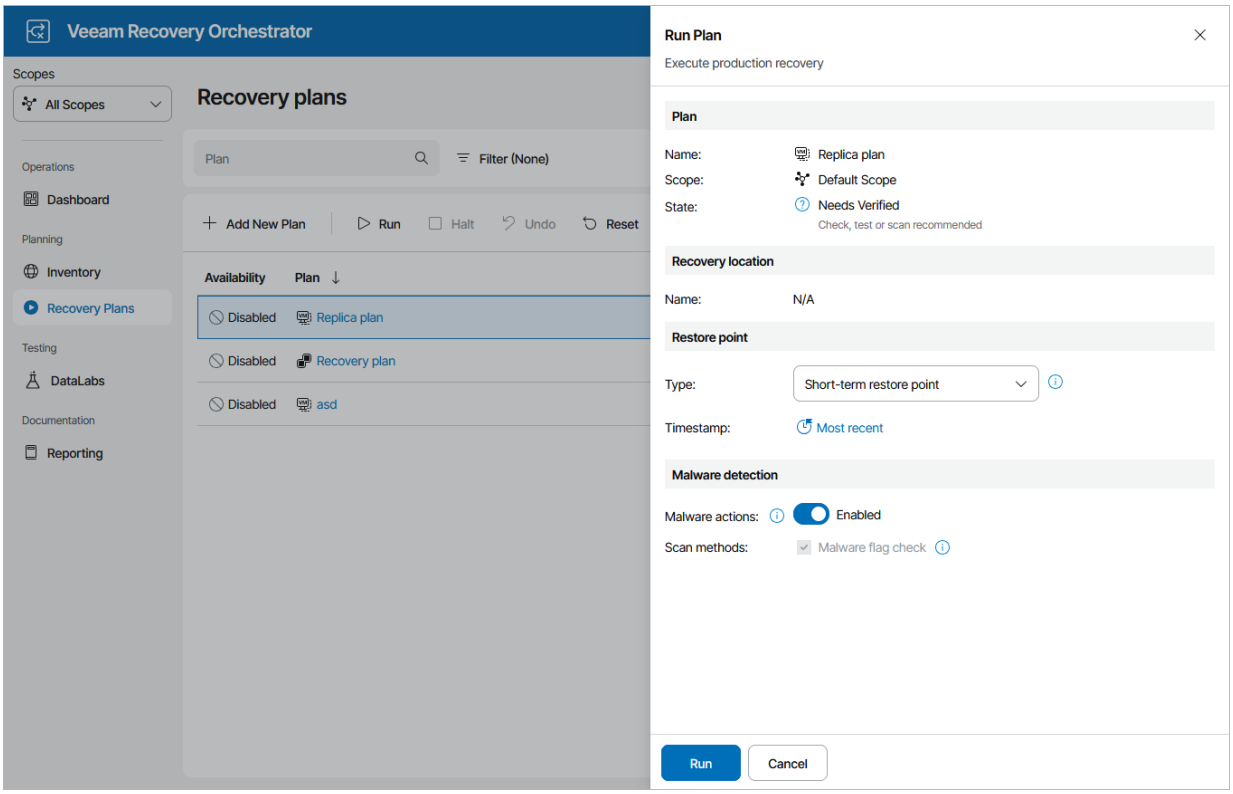

The plan goal is to reach the *FAILOVER* state. If any critical error is encountered, the plan will stop with the *HALTED* state. To learn how to work with *HALTED* replica plans, see Managing Halted Plans.

Halting Failover

The **Halt** action interrupts plan execution. Any currently executing steps will be completed, then the plan will enter the *HALTED* state. To learn how to work with *HALTED* replica plans, see Managing Halted Plans.

To stop a running replica plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Halt**.

3. In the **Halt Plan Execution** window, do the following:
    a. For security purposes, retype your password and click **Next**.
    b. Review configuration information and click **Halt**.

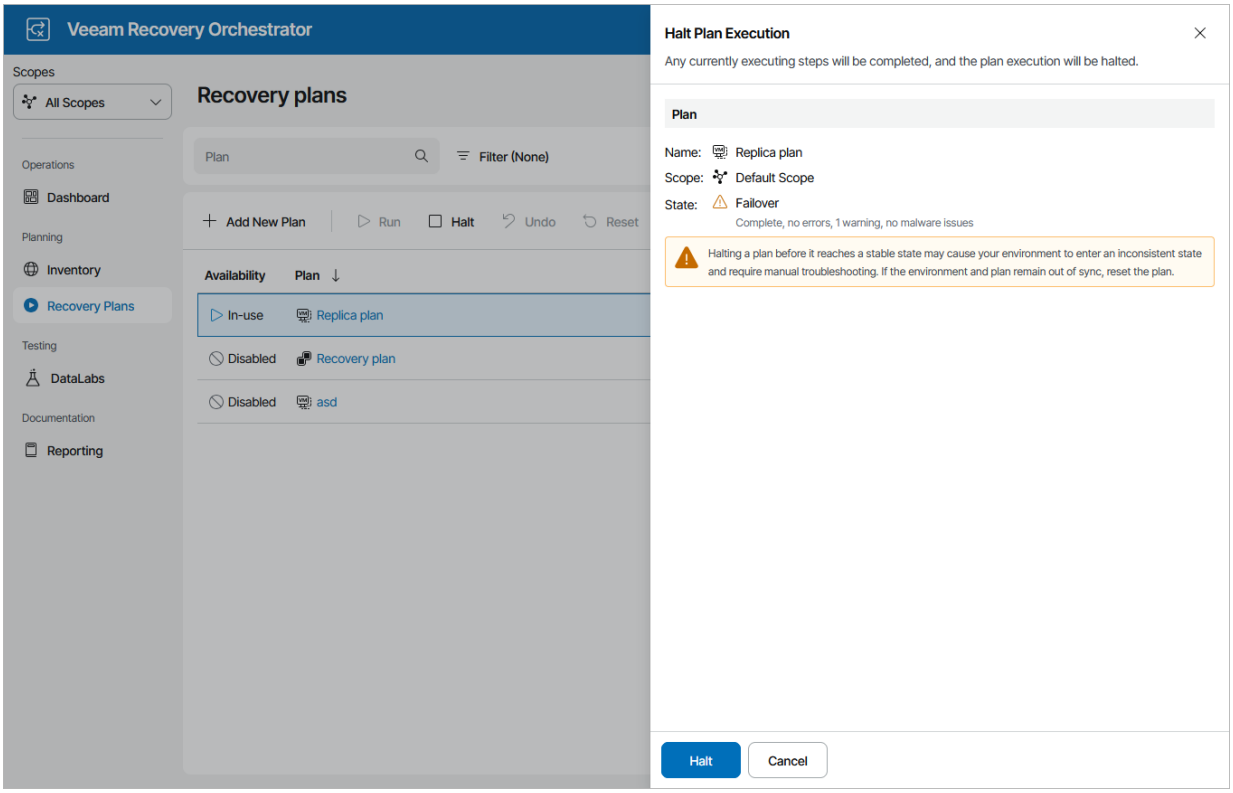

Finalizing Failover

Orchestrator provides you with a number of options to finalize failover to VM replicas:

- Perform permanent failover
- Perform failback
- Commit failback

For more information on the available options, see the Veeam Recovery Orchestrator User Guide, section Failover and Failback.

Running Permanent Failover

To perform permanent failover for a plan in the *FAILOVER* state:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Run**.
3. In the **Progress Plan** window, do the following:
    a. For security purposes, retype your password and click **Next**.
    b. In the **Recovery actions** section, select the **Permanent failover** option.
    c. Review configuration information and click **Run**.

> **NOTE**
>
> Failback will no longer be an option once the permanent failover process is complete.

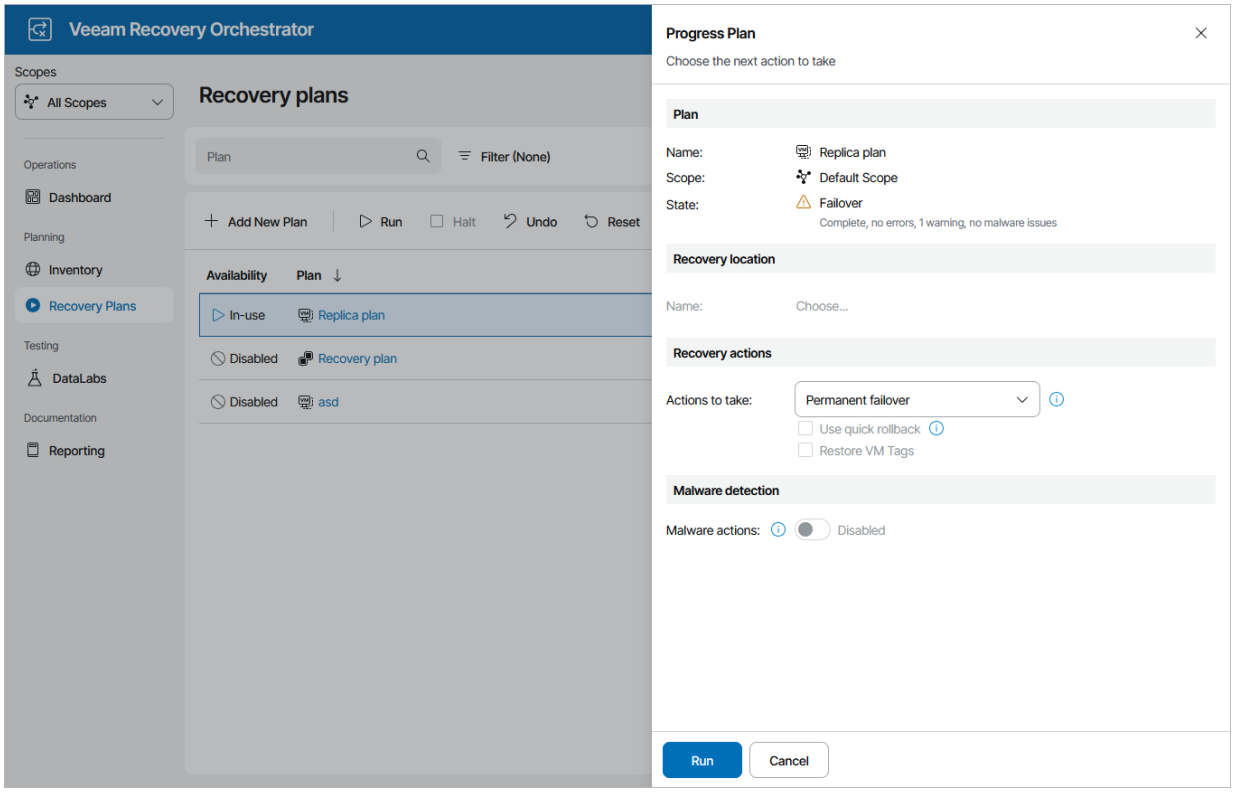

Running Failback

To perform failback for a plan in the *FAILOVER* state:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Run**.
3. In the **Progress Plan** window, do the following:
   a. For security purposes, retype your password and click **Next**.
   b. In the **Recovery Location** section, select a location to which VMs will be recovered.

      For a recovery location to be displayed in the list of available locations, it must be created and added to the list of inventory items available for the scope, as described in section Managing Recovery Locations.

> **NOTE**
>
> Orchestrator will perform failback using all settings configured for the location — except Instant VM Recovery and backup copy preference. These settings are not applicable to failback operations.

      If you want to fail back to a new recovery location and the selected location includes multiple hosts, datastores and networks, Orchestrator will use the round-robin algorithm to recover VMs. For more information, see the Veeam Recovery Orchestrator User Guide, section How Orchestrator Places VMs During Failback.

   c. In the **Recovery actions** section, do the following:

i. From the **Actions to take** drop-down list, select the **Prepare to failback** option for Orchestrator to switch from VM replicas to the source VMs when you run the plan next time.

ii. [Applies only if you have selected the Original VM Location] Select the **Use quick rollback** check box if you want to instruct Orchestrator to synchronize changed data blocks only — this may help you speed up the failback process significantly.

For more information on the quick rollback process, see the Veeam Backup & Replication User Guide, section Quick Rollback.

d. In the **Malware detection** section, choose whether you want to check restore points created for the VMs included in the plan for malware flags.

For more information on how Orchestrator performs malware scan, see the Veeam Recovery Orchestrator User Guide, section Overview.

e. Review configuration information and click **Run**.

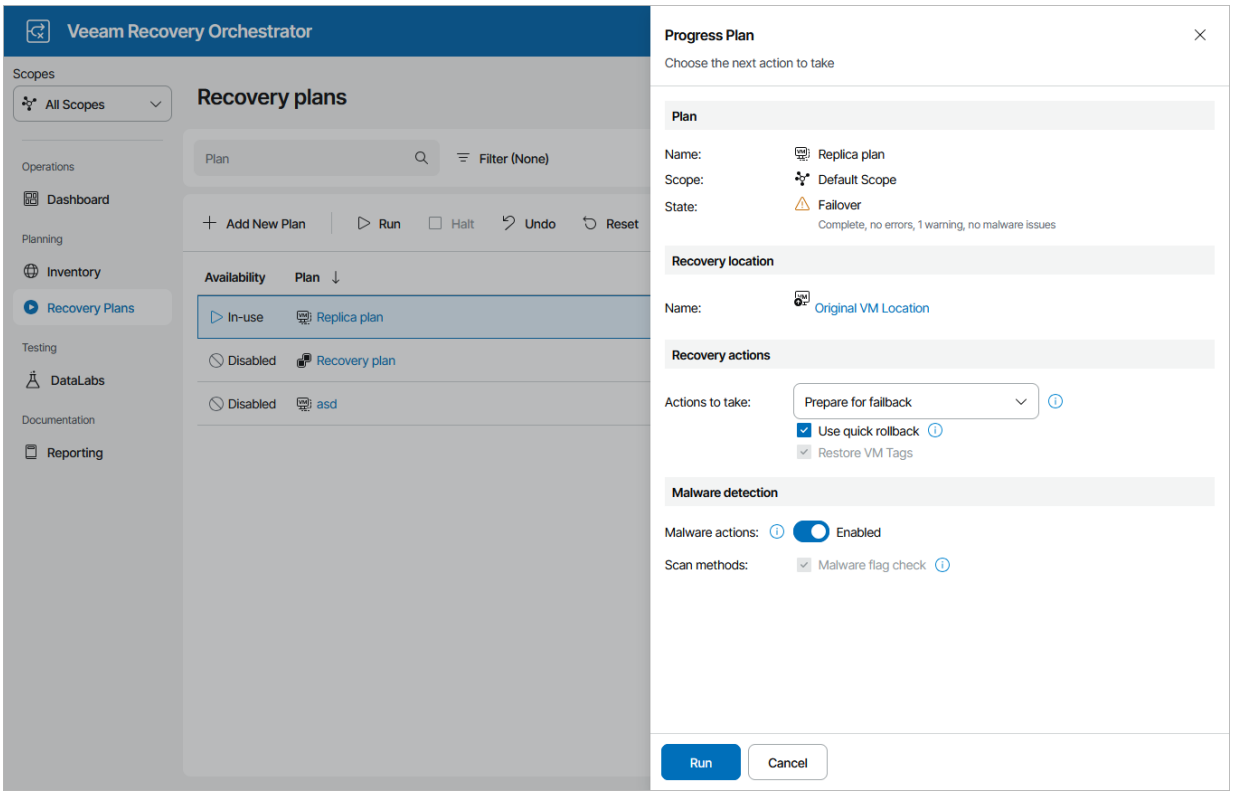

Committing Failback

To commit failback for a plan in the *FAILBACK* state:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Run**.
3. In the **Progress Plan** window, do the following:
   a. For security purposes, retype your password and click **Next**.
   b. In the **Recovery actions** section, select the **Failback now** option to switch from VM replicas to the source VMs immediately.

c. Review configuration information and click **Run**.

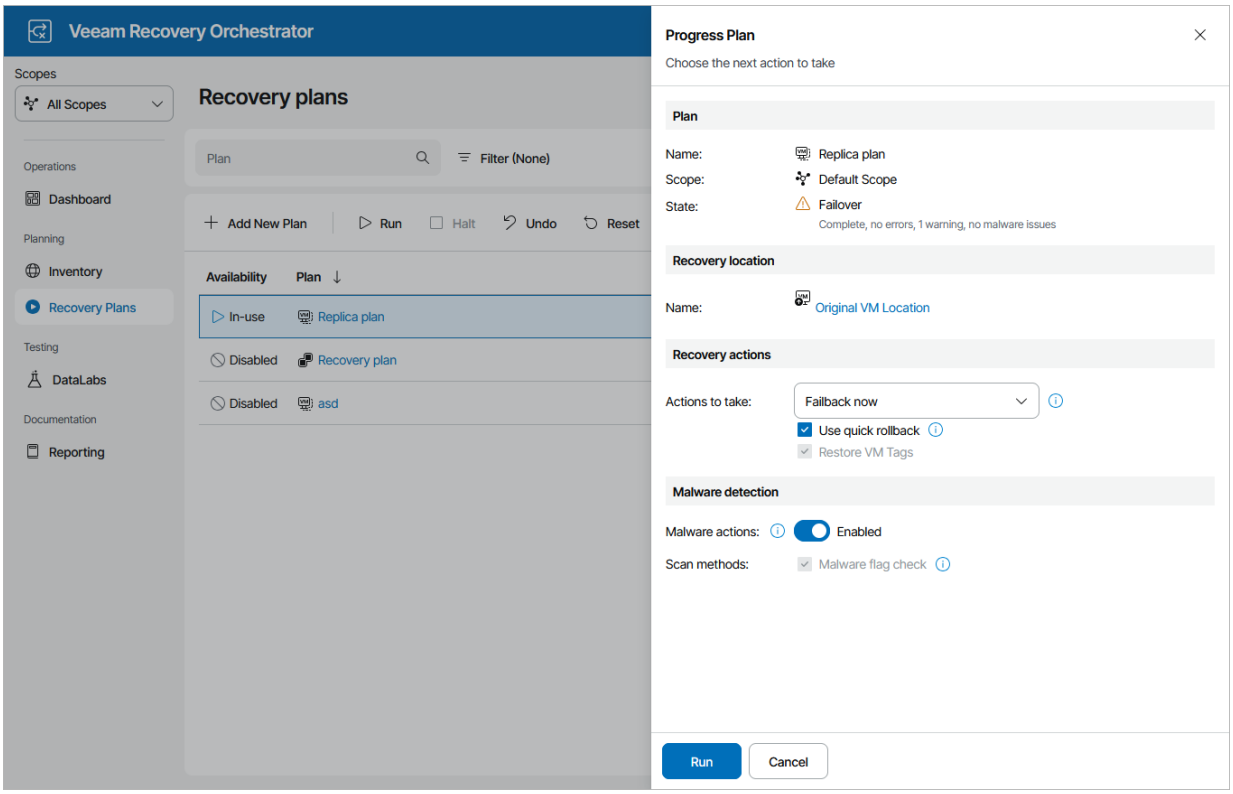

TIP

After the commit failback process completes, Orchestrator will leave the plan in the *IN-USE* mode. By design, this makes the results of the commit failback process accessible in the Orchestrator UI as long as required, and also prevents the plan from being modified by any automatic updates related to infrastructure changes.

If you want to perform any further actions with the plan (for example, to test the plan, to run readiness checks or to execute the plan again), reset the plan as described in section Resetting Replica Plans.

Undoing Failover

The **Undo Failover** action powers off VM replicas running on target hosts and rolls back to the VM state before failover. For more information on the undo failover operation, see the Veeam Backup & Replication User Guide, section Undo Failover.

To perform an undo operation for a plan in the *FAILOVER* state:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Undo**.
3. In the **Undo** window, do the following:
   a. For security purposes, retype your password and click **Next**.

b. Review configuration information and click **Undo**.

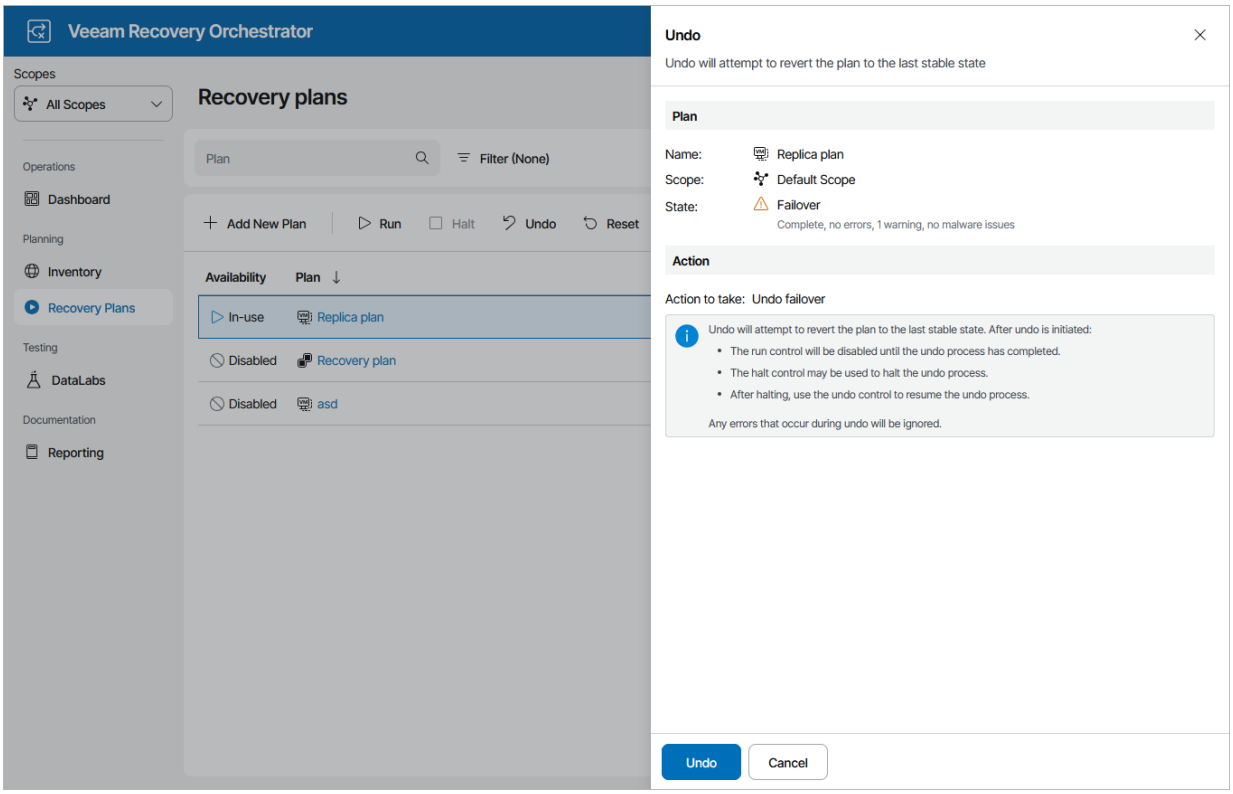

If the undo failover process encounters an error while being performed, it will not be halted automatically — the plan will proceed until the process completes. To terminate the undo failover process manually, use the **Halt** option to stop the currently running plan as described in section Halting Failover. To resume the undo failover process again, use the **Undo** option.

Undoing Failback

The **Undo Failback** action powers on VM replicas running on target hosts and switches from the production VMs back to the VM replicas — as a result, the plan acquires the *FAILOVER* state. For more information on the undo failback operation, see the Veeam Backup & Replication User Guide, section Undo Failback.

To perform an undo operation for a plan in the *PREPARE FOR FAILBACK* or *FAILBACK* state:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Undo**.
3. In the **Undo** window, do the following:
   a. For security purposes, retype your password and click **Next**.

b. Review configuration information and click **Undo**.

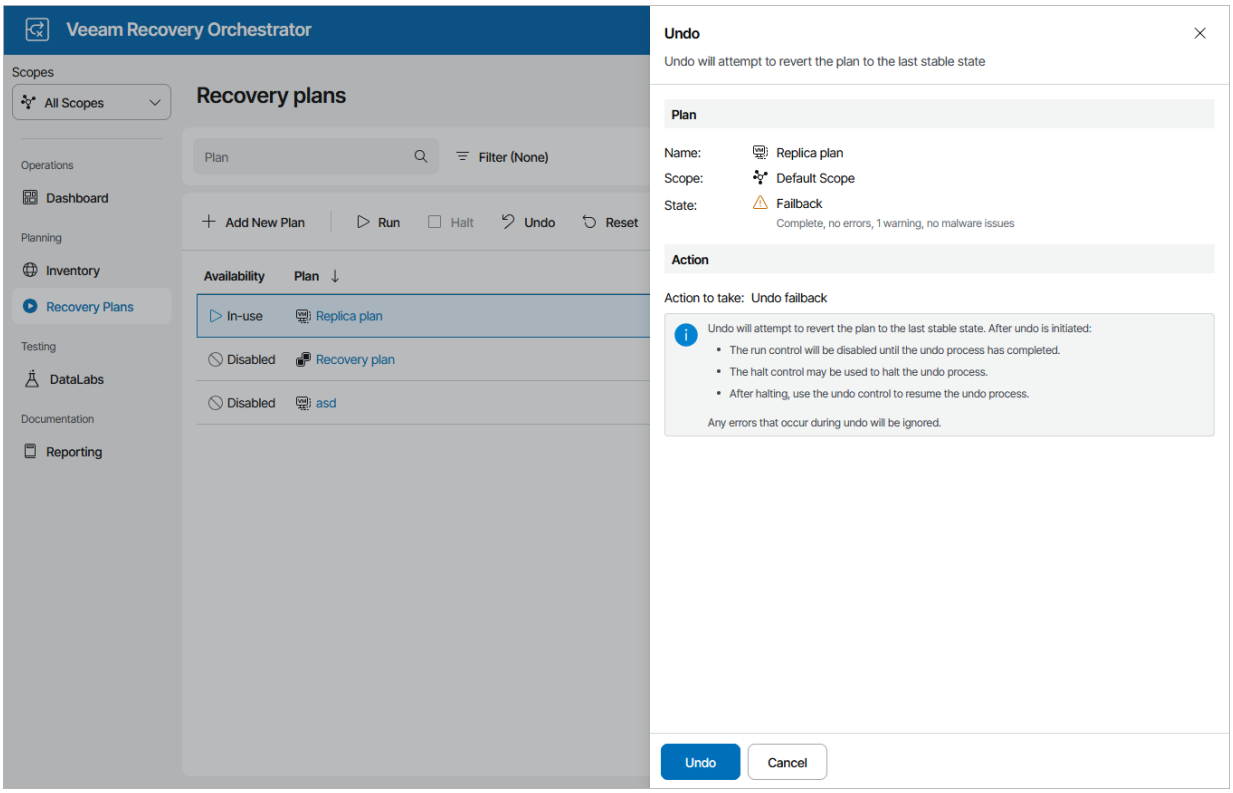

If the undo failback process encounters an error while being performed, it will not be halted automatically — the plan will proceed until the process completes. To terminate the undo failback process manually, use the **Halt** option to stop the currently running plan as described in section Halting Failover. To resume the undo failback process again, use the **Undo** option.

Resetting Replica Plans

If a replica plan becomes inconsistent with the virtual environment, you can reset the plan. This will return the plan to the *DISABLED* state, without making any changes to the external virtual infrastructure.

To reset a replica plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Reset**.
3. In the **Reset Plan** window, do the following:
   a. For security purposes, retype your password and click **Next**.
   b. Select the **Run a quick check after reset is complete** check box to run a readiness check after the reset.

c. Review configuration information and click **Reset**.

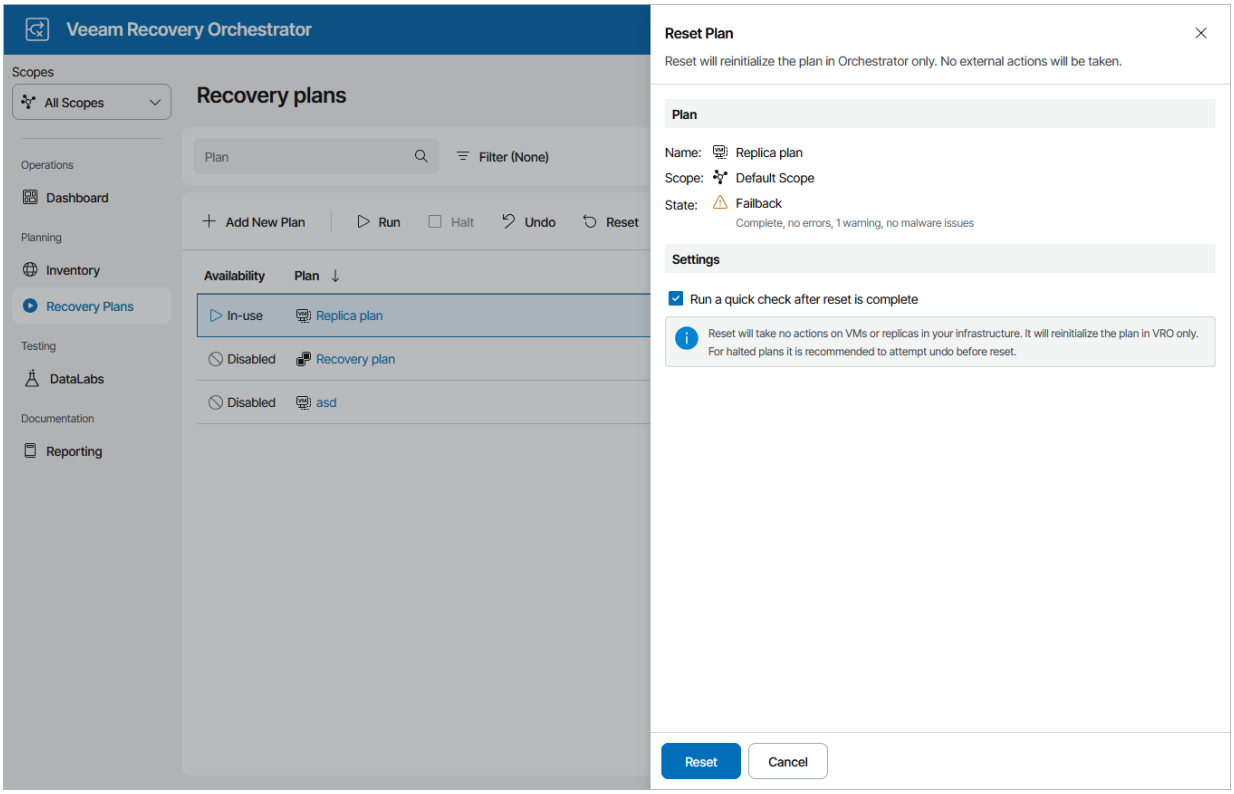

Halting Plan Testing

The **Halt** action interrupts plan testing. You may need to halt plan testing, for example, if you need to fix some environment-related issues and then proceed with testing later (in this case, VM replicas will still continue to run). Or you may need to stop the testing process completely, for example, if you no longer need to test the selected replica plan (in this case, VM replicas will be reverted to the latest restore point).

To halt testing of a replica plan:

1. Navigate to **Recovery Plans**.
2. Click the plan name to switch to the **Plan Details** page.
3. On the **Plan Details** page, select the plan and click **Halt**.
4. In the **Halt Plan Test** window, click **Halt Plan Test** to confirm the action.

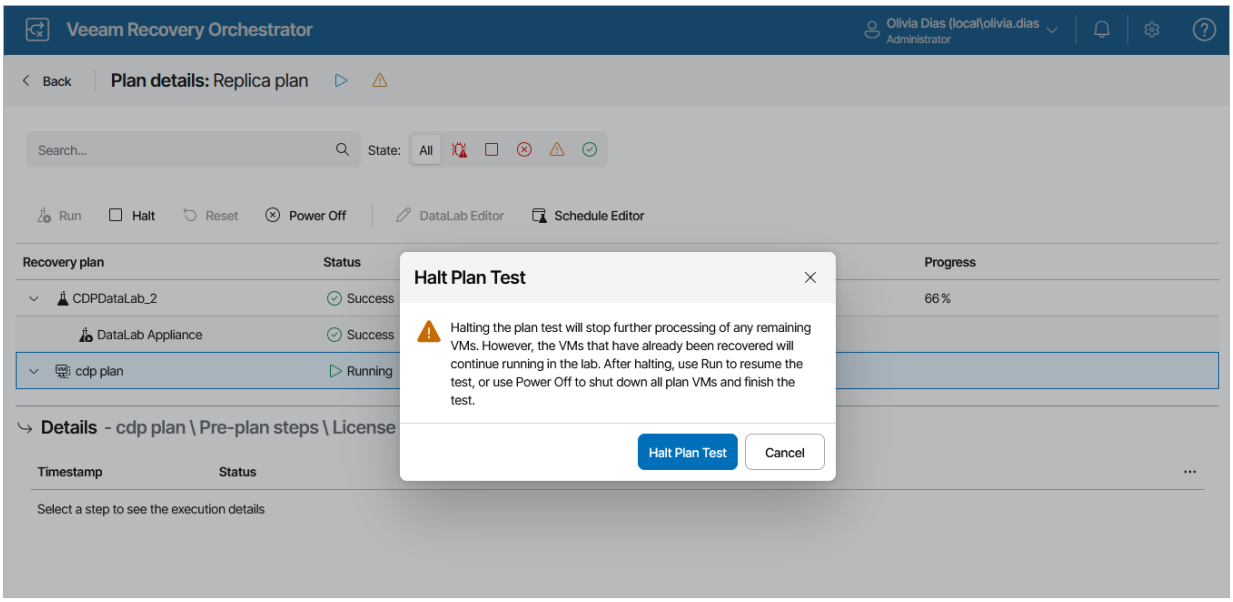

To cancel testing of a replica plan:

1. Navigate to **Recovery Plans**.
2. Click the plan name to switch to the **Plan Details** page.
3. On the **Plan Details** page, select the plan and click **Power Off**.
4. In the **Power Off Plan Testing** window, click **Power Off** to confirm the action.

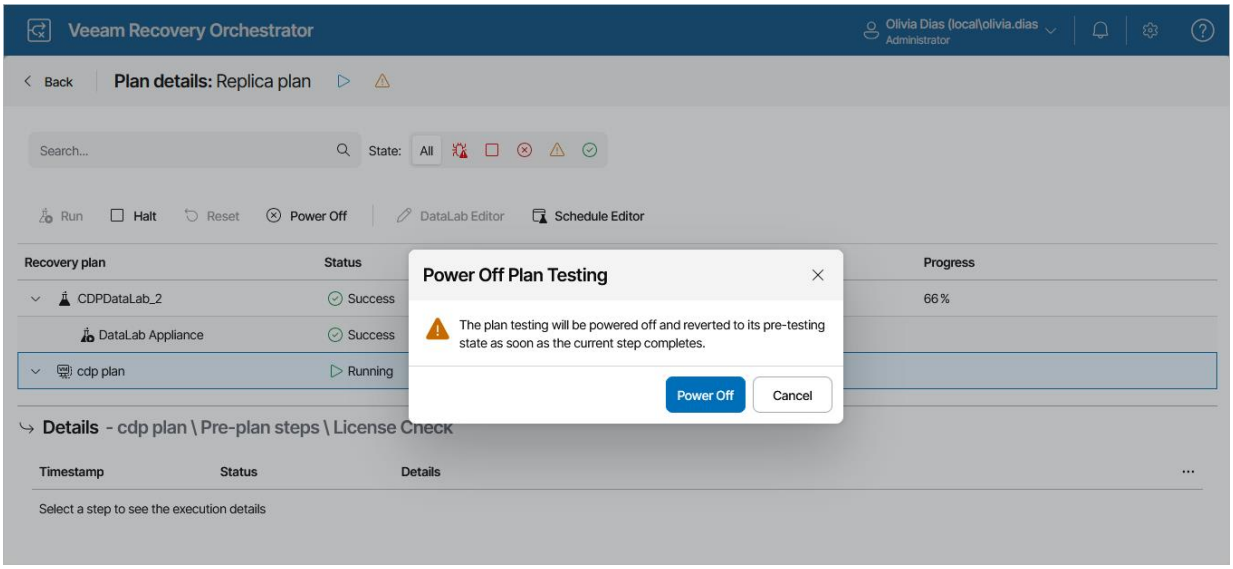

Managing Halted Replica Plans

If a critical step fails for a VM from a critical inventory group, the plan may enter the *HALTED* state. To troubleshoot reasons why a plan failed, use the **Plan Execution Report** generated as soon as the currently performed action completes. For more information on how to track plan performance history, see Viewing Plan Execution History.

After you eliminate the problem that caused the plan to become *HALTED*, you have the following options to resume the plan:

- Repeat the last failed step.
- Proceed to the next step.

Running Halted Replica Plans

To run a *HALTED* replica plan:

1. Navigate to **Recovery Plans**.
2. Select the halted plan and click **Run**.
3. In the **Continue Execution** window, do the following:
   a. For security purposes, retype your password and click **Next**.
   b. In the **Retry settings** section, select an option to resume plan execution.

      Choose whether you want to proceed with plan execution from the next plan step or to retry the failed step.

   c. Review configuration information and click **Continue**. The failover process will be started.

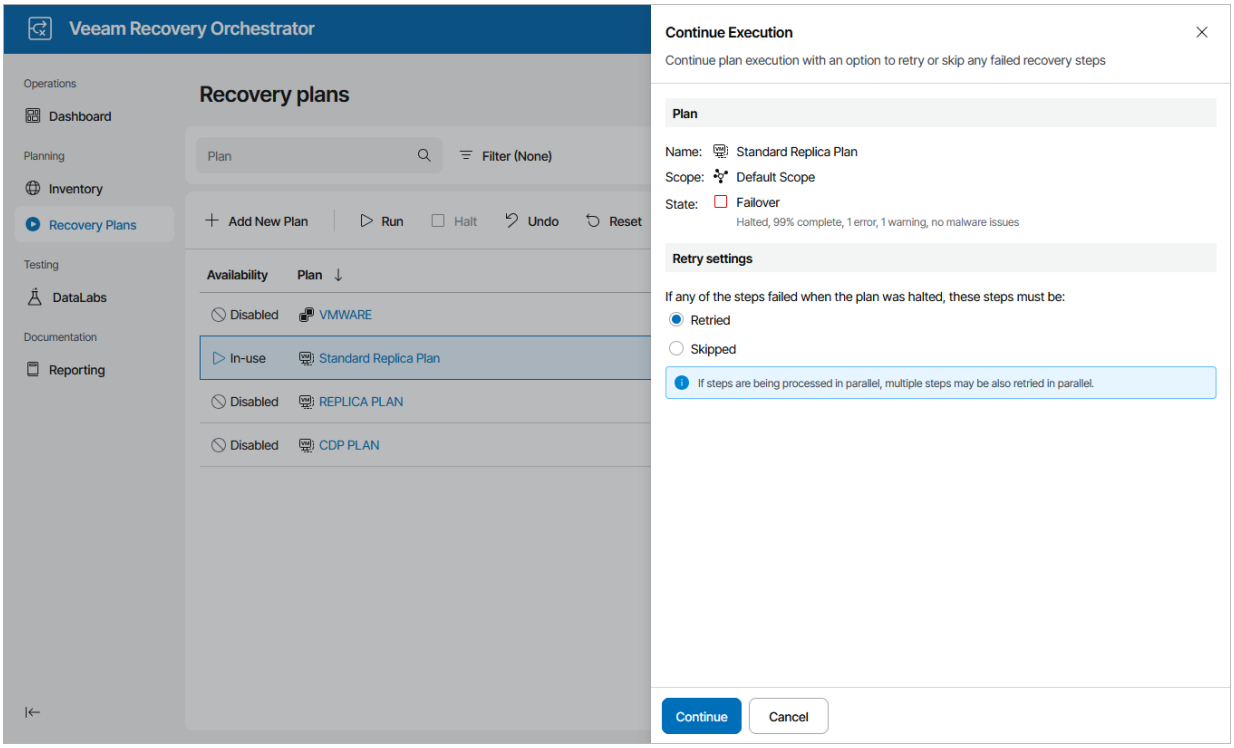

Resuming Plan Testing

To start the *HALTED* plan testing process:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Test**.

3. In the **Resume DataLab Test** window, choose whether you want to proceed with test execution from the next plan step or to retry the failed step, and then click **Continue**. The testing process will be started.

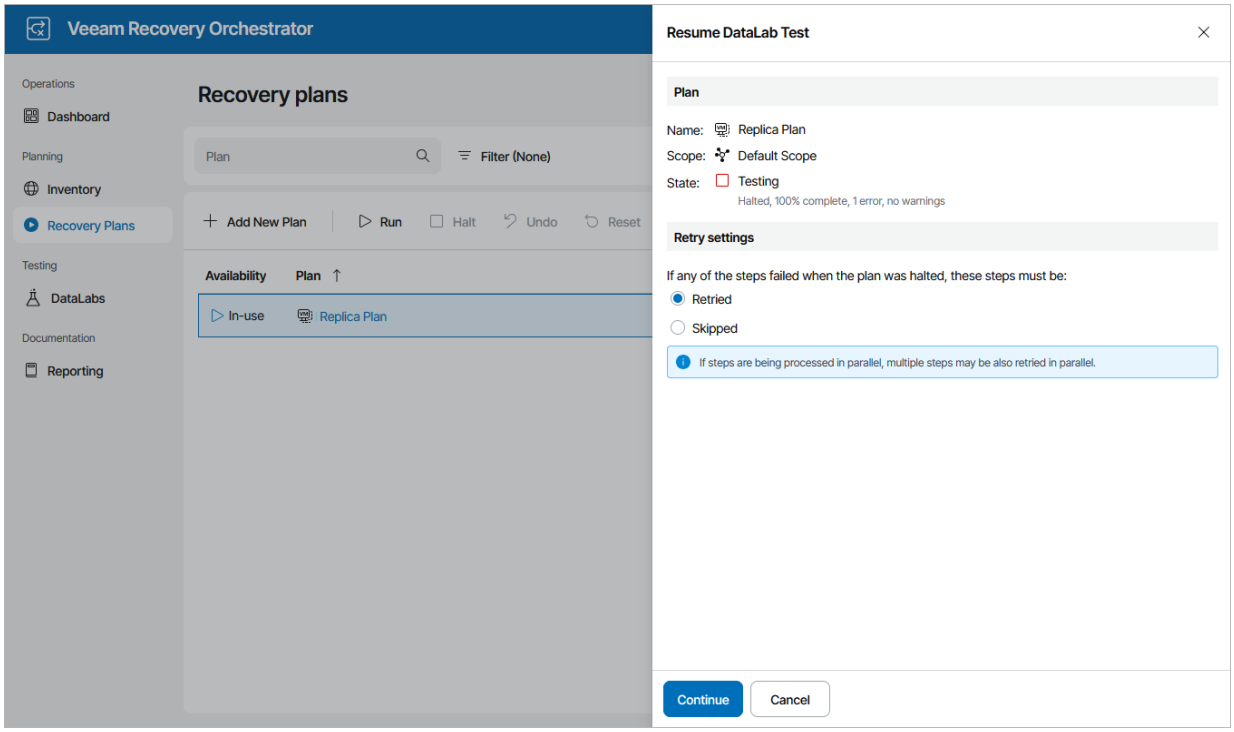

Undoing Halted Replica Plans

To perform an undo operation for a *HALTED* replica plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Undo**.
3. In the **Undo** window, do the following:
    a. For security purposes, retype your password and click **Next**.

b. Review configuration information and click **Undo**. The failover process will be started.

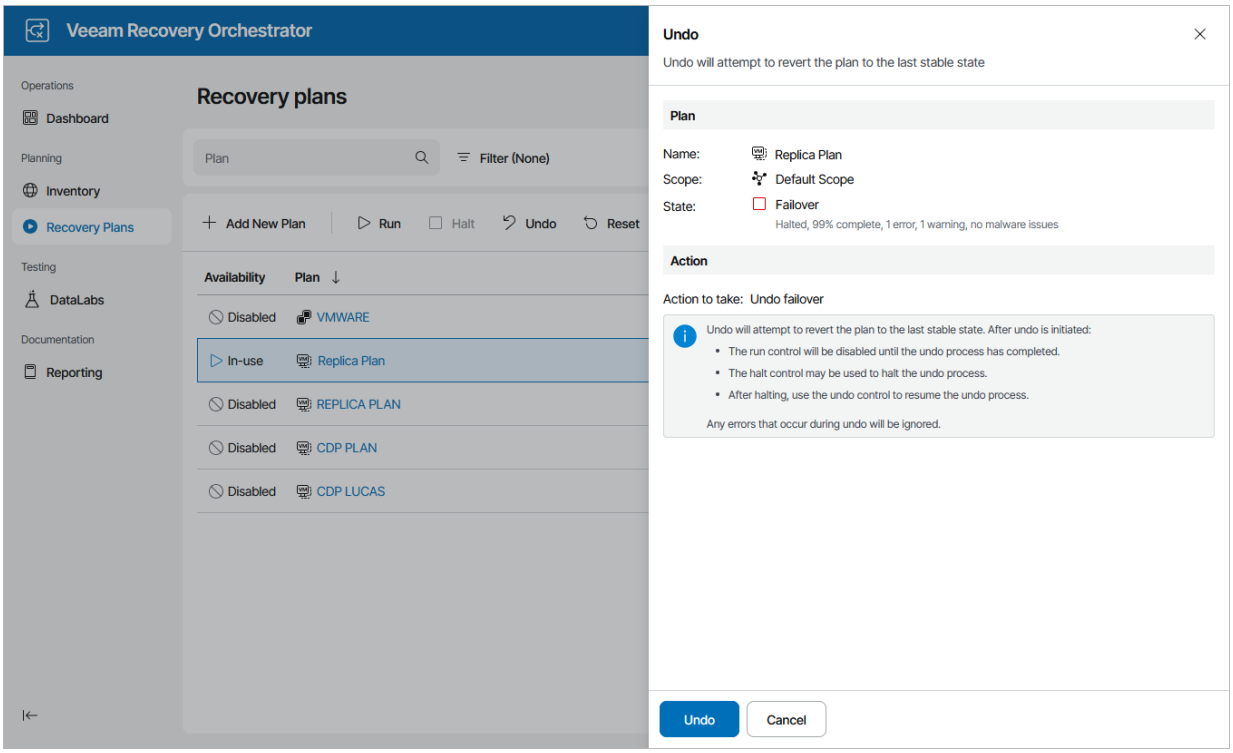

If a plan repeatedly enters the *HALTED* state due to misconfiguration or changes in the external environment, the only option left may be to RESET the plan.

Resetting Halted Replica Plans

To reset a *HALTED* replica plan, follow the instructions provided in section Resetting Replica Plans.

> **NOTE**
>
> When you reset a replica plan, Orchestrator returns it to the *DISABLED* state without making any changes to the external virtual infrastructure. You may need to deal with any infrastructure reconfiguration manually.

Working with CDP Replica Plans

The type of a recovery plan you create depends on whether you intend to use Orchestrator to switch to VM replicas, to restore machines from backups or backup copies, or to serve data from a destination (NetApp) or secondary (HPE) volume in case a disaster strikes.

If you want to recover vSphere VMs protected by Veeam CDP policies, create a CDP replica plan.

Creating CDP Replica Plans

To create a CDP replica plan:

1. Navigate to **Recovery Plans**.
2. Click **Add New Plan**.
3. Complete the **New Recovery Plan** wizard:
    a. Choose a type of the plan.
    b. Choose the necessary replica type.
    c. Choose a scope for the plan.
    d. Specify a plan name and description.
    e. Specify the target RTO and RPO.
    f. Select a template for plan reports.
    g. Finish working with the wizard.

Step 1. Choose Plan Type

At the **Plan Type** step of the wizard, select the **Replica** option.

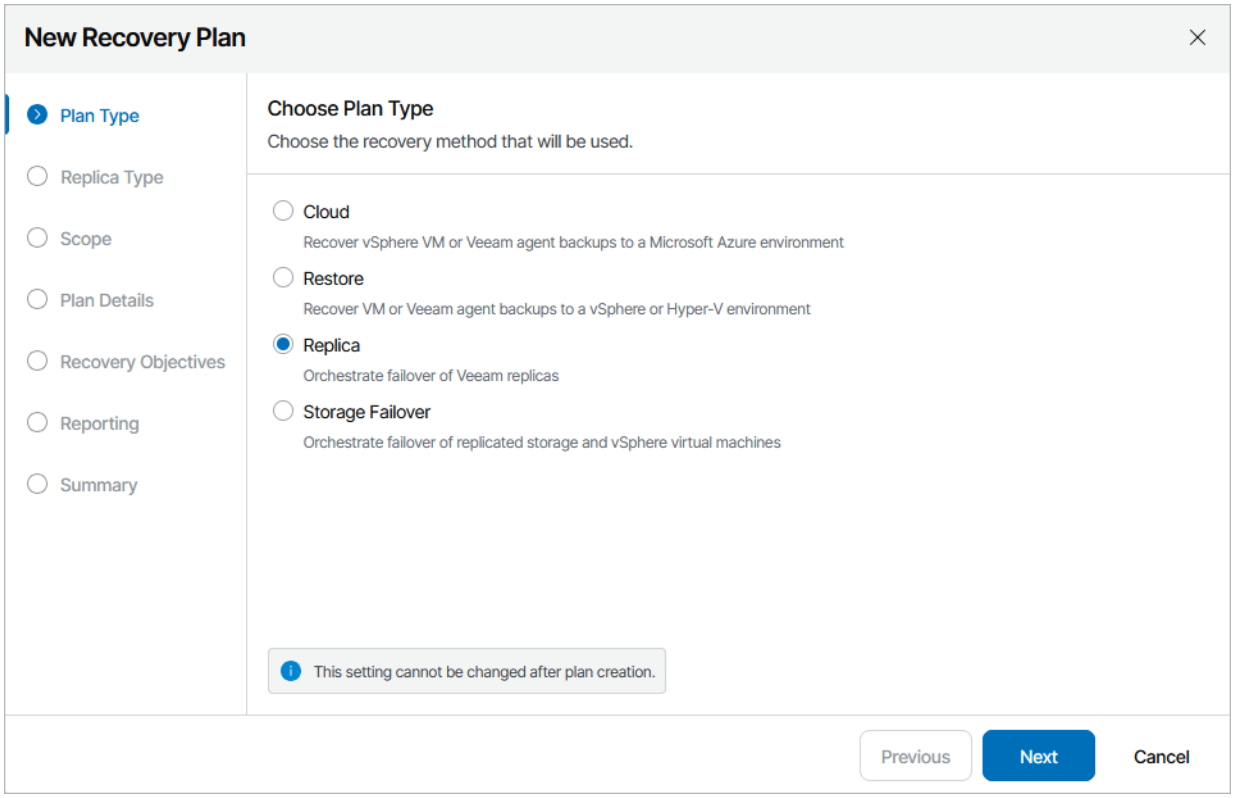

Step 2. Choose Replica Type

At the **Replica Type** step of the wizard, select the **Continuous Data Protection (CDP Replica)** option.

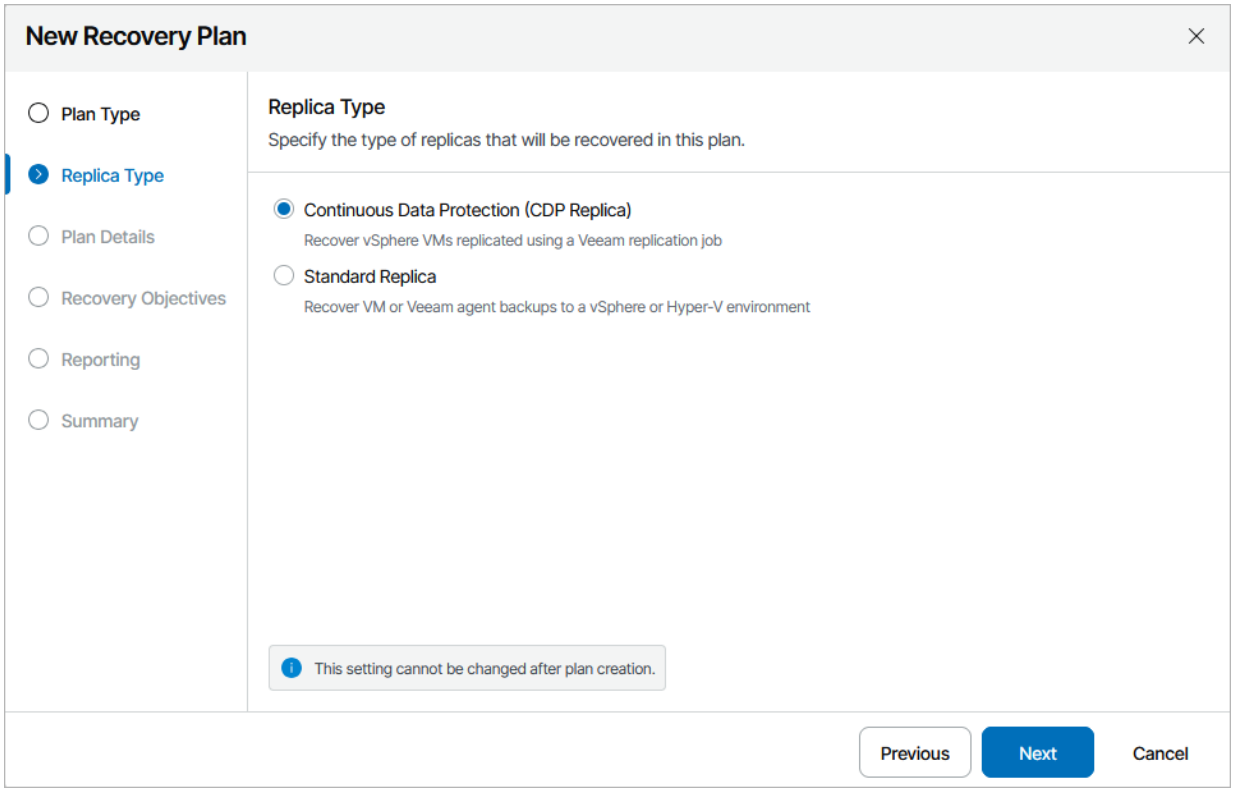

Step 3. Choose Plan Scope

At the **Scope** step of the wizard, select a scope for which you want to create the plan.

For a scope to be displayed in the **Available Scopes** list, it must be created and customized as described in section Managing Scopes.

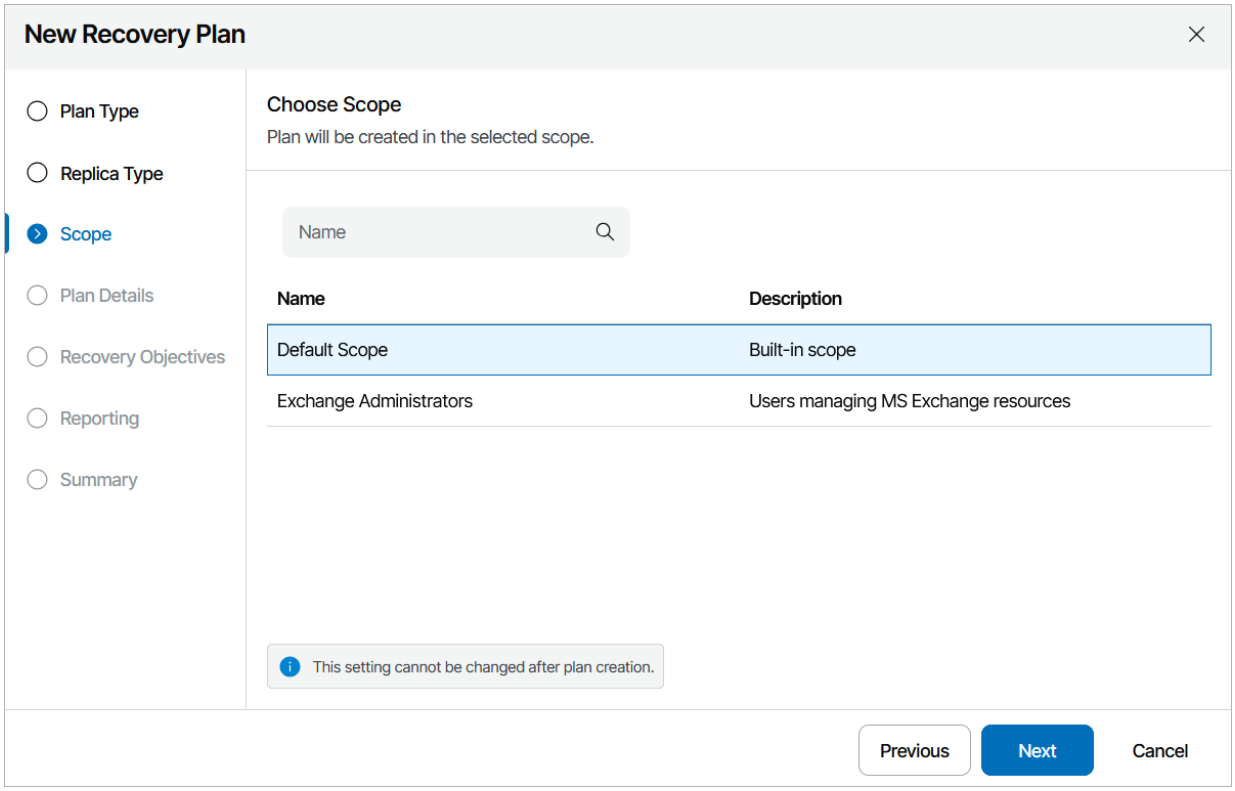

Step 4. Specify Plan Name and Description

At the **Plan Details** step of the wizard, use the **Plan name** and **Description** fields to enter a name for the new plan and to provide a description for future reference. The maximum length of the plan name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

You can also provide a contact name, email and telephone number of a person responsible for the plan.

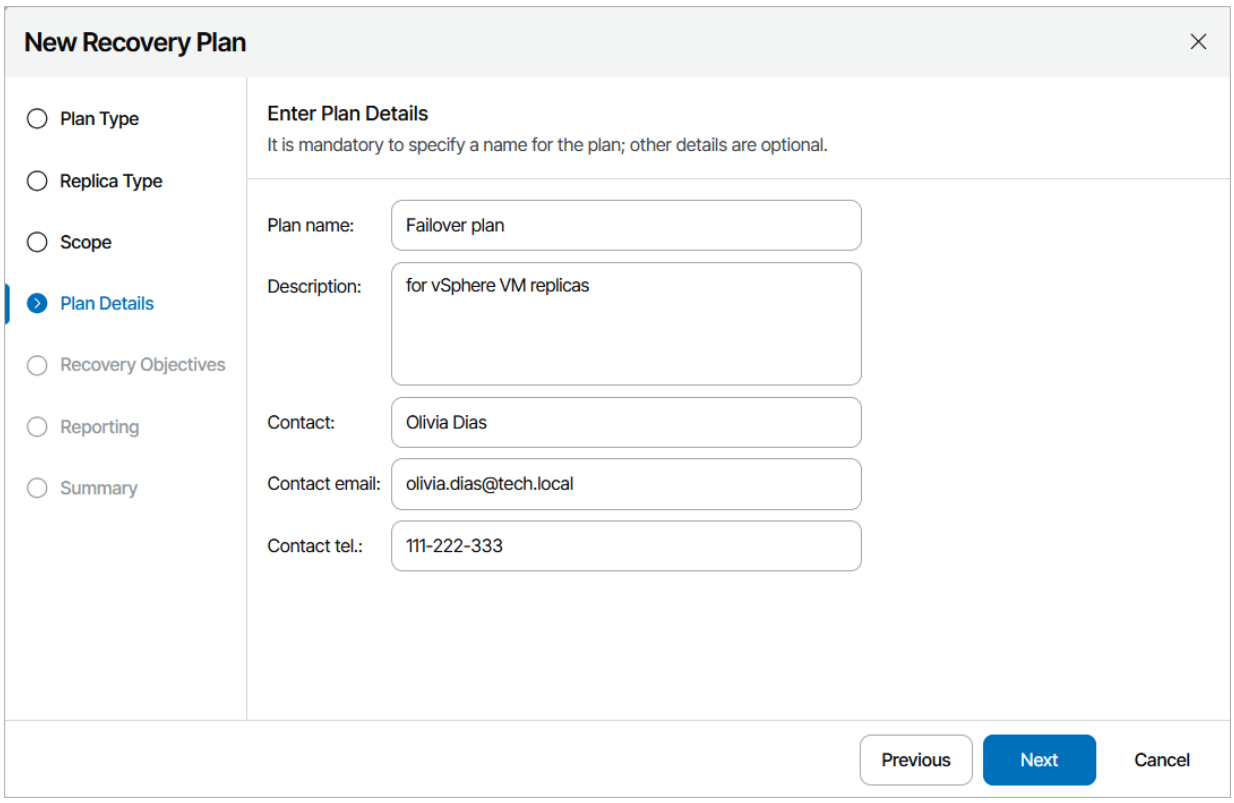

Step 5. Specify Target RTO and RPO

At the **Recovery Objectives** step of the wizard, define your Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for the plan:

- The **RPO** defines the maximum acceptable period of data loss.
- The **RTO** represents the amount of time it should take to recover from an incident.

RTO and RPO performance will be recorded in the Plan Readiness Check, Plan Execution and Plan Audit reports, and you will be able to track the achieved RTO and RPO objectives for each plan on the Home Page Dashboard.

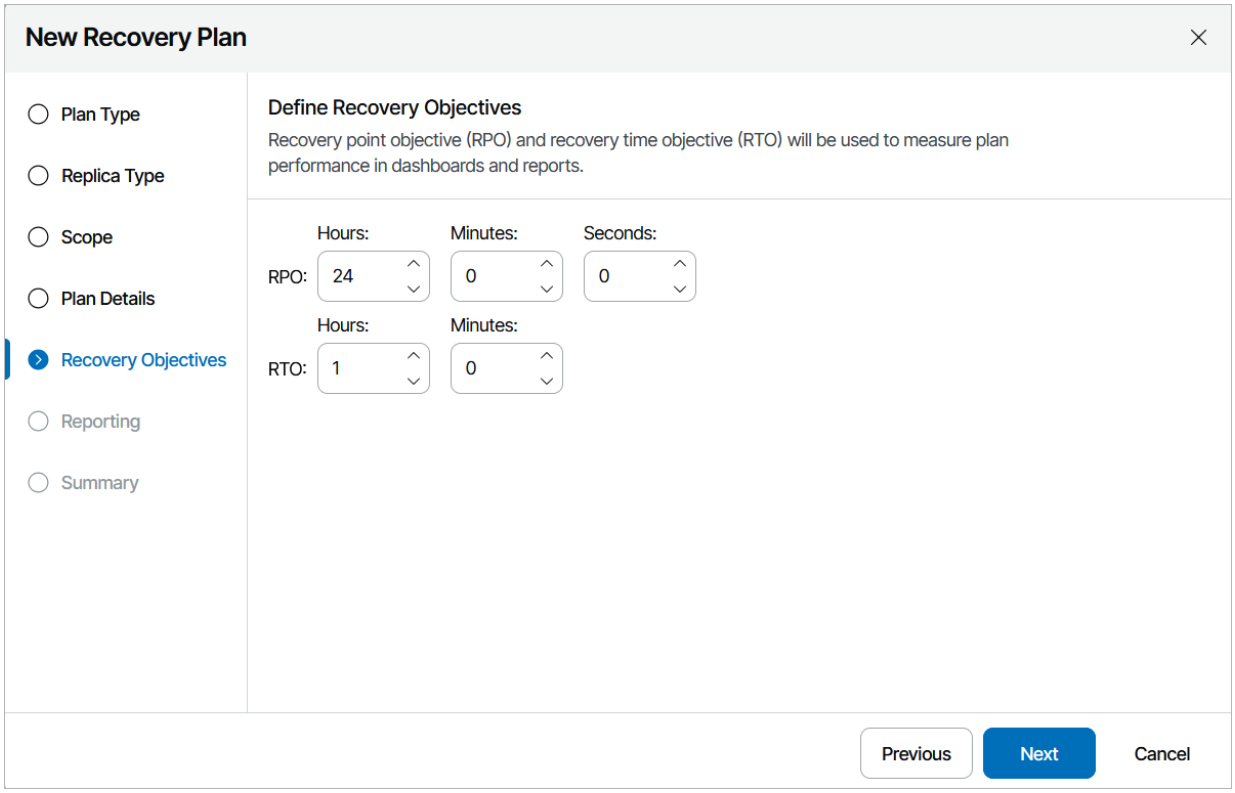

Step 6. Select Report Template

At the **Reporting** step of the wizard, select a document template that will be used as the cover page for all Orchestrator reports.

For a custom document template to be displayed in the list, it must be created and customized as described in section Managing Templates.

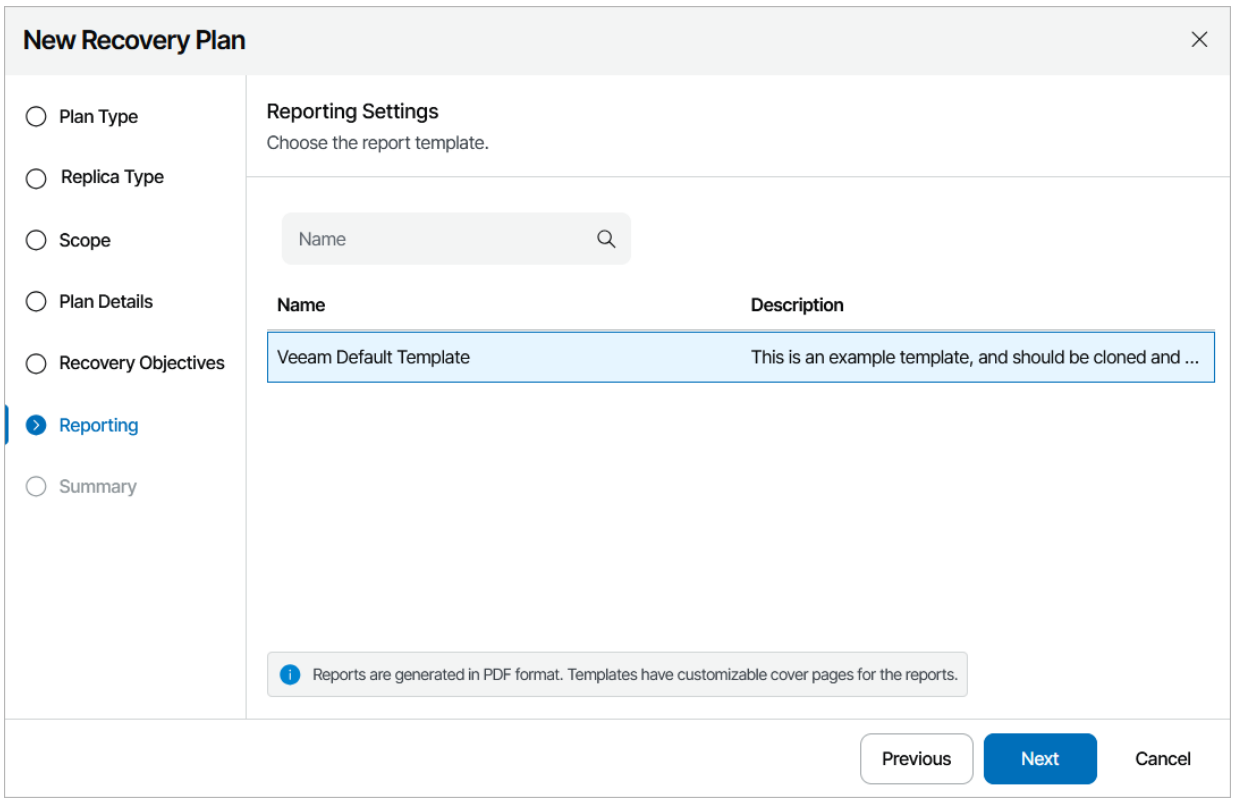

Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.

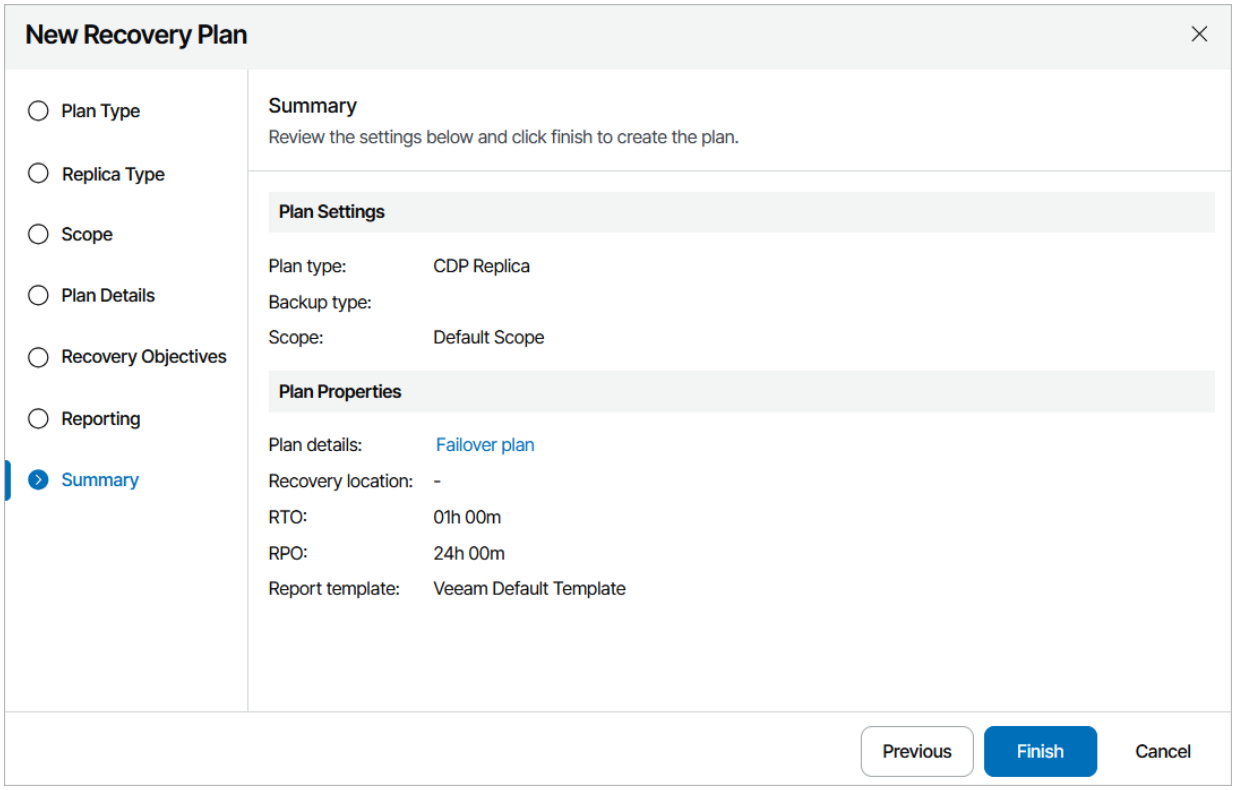

Editing CDP Replica Plans

If you want to specify granular settings not provided in the New Recovery Plan wizard, the Orchestrator UI allows you to customize CDP replica plans and configure the settings for groups, recovered VMs, plan steps and step parameters.

The procedures to edit replica, CDP replica, restore, storage and cloud plans are almost identical. For more information, see Editing Recovery Plans.

Testing CDP Replica Plans

You can start on-demand plan testing and configure test scheduling for any replica plan. There is almost no difference between the procedures performed for replica, CDP replica, restore and storage plans. For more information, see Testing Recovery Plans.

Scanning CDP Replica Plans

You can start on-demand plan scanning for malware and configure scan scheduling for any CDP replica plan. There is almost no difference between the procedures performed for replica, CDP replica, restore and cloud plans. For more information, see Scanning Recovery Plans.

Running and Scheduling CDP Replica Plans

To run a CDP replica plan, it must be *ENABLED*. To enable a plan:

1. Navigate to **Recovery Plans**.
2. Select the plan.
3. From the **Manage** menu, select **Properties**.

   OR-

   Right-click the plan name and select **Manage** > **Properties**.
4. Set the **Availability** toggle to *Enabled*.
5. Click **Save**.

If you do not enable a plan before you run it, the Run Plan wizard will force you to do that as soon as you try running the plan.

> **NOTES**
>
> 1. An Orchestrator Administrator or Plan Operator can force-enable a plan in the **Run Plan** wizard. However, a Plan Operator will not be able to run a disabled CDP replica plan.
>
>   For more information on roles that can be assigned to users and user groups working with the Orchestrator UI, see Managing User Accounts.
> 2. For security purposes, all 'real-world' actions associated with CDP replica plans (such as failover and failback) require password confirmation.

Scheduling Failover

You can schedule a time for a CDP replica plan to execute. Only the failover process can be scheduled — all other operations (failback, undo failover and so on) must be performed manually in the Orchestrator UI.

To schedule a CDP replica plan:

1. Navigate to **Recovery Plans**.
2. Select the plan. From the **Manage** menu, select **Schedule**.

   -OR-

   Right-click the plan name and select **Manage** > **Schedule**.
3. In the **Scheduled Tasks** window, do the following:

   a. Set the **Schedule plan execution** toggle to *Enabled*.

   b. Click the **Configure schedule** link and choose whether you want to run the plan on schedule or after any other plan:

   - If you want to run the plan at a specific time, click the **Schedule** icon in the **Run on** field, set the desired date and time, and click **Apply**.
   - If you want to run the plan after another plan, select the **Schedule after plan** check box and click **Choose plan**. Then, in the **Select Plan** window, select the necessary plan and click **Apply**.

     For a plan to be displayed in the list, it must be *ENABLED* as described in section Running and Scheduling Restore Plans.

c. Set the **Malware actions** toggle to *Enabled* if you want to check restore points created for machines included in the plan for malware flags.

   For more information on how Orchestrator performs malware scan, see the Veeam Recovery Orchestrator User Guide, section Overview.

d. Review the configuration information and click **Save**.

> **TIP**
>
> You can also scan a recovery plan for possible malware without scheduling the plan execution. To do that, follow the instructions provided in section Scanning Recovery Plans.

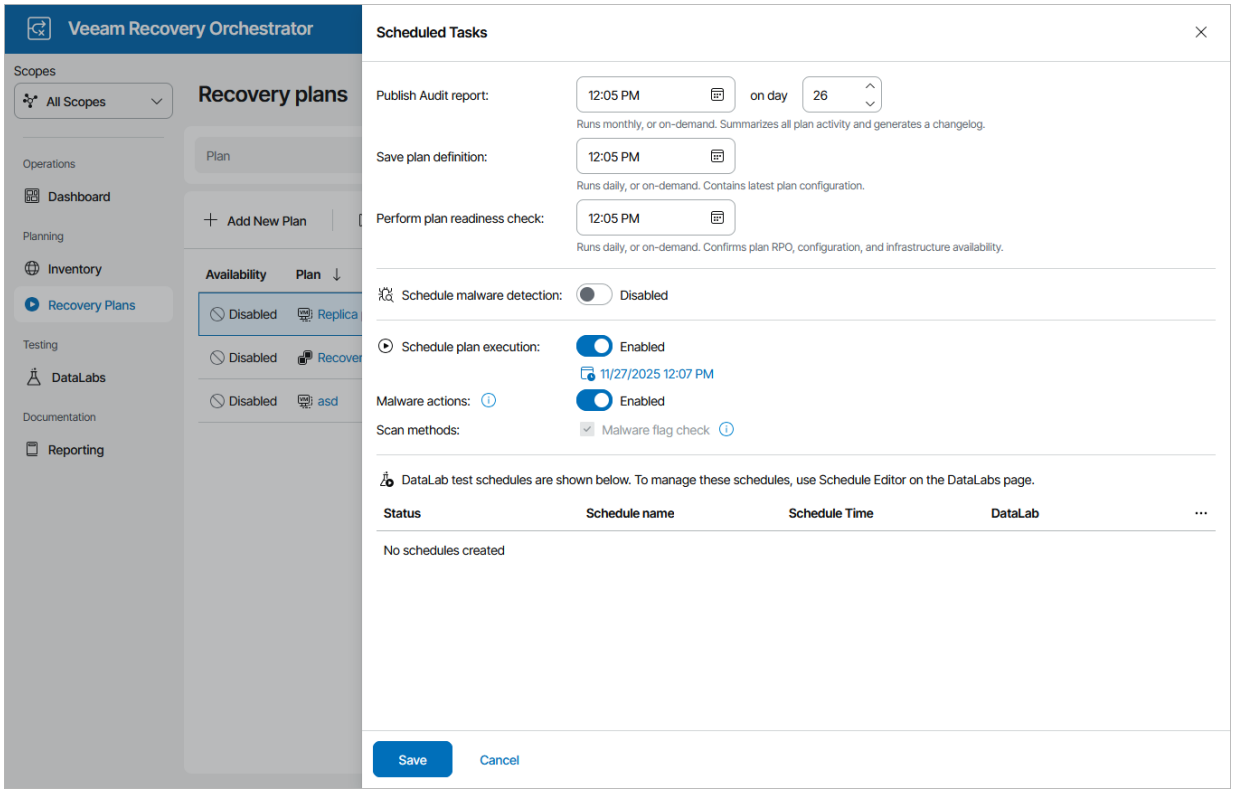

> **TIP**
>
> You can disable a configured schedule if you no longer need it. To do that, set the **Schedule plan execution** toggle to *Disabled* in the **Scheduled Tasks** window.

Running Failover

The **Run** action causes VMs in a plan to fail over to their replicas. For more information on the failover process, see the Veeam Backup & Replication User Guide, section Failover.

To run a CDP replica plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Run**.

3. In the **Run Plan** window, do the following:

   a. For security purposes, retype your password and click **Next**.

      You must also select the **Force-enable the plan** check box if you have not enabled the plan yet.

   b. In the **Restore point** section, choose whether you want to use a short-term or long-term restore point to recover VM replicas:

      - Short-term restore points are replicated states that are created with the shortest RPO (several seconds or minutes) and stored according to the short-term retention settings (no longer than several hours).
      - Long-term restore points are restore points that are created with a longer RPO (several hours) and stored according to the long-term retention settings (up to several days). Depending on the specified CDP policy settings, long-term restore points can be application-consistent and crash-consistent.

      For more information on CDP retention policies, see the Veeam Backup & Replication User Guide, section Creating CDP Policies.

   c. In the **Timestamp** field, choose a restore point that will be used to recover VM replicas.

   d. In the **Malware actions** field, choose whether you want to check restore points created for machines included in the plan for malware flags.

      For more information on how Orchestrator performs malware scan, see the Veeam Recovery Orchestrator User Guide, section Overview.

   e. Review configuration information and click **Run**.

> **TIP**
>
> You can also scan a restore plan for possible malware without running the plan. To do that, follow the instructions provided in section Scanning Recovery Plans.

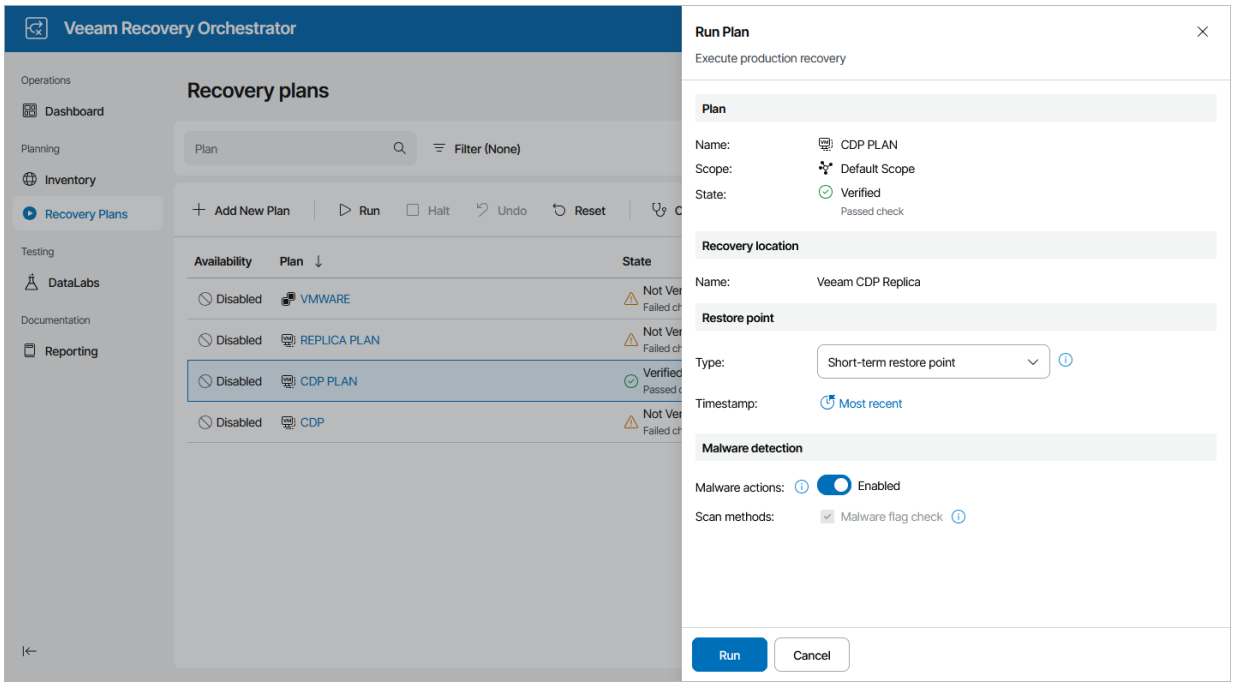

The plan goal is to reach the *FAILOVER* state. If any critical error is encountered, the plan will stop with the *HALTED* state. To learn how to work with *HALTED* CDP replica plans, see Managing Halted Plans.

Halting Failover

The **Halt** action interrupts plan execution. Any steps currently executing will be completed, then the plan will enter the *HALTED* state. To learn how to work with *HALTED* CDP replica plans, see Managing Halted Plans.

To stop a running CDP replica plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Halt**.
3. In the **Halt Plan Execution** window, do the following:
   a. For security purposes, retype your password and click **Next**.
   b. Review configuration information and click **Halt**.

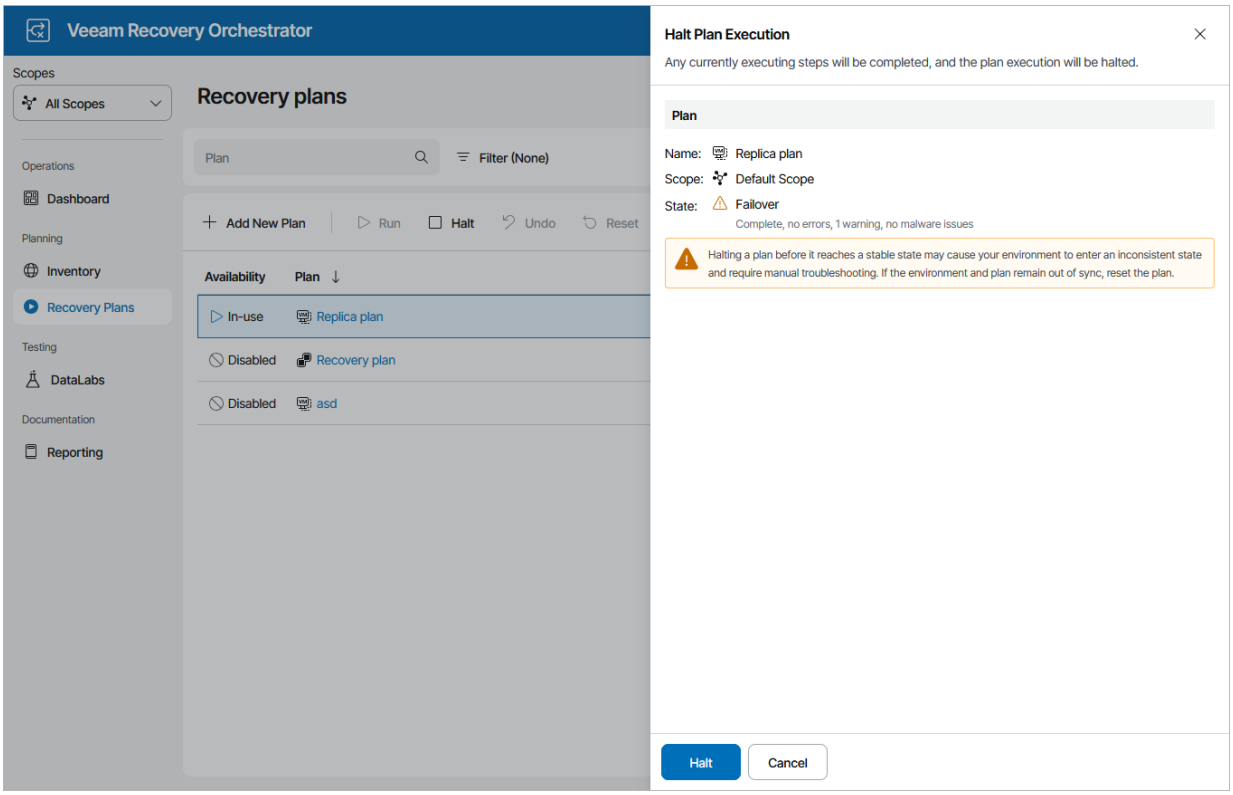

Finalizing Failover

Orchestrator provides you with a number of options to finalize failover to VM replicas:

- Perform permanent failover
- Perform failback
- Commit failback

For more information on the available options, see the Veeam Recovery Orchestrator User Guide, section Failover and Failback.

Running Permanent Failover

To perform permanent failover for a plan in the *FAILOVER* state:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Run**.
3. In the **Progress Plan** window, do the following:
    a. For security purposes, retype your password and click **Next**.
    b. In the **Recovery actions** section, select the **Permanent failover** option.
    c. Review configuration information and click **Run**.

> **NOTE**
>
> Failback will no longer be an option once the permanent failover process is complete.

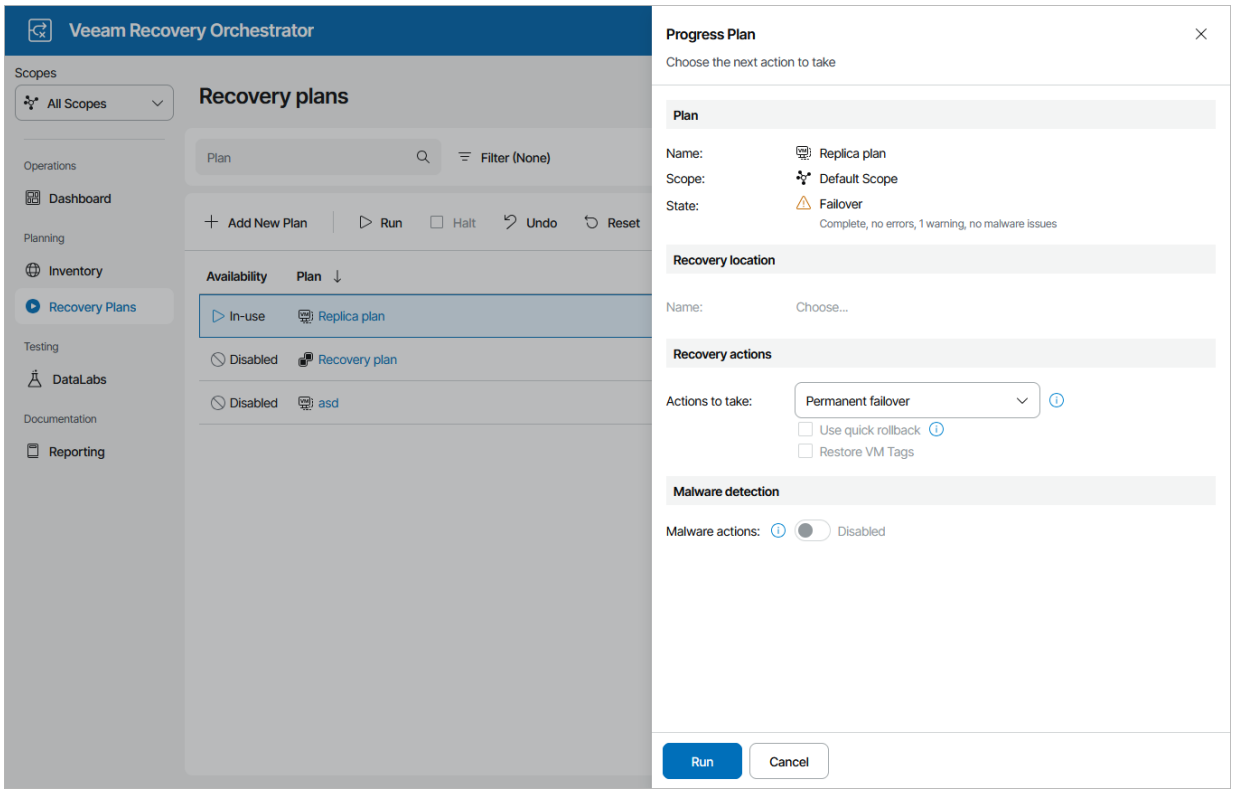

Running Failback

To start failback for a plan in the *FAILOVER* state:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Run**.
3. In the **Progress Plan** window, do the following:
1. In the **Progress Plan** window, do the following:
    a. For security purposes, retype your password and click **Next**.

b. In the **Recovery Location** section, select a location to which VMs will be recovered.

For a recovery location to be displayed in the list of available locations, it must be created and added to the list of inventory items available for the scope, as described in section Managing Recovery Locations.

> **NOTE**
>
> Orchestrator will perform failback using all settings configured for the location — except Instant VM Recovery and backup copy preference. These settings are not applicable to failback operations.

If you want to fail back to a new recovery location and the selected location includes multiple hosts, datastores and networks, Orchestrator will use the round-robin algorithm to recover VMs. For more information, see the Veeam Recovery Orchestrator User Guide, section How Orchestrator Places VMs During Failback.

c. In the **Recovery actions** section, do the following:

i. From the **Actions to take** drop-down list, select the **Prepare to failback** option for Orchestrator to switch from VM replicas to the source VMs when you run the plan next time.

ii. [Applies only if you have selected the Original VM Location] Select the **Use quick rollback** check box if you want to instruct Orchestrator to synchronize changed data blocks only — this may help you speed up the failback process significantly.

For more information on the quick rollback process, see the Veeam Backup & Replication User Guide, section Quick Rollback.

d. In the **Malware detection** section, choose whether you want to check restore points created for the VMs included in the plan for malware flags.

For more information on how Orchestrator performs malware scan, see the Veeam Recovery Orchestrator User Guide, section Overview.

e. Review configuration information and click **Run**.

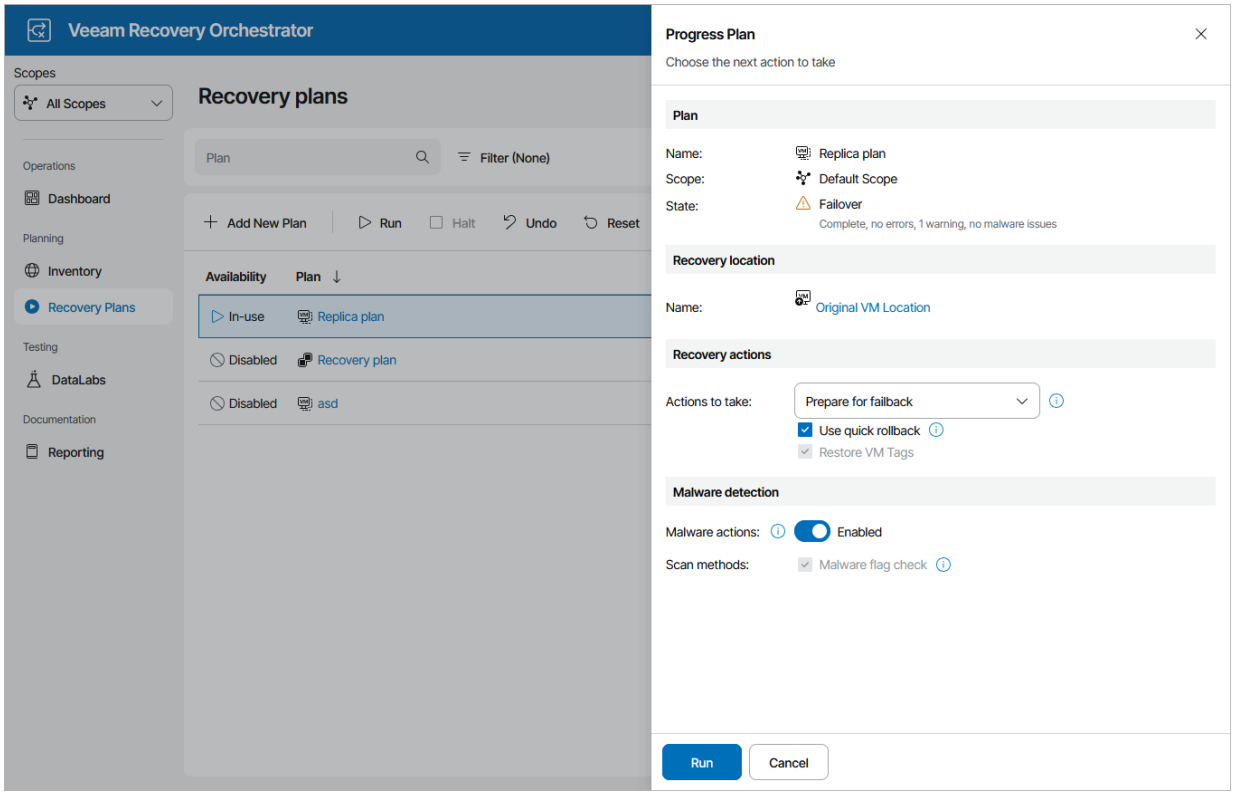

Committing Failback

To commit failback for a plan in the *FAILBACK* state:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Run**.
3. In the **Progress Plan** window, do the following:
   a. For security purposes, retype your password and click **Next**.
   b. In the **Recovery actions** section, select the **Failback now** option to switch from VM replicas to the source VMs immediately.
   c. Review configuration information and click **Run**.

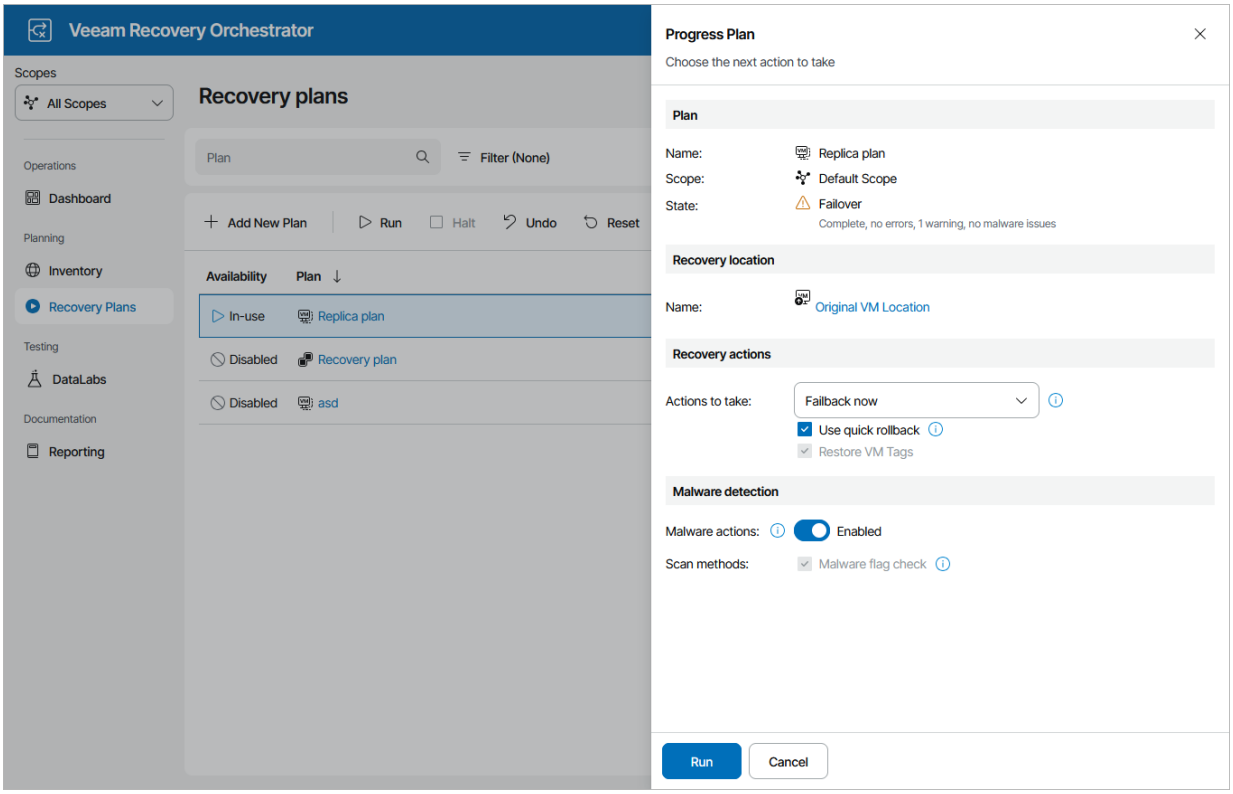

TIP

After the commit failback process completes, Orchestrator will leave the plan in the *IN-USE* mode. By design, this makes the results of the commit failback process accessible in the Orchestrator UI as long as required, and also prevents the plan from being modified by any automatic updates related to infrastructure changes.

If you want to perform any further actions with the plan (for example, to run readiness checks or to execute the plan again), reset the plan as described in section Resetting CDP Replica Plans.

Undoing Failover

The **Undo Failover** action powers off VM replicas running on target hosts and rolls back to the VM state before failover. For more information on the undo failover operation, see the Veeam Backup & Replication User Guide, section Undo Failover.

To perform an undo operation for a plan in the *FAILOVER* state:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Undo**.
3. In the **Undo** window, do the following:
    a. For security purposes, retype your password and click **Next**.
    b. Review configuration information and click **Undo**.

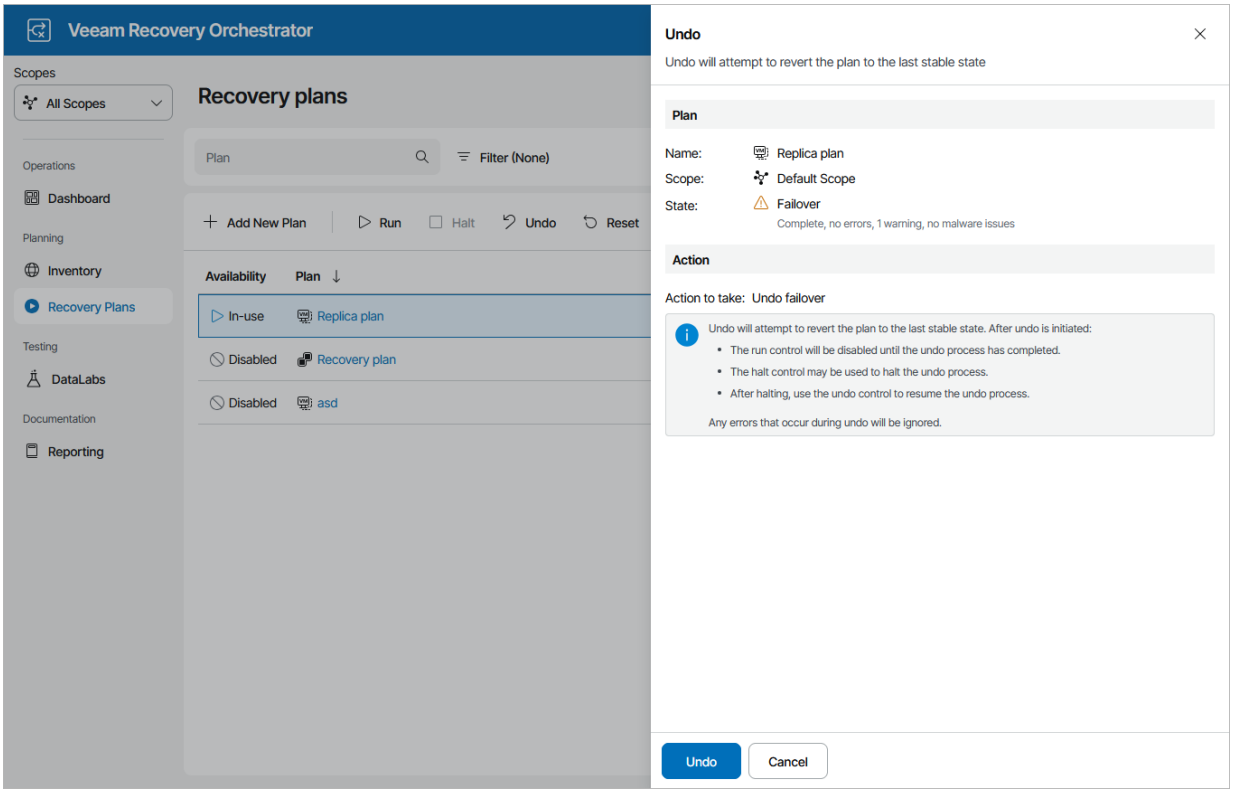

If the undo failover process encounters an error while being performed, it will not be halted automatically — the plan will proceed until the process completes. To terminate the undo failover process manually, use the **Halt** option to stop the currently running plan as described in section Halting Failover. To resume the undo failover process again, use the **Undo** option.

Undoing Failback

The **Undo Failback** action powers on VM replicas running on target hosts and switches from the production VMs back to the VM replicas — as a result, the plan acquires the *FAILOVER* state. For more information on the undo failback operation, see the Veeam Backup & Replication User Guide, section Undo Failback.

To perform an undo operation for a plan in the *PREPARE FOR FAILBACK* or *FAILBACK* state:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Undo**.
3. In the **Undo** window, do the following:
    a. For security purposes, retype your password and click **Next**.

b. Review configuration information and click **Undo**.

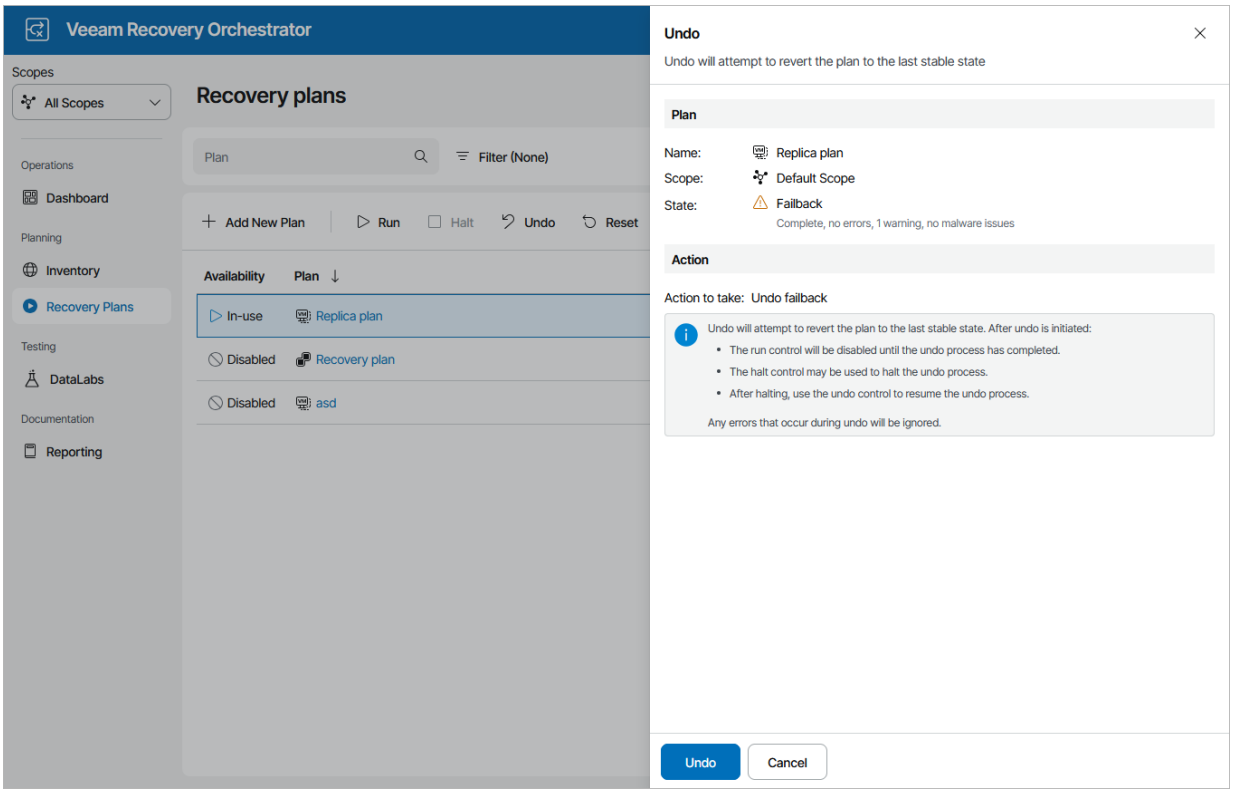

If the undo failback process encounters an error while being performed, it will not be halted automatically — the plan will proceed until the process completes. To terminate the undo failback process manually, use the **Halt** option to stop the currently running plan as described in section Halting Failover. To resume the undo failback process again, use the **Undo** option.

Resetting CDP Replica Plans

If a CDP replica plan becomes inconsistent with the virtual environment, you can reset the plan. This will return the plan to the *DISABLED* state, without making any changes to the external virtual infrastructure.

To reset a CDP replica plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Reset**.
3. In the **Reset Plan** window, do the following:
   a. For security purposes, retype your password and click **Next**.
   b. Select the **Run a quick check after reset is complete** check box to run a readiness check after the reset.

c. Review configuration information and click **Reset**.

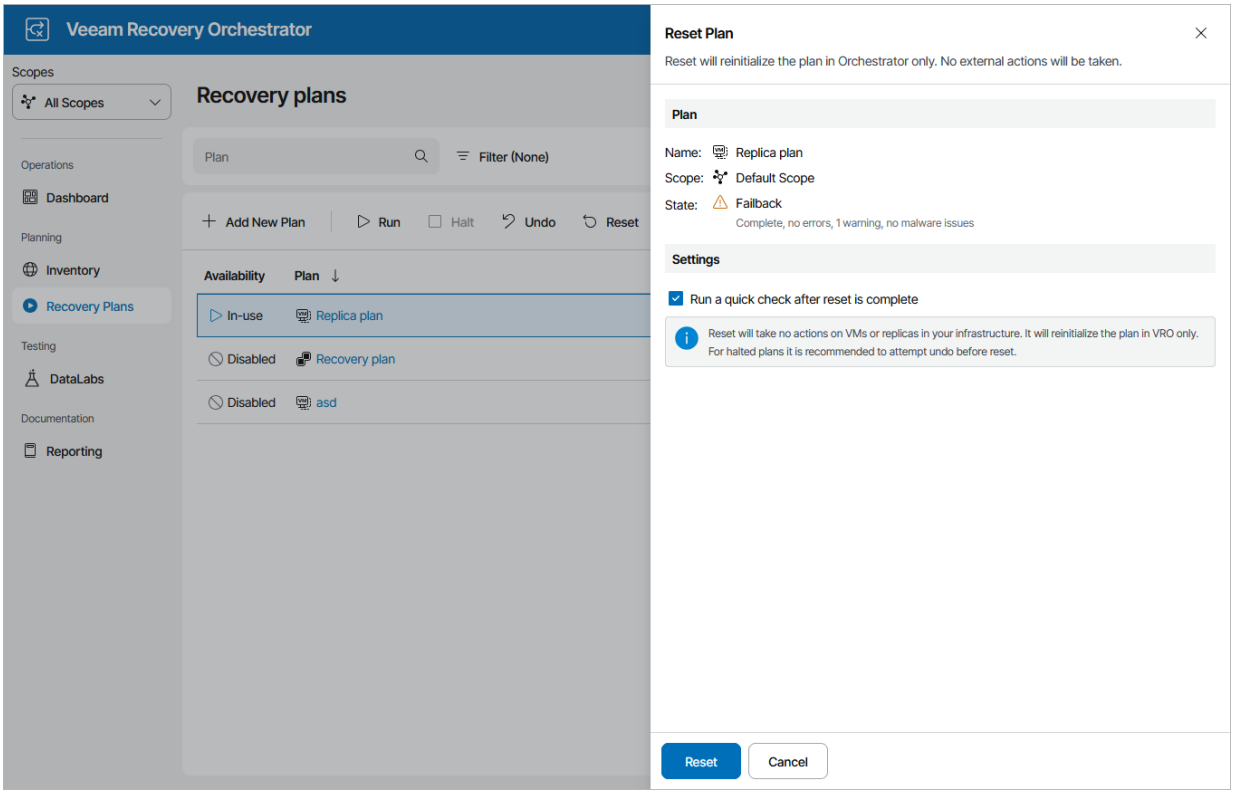

Halting Plan Testing

The **Halt** action interrupts plan testing. You may need to halt plan testing, for example, if you need to fix some environment-related issues and then proceed with testing later (in this case, VM replicas will still continue to run). Or you may need to stop the testing process completely, for example, if you no longer need to test the selected CDP replica plan (in this case, VM replicas will be reverted to the latest restore point).

To halt testing of a CDP replica plan:

1. Navigate to **Recovery Plans**.
2. Click the plan name to switch to the **Plan Details** page.
3. On the **Plan Details** page, select the plan and click **Halt**.
4. In the **Halt Plan Test** window, click **Halt Plan Test** to confirm the action.

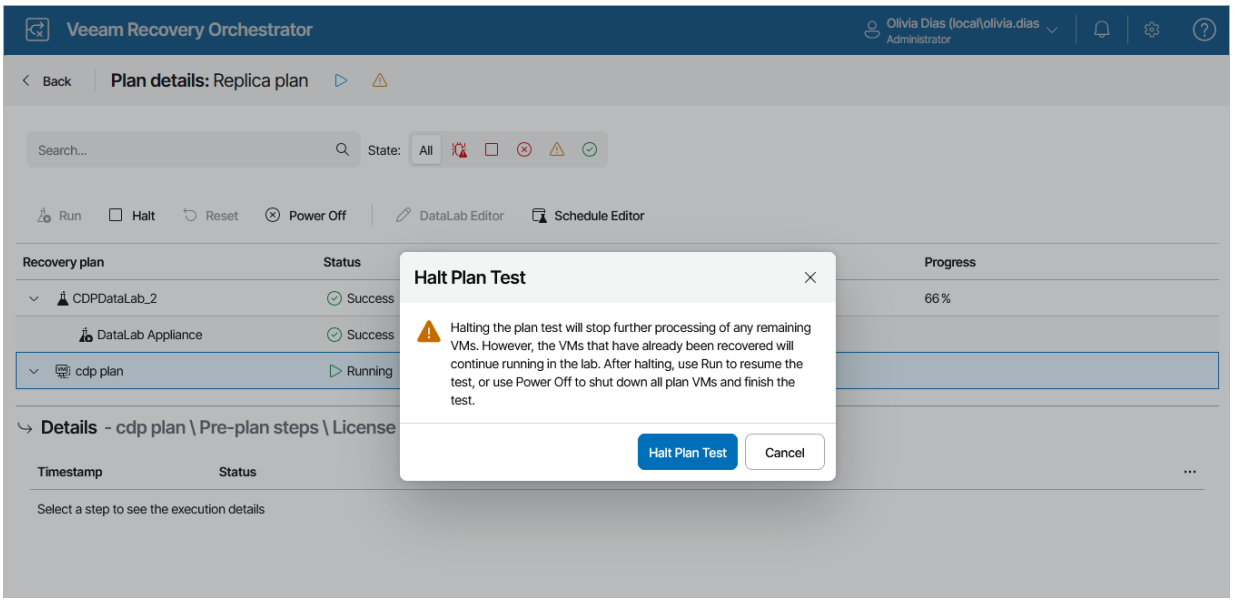

To cancel testing of a replica plan:

1. Navigate to **Recovery Plans**
2. Click the plan name to switch to the **Plan Details** page.
3. On the **Plan Details** page, select the plan and click **Power Off**.
4. In the **Power Off Plan Testing** window, click **Power Off** to confirm the action.

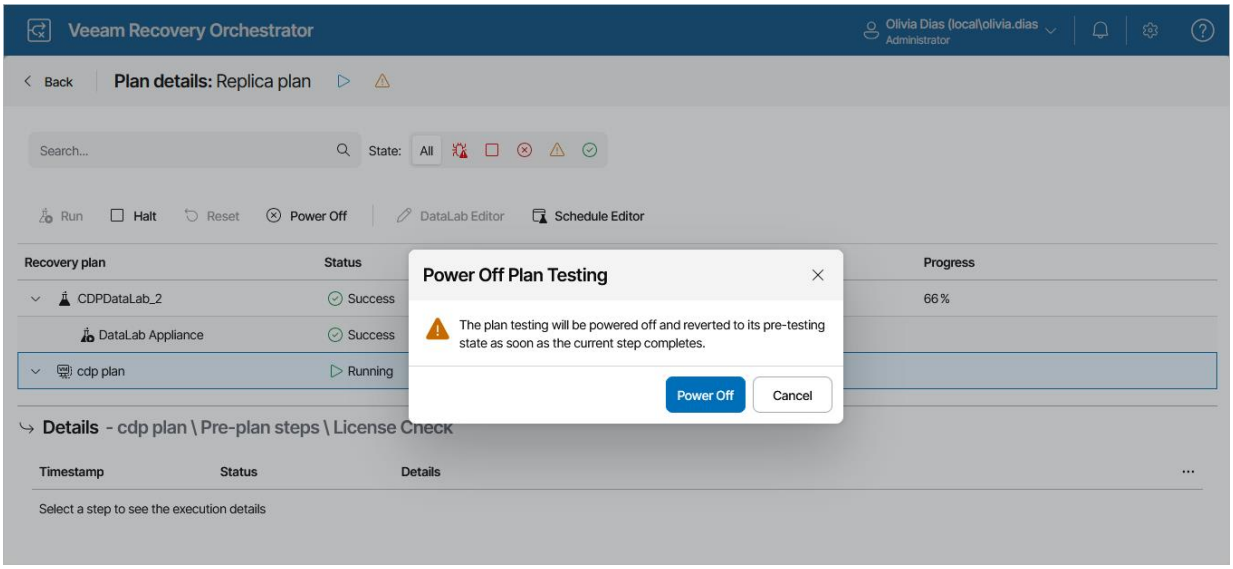

Managing Halted CDP Replica Plans

If a critical step fails for a VM from a critical inventory group, the plan may enter the *HALTED* state. To troubleshoot reasons why a plan failed, use the **Plan Execution Report** generated as soon as the currently performed action completes. For more information on how to track plan performance history, see Viewing Plan Execution History.

After you eliminate the problem that caused the plan to become *HALTED*, you have the following options to resume the plan:

- Repeat the last failed step.
- Proceed to the next step.

Running Halted CDP Replica Plans

To run a *HALTED* CDP replica plan:

1. Navigate to **Recovery Plans**.
2. Select the halted plan and click **Run**.
3. In the **Continue Execution** window, do the following:
   a. For security purposes, retype your password and click **Next**.
   b. In the **Retry settings** section, select an option to resume plan execution.

      Choose whether you want to proceed with plan execution from the next plan step or to retry the failed step.

   c. Review configuration information and click **Continue**. The failover process will be started.

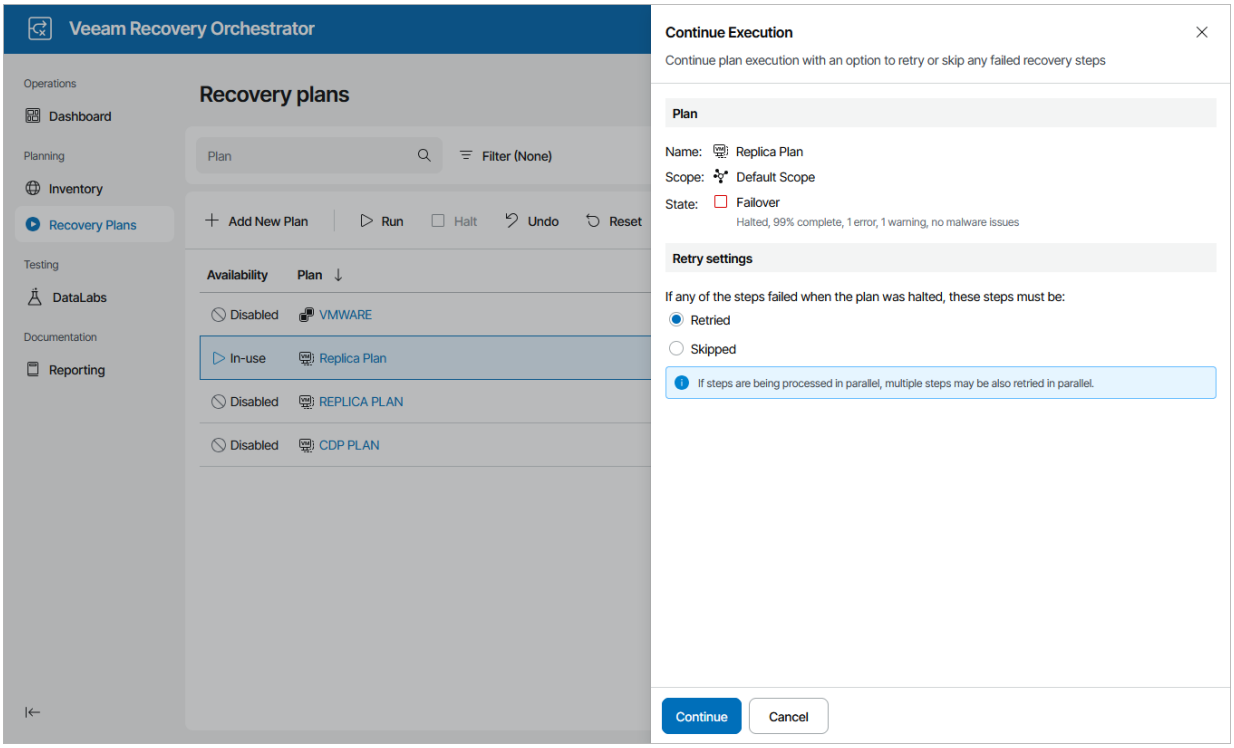

Resuming Plan Testing

To start the *HALTED* plan testing process:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Test**.

3. In the **Resume DataLab Test** window, choose whether you want to proceed with test execution from the next plan step or to retry the failed step, and then click **Continue**. The testing process will be started.

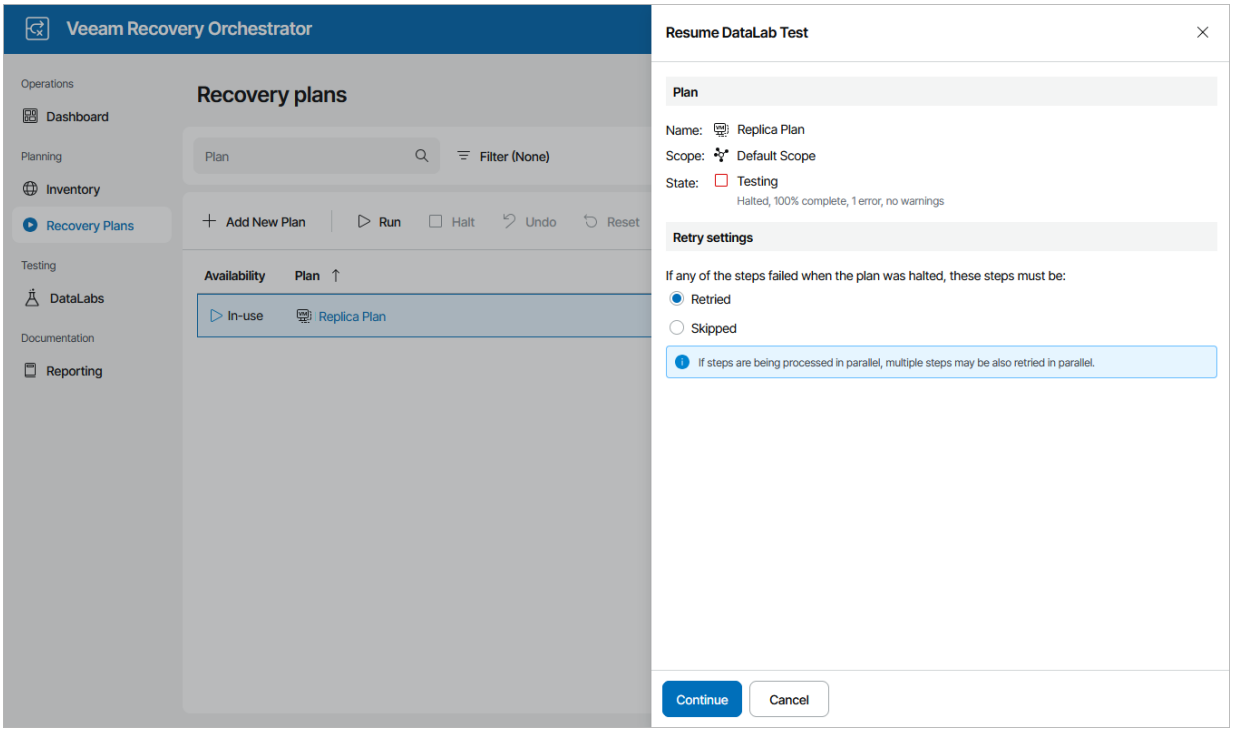

Undoing Halted CDP Replica Plans

To perform an undo operation for a *HALTED* CDP replica plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Undo**.
3. In the **Undo** window, do the following:
   a. For security purposes, retype your password and click **Next**.

b. Review configuration information and click **Undo**. The failover process will be started.

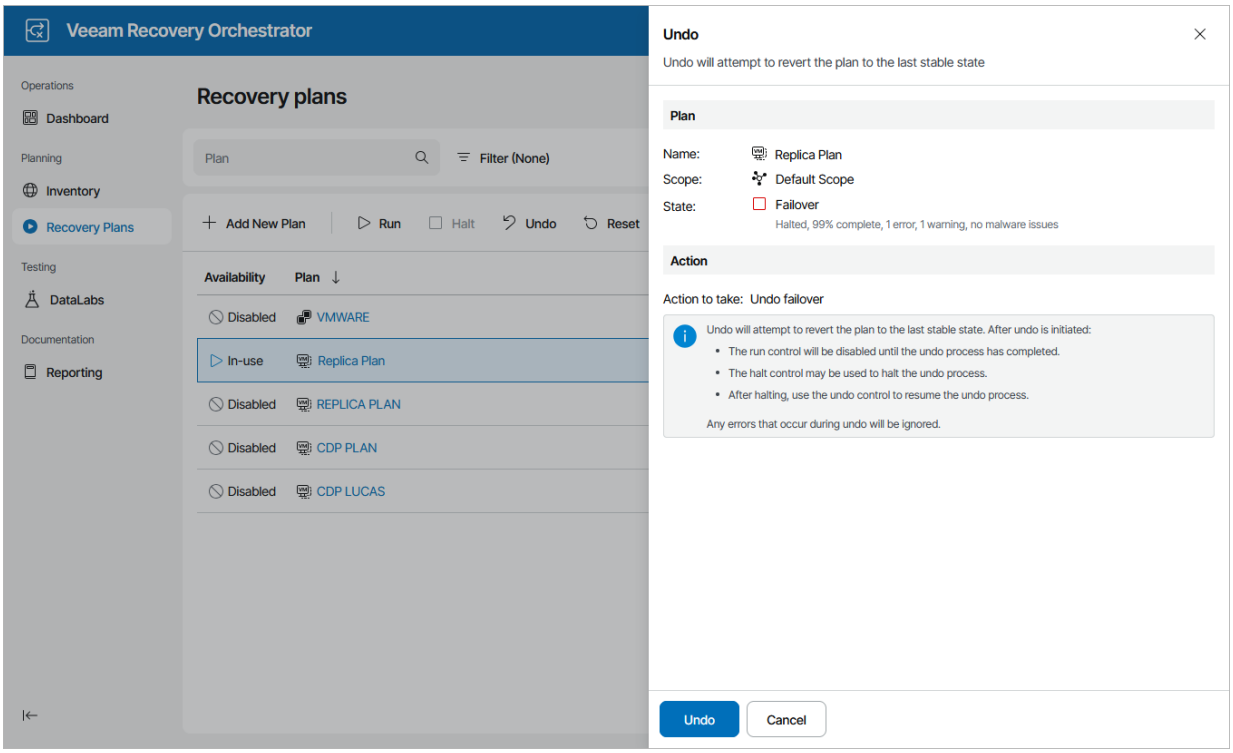

If a plan repeatedly enters the *HALTED* state due to misconfiguration or changes in the external environment, the only option left may be to RESET the plan.

Resetting Halted CDP Replica Plans

To reset a *HALTED* CDP replica plan, follow the instructions provided in section Resetting CDP Replica Plans.

> **NOTE**
>
> When you reset a CDP replica plan, Orchestrator returns it to the *DISABLED* state without making any changes to the external virtual infrastructure. You may need to deal with any infrastructure reconfiguration manually.

Working with Restore Plans

The type of a recovery plan you create depends on whether you intend to use Orchestrator to switch to VM replicas, to restore machines from backups or backup copies, or to serve data from a destination (NetApp) or secondary (HPE) volume in case a disaster strikes.

If you want to recover vSphere VM, Hyper-V VM or Veeam agent backups to a VMware vSphere or Microsoft Hyper-V environment, create a restore plan.

Creating Restore Plans

To create a restore plan:

1. Navigate to **Recovery Plans**.
2. Click **Add New Plan**.
3. Complete the **New Recovery Plan** wizard:
    a. Choose a type of the plan.
    b. Choose the necessary backup type.
    c. Choose a scope for the plan.
    d. Specify a plan name and description.
    e. Specify the target RTO and RPO.
    f. Select a recovery location for the plan.
    g. Select a template for plan reports.
    h. Finish working with the wizard.

Step 1. Choose Plan Type

At the **Plan Type** step of the wizard, select the **Restore** option.

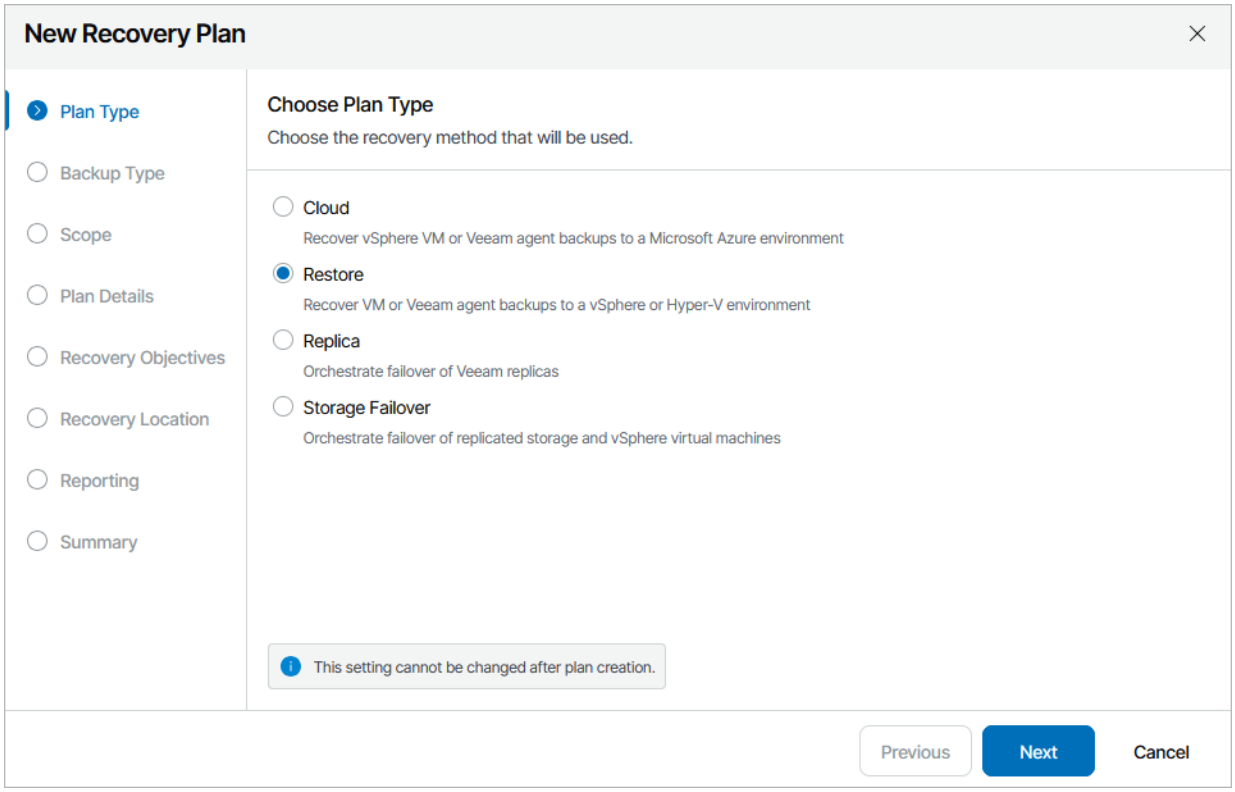

Step 2. Choose Backup Type

At the **Backup Type** step of the wizard, choose whether you want to recover machines from vSphere backups, Veeam Agent backups or Hyper-V backups.

Note that one restore plan can contain inventory groups of one type only (either vSphere, Veeam Agent or Hyper-V). To recover workloads added to inventory groups of different types, create separate restore plans.

IMPORTANT

For Orchestrator to be able to recover a machine to a VMware vSphere environment, the machine must have VMware Tools installed:

- For VMs recovered from Veeam agent backups, Orchestrator automatically verifies whether VMware Tools are installed on all machines included in a plan when running a readiness check or a DataLab test for the plan. However, this verification is supported for Windows-based machines only. For Linux-based machines, you must perform the verification manually.
- For VMs recovered from vSphere backups, the verification is performed automatically on the vCenter Server side — for both Windows-based and Linux-based VMs. To know how to install and upgrade VMware Tools in vSphere, see this VMware KB article.

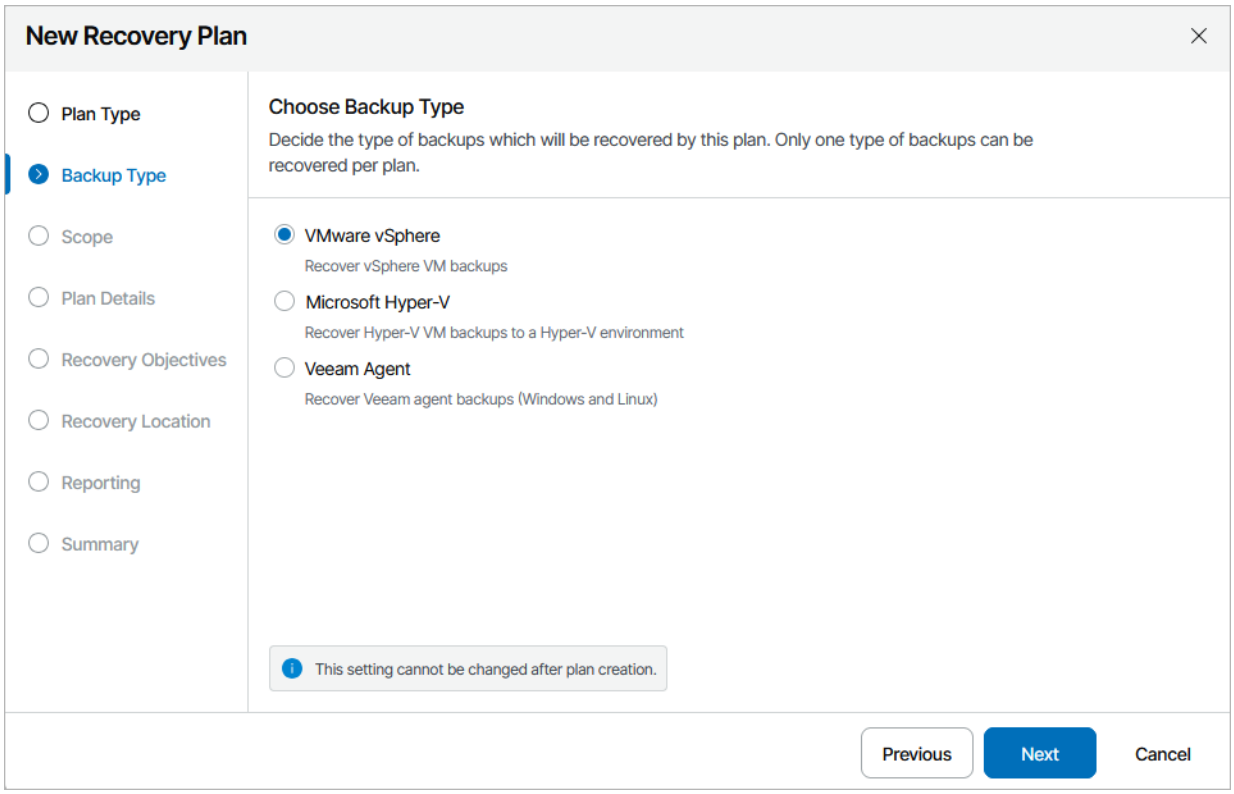

Step 3. Choose Plan Scope

At the **Scope** step of the wizard, select a scope for which you want to create the plan.

For a scope to be displayed in the list, it must be created and customized as described in section Managing Scopes.

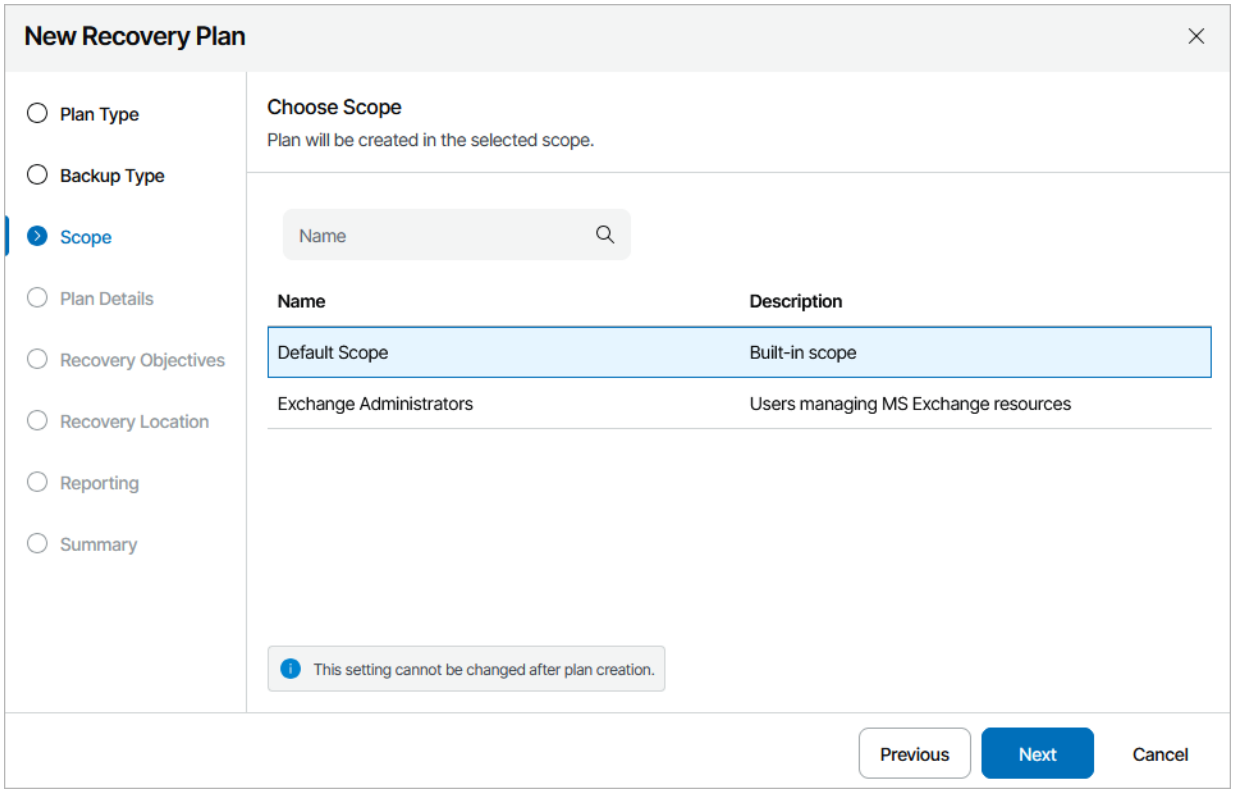

Step 4. Specify Plan Name and Description

At the **Plan Details** step of the wizard, use the **Plan name** and **Description** fields to enter a name for the new plan and to provide a description for future reference. The maximum length of the plan name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

You can also provide a contact name, email and telephone number of a person responsible for the plan.

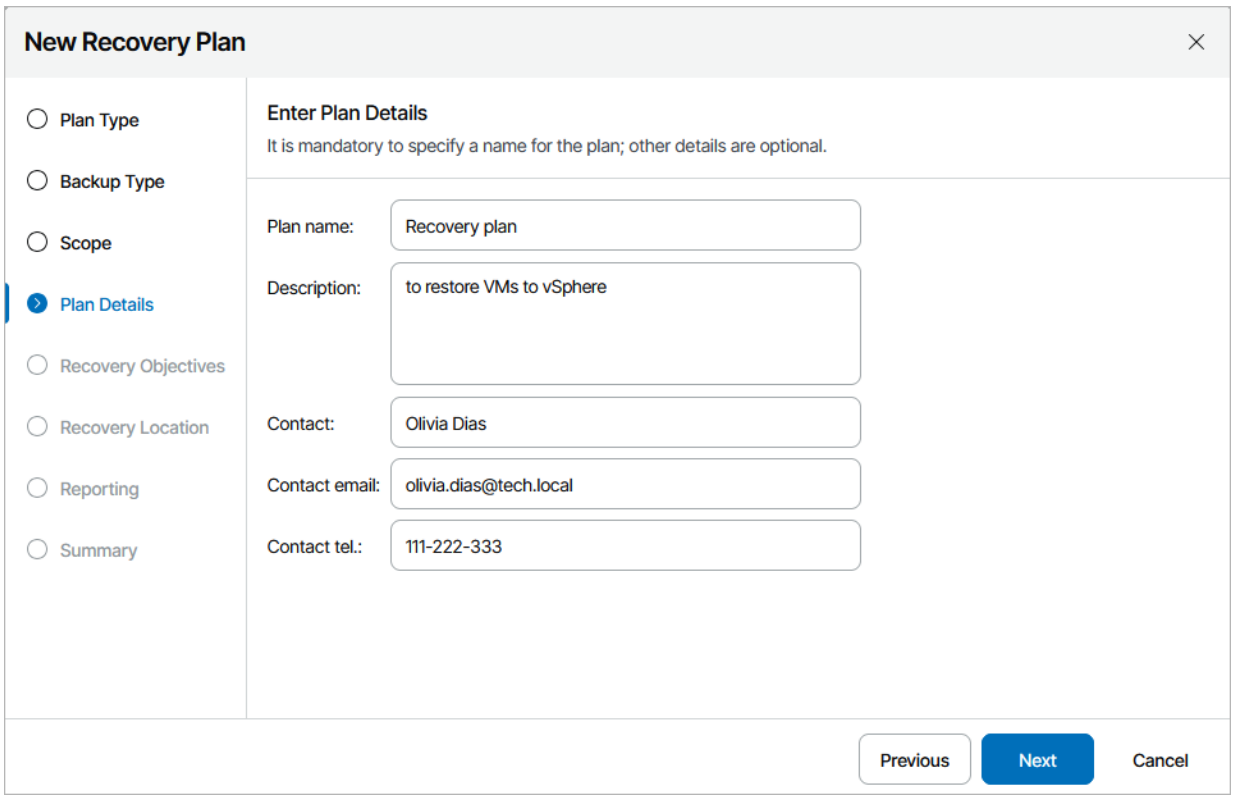

Step 5. Specify Target RTO and RPO

At the **Recovery Objectives** step of the wizard, define your Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for the plan:

- The **RPO** defines the maximum acceptable period of data loss.
- The **RTO** represents the amount of time it should take to recover from an incident.

> **NOTE**
>
> If you choose to perform malware scan while running the plan, Orchestrator will scan only one disk per mount server at a time. This process may take a while, affecting the plan RTO.

RTO and RPO performance will be recorded in the Plan Readiness Check, Plan Execution, Plan Audit and DataLab Test reports, and you will be able to track the achieved RTO and RPO objectives for each plan on the Home Page Dashboard.

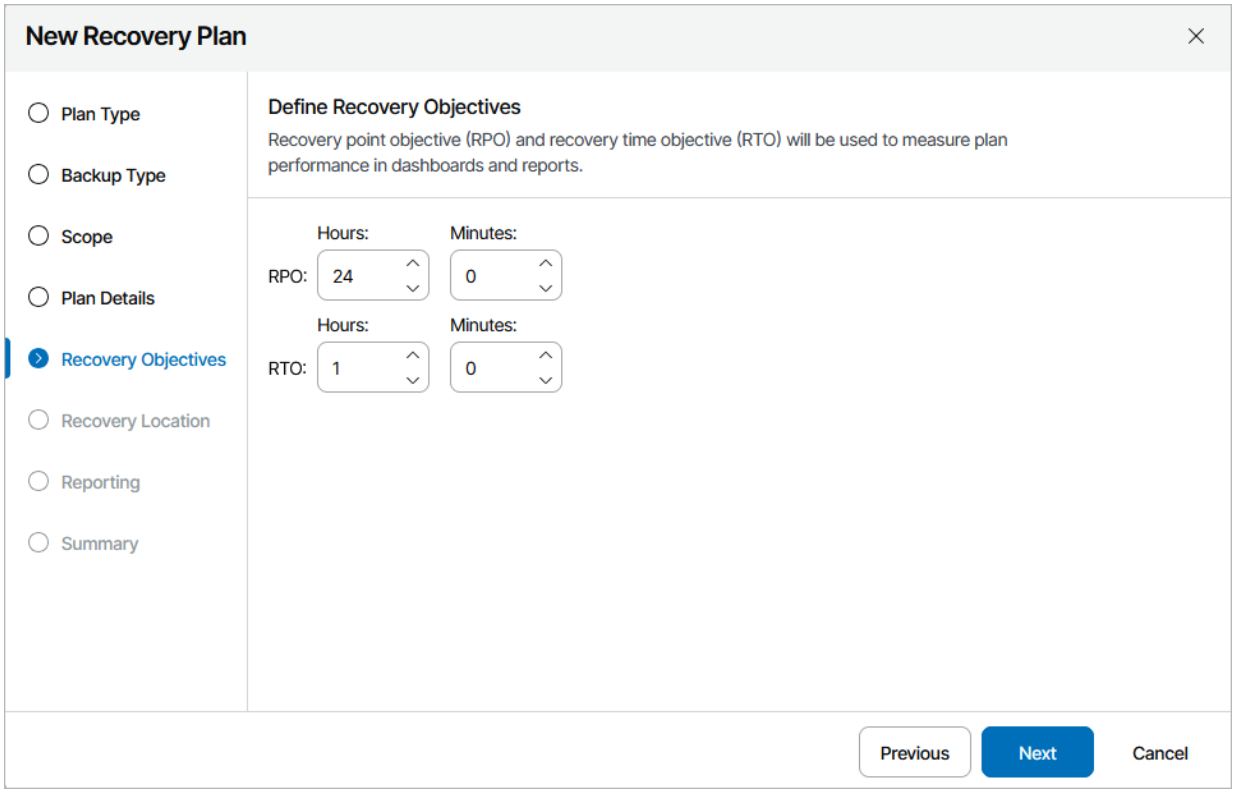

Step 6. Select Recovery Location

At the **Recovery Location** step of the wizard, select a location to which inventory groups added to the plan will be restored.

For a recovery location to be displayed in the list of available recovery locations, it must be created and added to the list of inventory items available for the scope, as described in section Managing Recovery Locations.

> **NOTE**
>
> When selecting a recovery location, you must make sure that target hosts specified when creating the location run a hardware version that is compatible with hardware versions of VMs included in the plan groups. For more information on version compatibility for vSphere VMs, see VMware Docs; for more information on version compatibility for Hyper-V VMs, see Microsoft Docs.

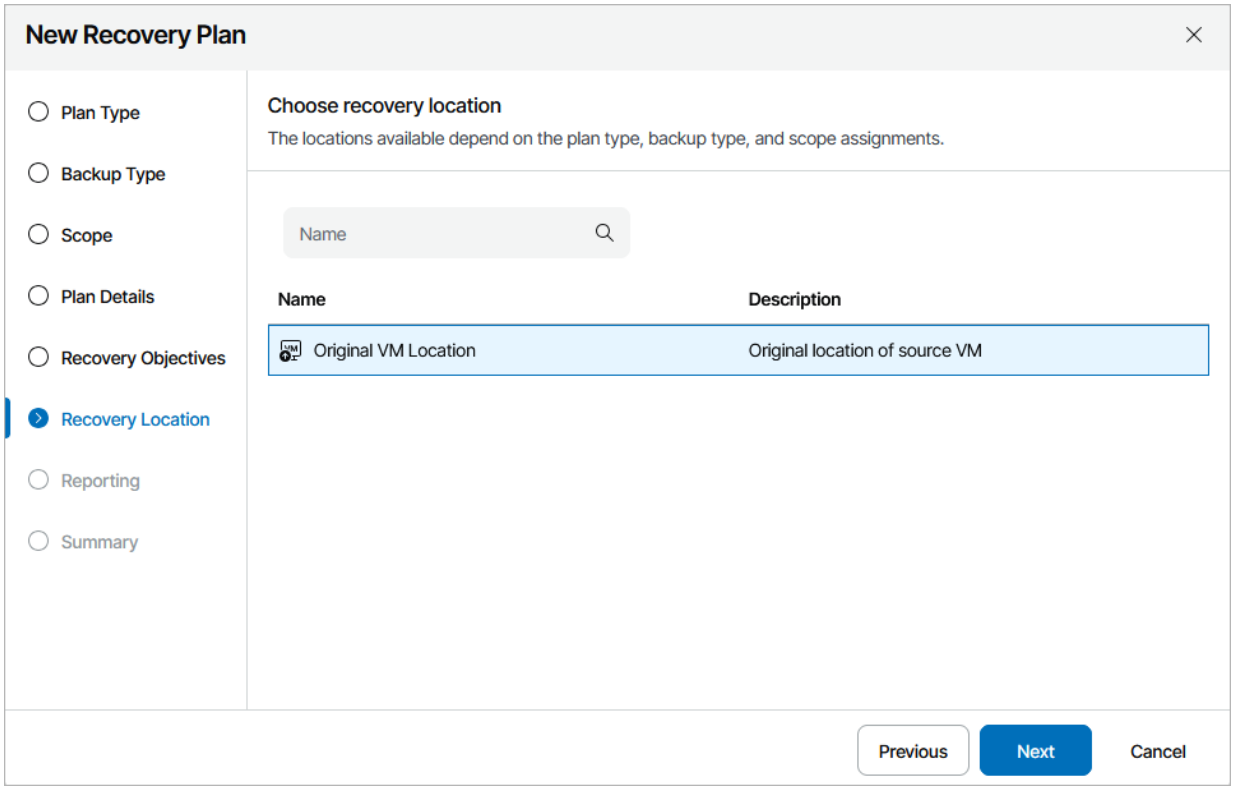

Step 7. Select Report Template

At the **Reporting** step of the wizard, select a document template that will be used as the cover page for all Orchestrator reports.

For a custom document template to be displayed in the list, it must be created and customized as described in section Managing Templates.

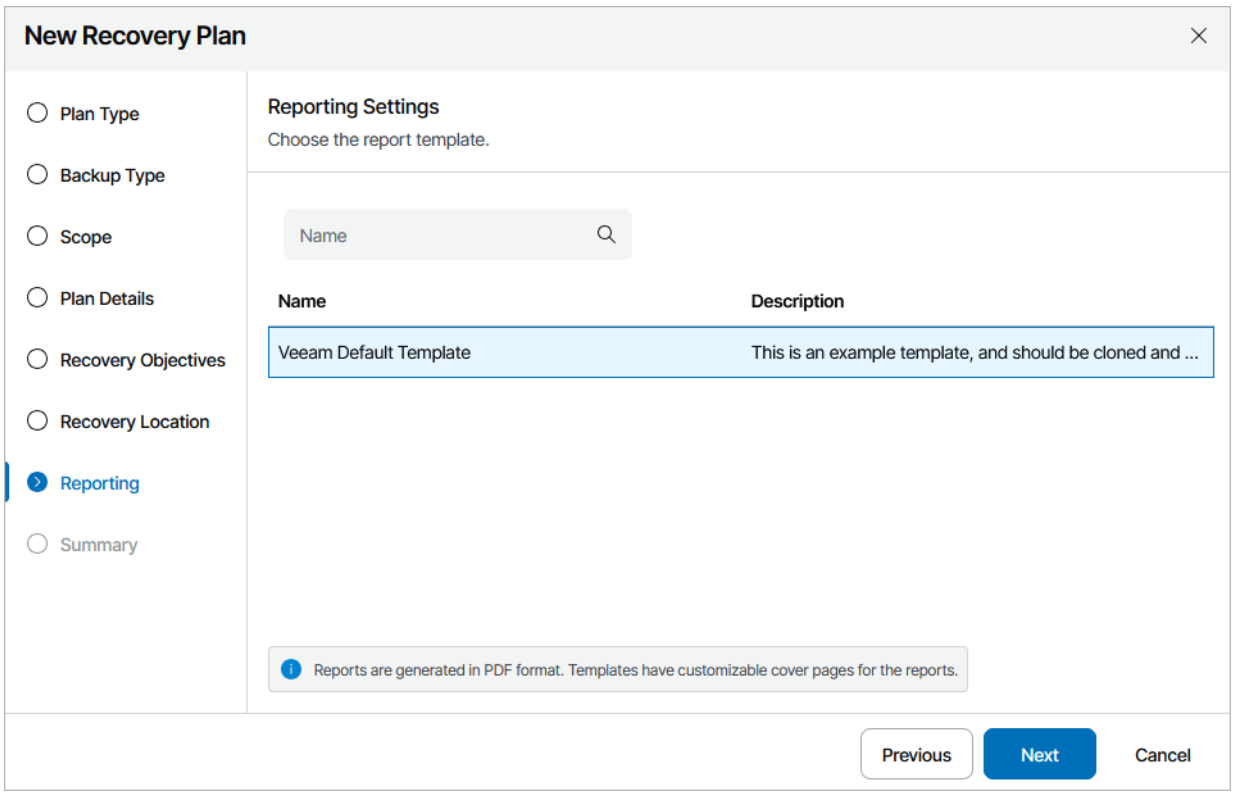

Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.

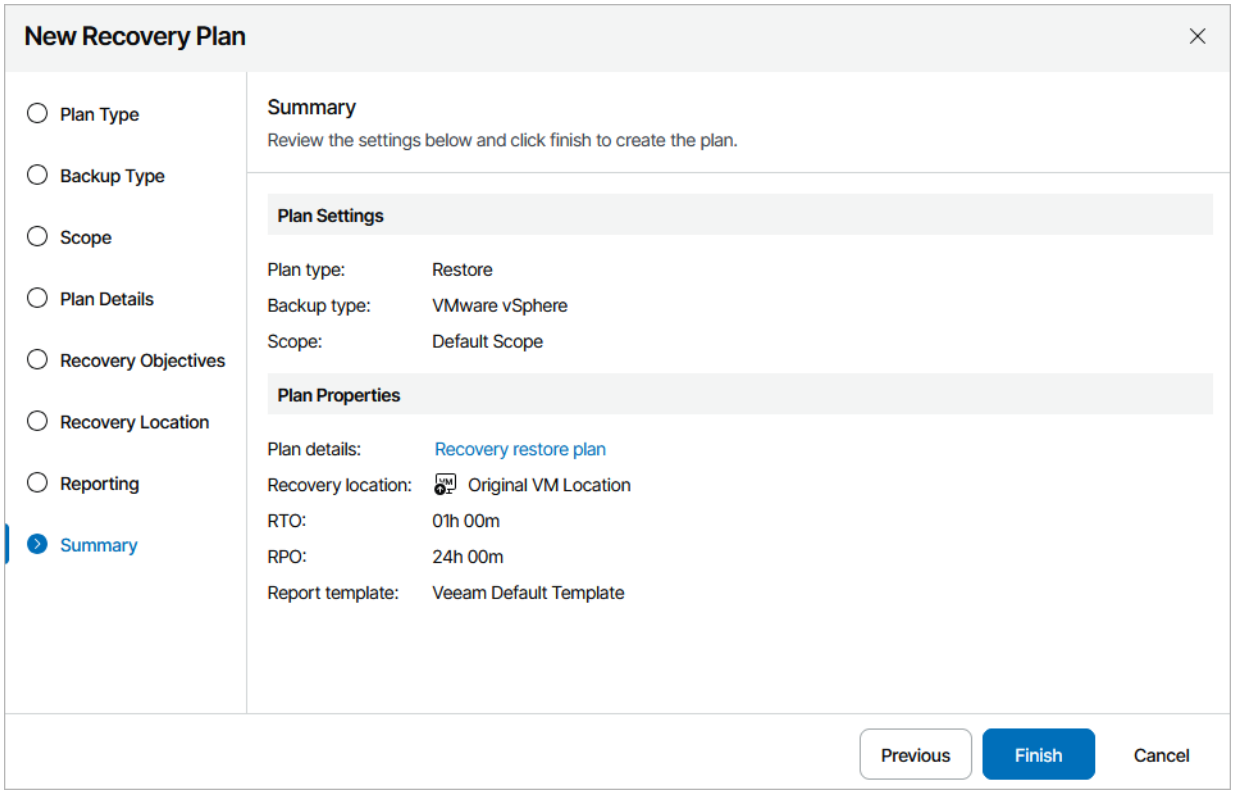

Editing Restore Plans

If you want to specify granular settings not provided in the New Recovery Plan wizard, the Orchestrator UI allows you to customize restore plans and configure the settings for groups, recovered VMs, plan steps and step parameters.

The procedures to edit replica, CDP replica, restore, storage and cloud plans are almost identical. For more information, see Editing Recovery Plans.

Testing Restore Plans

You can start on-demand plan testing and configure test scheduling for any restore plan. There is almost no difference between the procedures performed for replica, CDP replica, restore and storage plans. For more information, see Testing Recovery Plans.

Scanning Restore Plans

You can start on-demand plan scanning for malware and configure scan scheduling for any restore plan. There is almost no difference between the procedures performed for replica, CDP replica, restore and cloud plans. For more information, see Scanning Recovery Plans.

Running and Scheduling Restore Plans

> **IMPORTANT**
>
> When restoring a VM to a Hyper-V environment while the source VM still exists in the original location, Orchestrator may create a new VM using the same MAC address as the source VM. To avoid network issues and MAC address conflicts, it is recommended that you power the source VM off before you run the restore plan.

To run a restore plan, it must be *ENABLED*. To enable a plan:

1. Navigate to **Recovery Plans**.
2. Select the plan.
3. From the **Manage** menu, select **Properties**.

   OR-

   Right-click the plan name and select **Manage** > **Properties**.
4. Set the **Availability** toggle to *Enabled*.
5. Click **Save**.

If you do not enable a plan before you run it, the Run Plan wizard will force you to do that as soon as you try running the plan.

> **NOTES**
>
> 1. An Orchestrator Administrator or Plan Operator can force-enable a plan in the **Run Plan** wizard. However, a Plan Operator will not be able to run a disabled restore plan.
>
>   For more information on roles that can be assigned to users and user groups working with the Orchestrator UI, see Managing User Accounts.
> 2. For security purposes, all 'real-world' actions associated with restore plans require password confirmation.

Scheduling Restore

You can schedule a time for a restore plan to execute. To do that:

1. Navigate to **Recovery Plans**.
2. Select the plan. From the **Launch** menu, select **Manage** > **Schedule**.

   -OR-

   Right-click the plan name and select **Manage** > **Schedule**.
3. In the **Scheduled Tasks** window, do the following:
   a. Set the **Schedule plan execution** toggle to *Enabled*.
   b. Click the **Configure schedule** link and choose whether you want to run the plan on schedule or after any other plan:
      - If you want to run the plan at a specific time, click the **Schedule** icon in the **Run on** field, set the desired date and time, and click **Apply**.

- If you want to run the plan after another plan, select the **Schedule after plan** check box and click **Choose plan**. Then, in the **Select Plan** window, select the necessary plan and click **Apply**.

  For a plan to be displayed in the list of available plans, it must be *ENABLED* as described in section Running and Scheduling Restore Plans.

c. Set the **Malware actions** toggle to *Enabled* if you want to check restore points created for machines included in the plan for malware flags. When restoring to a VMware vSphere environment, you can also decide whether you want to scan these restore points with antivirus software, YARA rules or both. In this case, you must as the default Microsoft Windows-based or Linux-based mount server.

   By default, Orchestrator checks the most recent restore point on each machine. If all the restore points are infected, Orchestrator restores the machine to the selected recovery location without connecting it to any network. However, you can instruct Orchestrator to halt the plan and cancel the restore operation if no clean restore point is found.

   For more information on how Orchestrator performs malware scan, see the Veeam Recovery Orchestrator User Guide, section Overview.

d. Review the configuration information and click **Save**.

TIP

You can also scan a recovery plan for possible malware without scheduling the plan execution. To do that, follow the instructions provided in section Scanning Recovery Plans.

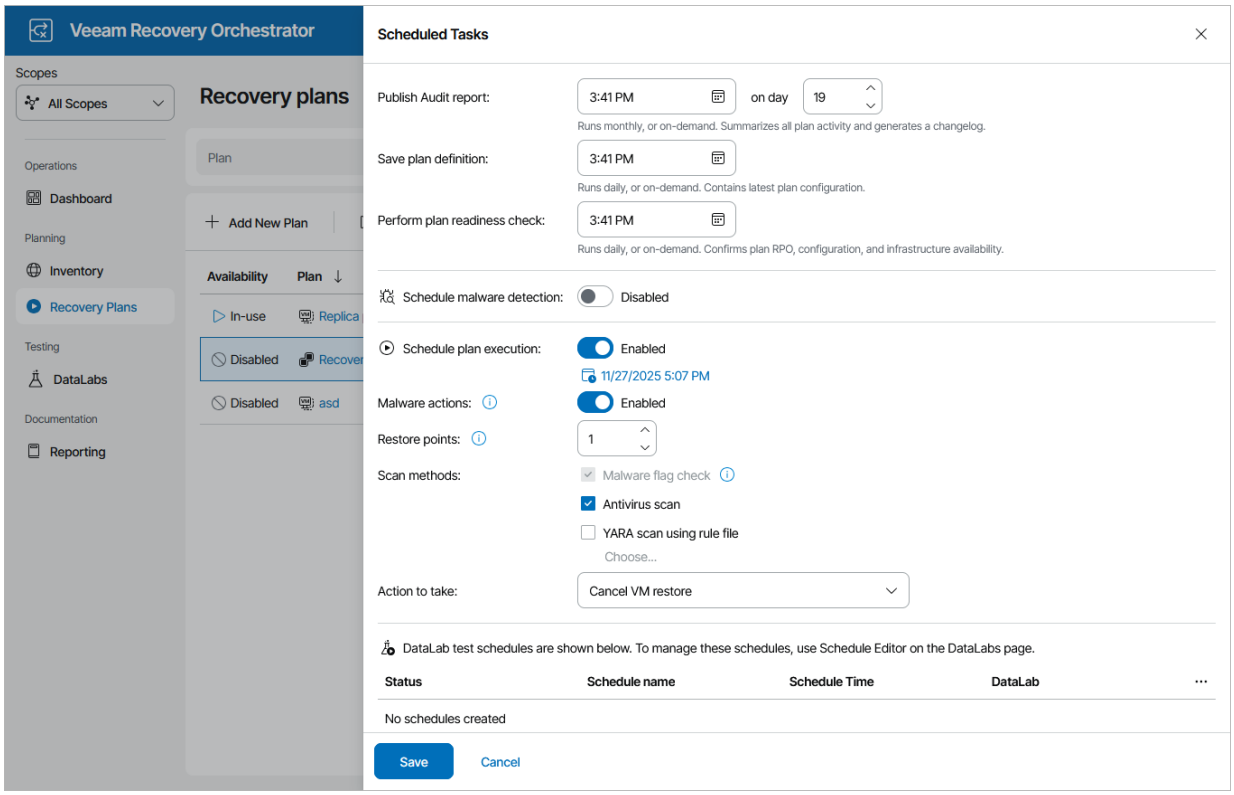

TIP

You can disable a configured schedule if you no longer need it. To do that, set the **Schedule plan execution** toggle to *Disabled* in the **Scheduled Tasks** window.

Running Restore

The **Run** action causes machines in a plan to recover from their backup files. For more information on the data recovery process, see the Veeam Backup & Replication User Guide, section Data Recovery.

> **IMPORTANT**
>
> For Orchestrator to be able to perform restore in a clean roomhttps://helpcenter.veeam.com/docs/vro/userguide/infrastructure_clean_room.html?ver=13 in case the production Veeam Backup & Replication server becomes unavailable, you must do the following before running a restore plan:
>
> 1. Add the production backup repository to the embedded Veeam Backup & Replication server as described in the Veeam Backup & Replication User Guide, section Backup Repositories.
> 2. Rescan the backup repository as described in the Veeam Backup & Replication User Guide, section Rescanning Backup Repositories. Note that this process may take several minutes to complete.
> 3. Run the restore plan.

To run a restore plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Run**.
3. In the **Run Plan** window, do the following:
   a. For security purposes, retype your password and click **Next**.

      You must also select the **Force-enable the plan** check box if you have not enabled the plan yet.

   b. In the **Recovery location** section, click the link and select a location to which inventory groups included in the plan will be restored. For a recovery location to be displayed in the list of available locations, it must be created and added to the list of inventory items available for the scope, as described in section Managing Recovery Locations.

      If the selected recovery location includes multiple hosts, datastores and networks, Orchestrator will use the round-robin algorithm to restore machines added to the plan. For more information, see the Veeam Recovery Orchestrator User Guide, section How Orchestrator Places VMs During Restore.

   c. In the **Restore point** section, choose a restore point that will be used to recover machines.

      Keep in mind that recovering data from the archive tier is not supported. If you select the **Most recent** option, make sure to choose a restore point that is stored in either the capacity or the performance tier. For more information on Veeam Backup & Replication tiering options, see the Veeam Backup & Replication User Guide, section Scale-Out Backup Repository.

   d. In the **Malware detection** section, choose whether you want to check restore points created for machines included in the plan for malware flags. When restoring to a VMware vSphere environment, you can also decide whether you want to scan these restore points with antivirus software, YARA rules or both.

      By default, Orchestrator checks the most recent restore point on each machine. If all the restore points are infected, Orchestrator restores the machine to the selected recovery location without connecting it to any network. However, you can instruct Orchestrator to halt the plan and cancel the restore operation if no clean restore point is found.

      For more information on how Orchestrator performs malware scan, see the Veeam Recovery Orchestrator User Guide, section Overview.

   e. Review configuration information and click **Run**.

TIPS

- If you select a Hyper-V recovery location and if the restore plan contains a VM that still exists in this location, Orchestrator will make an attempt to replace the original VM with the restored one. If the target VM is powered on, the restore operation will complete with an error. To work around the issue, power all the plan VMs off — and then try running the plan again.
- You can also scan a restore plan for possible malware without running the plan. To do that, follow the instructions provided in section Scanning Recovery Plans.

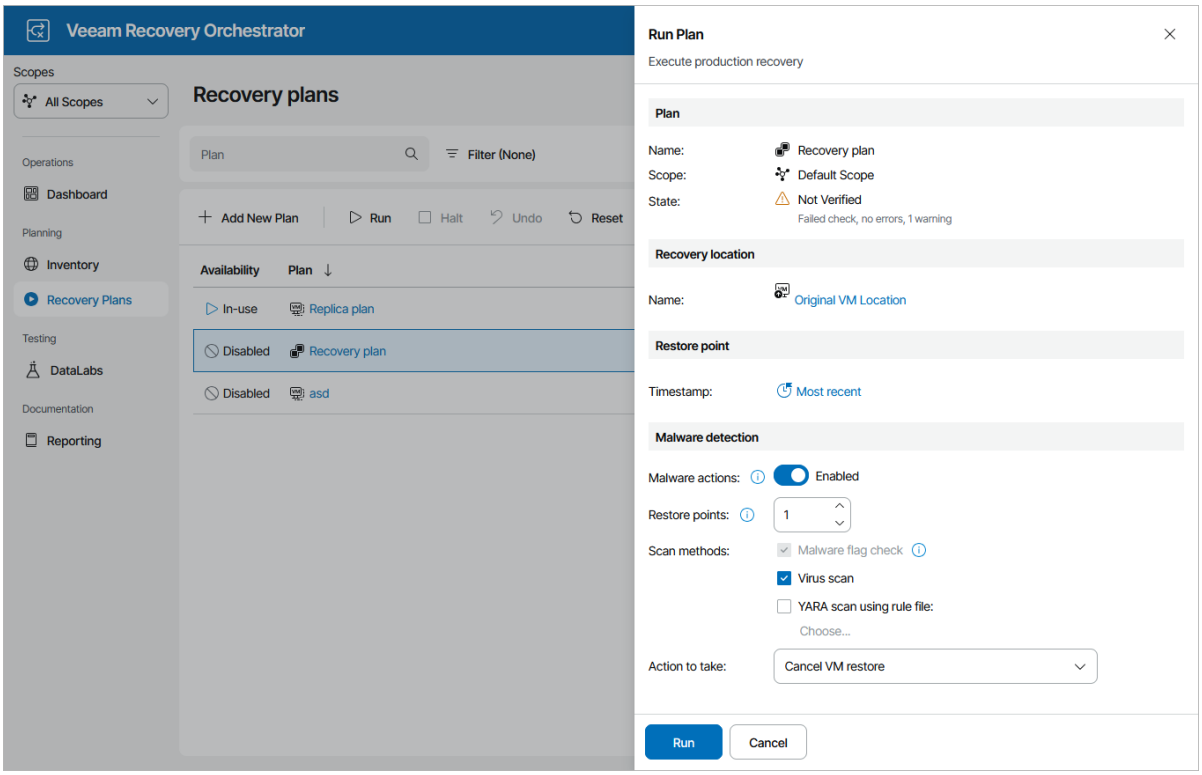

The plan goal is to reach the *RESTORED* state. If any critical error is encountered, the plan will stop with the *HALTED* state. To learn how to work with *HALTED* restore plans, see Managing Halted Plans.

Halting Restore

The **Halt** action interrupts plan execution. Any steps currently executing will be completed, then the plan will enter the *HALTED* state. To learn how to work with *HALTED* restore plans, see Managing Halted Plans.

To stop a running restore plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Halt**.
3. In the **Halt Plan Execution** window, do the following:
    a. For security purposes, retype your password and click **Next**.

b. Review configuration information and click **Halt**.

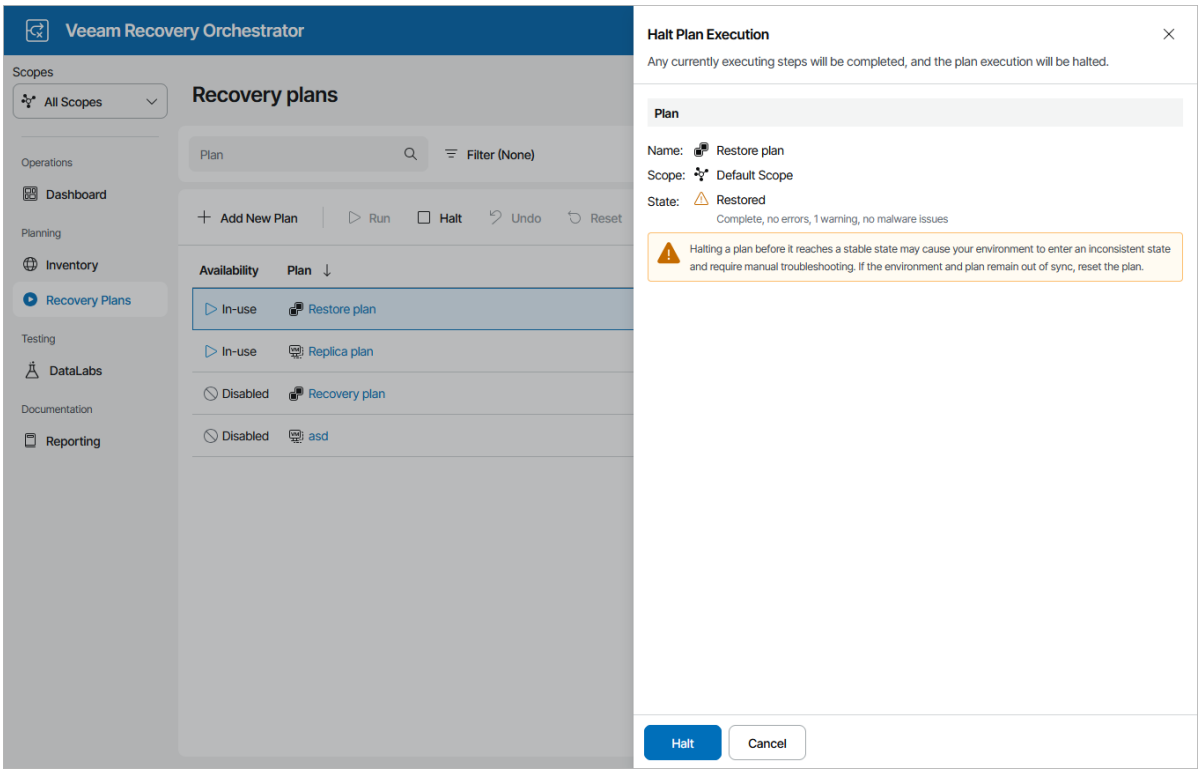

Resetting Restore Plans

If a restore plan becomes inconsistent with the virtual environment, you can reset the plan. This will return the plan to the *DISABLED* state, without making any changes to the external virtual infrastructure.

To reset a restore plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Reset**.
3. In the **Reset Plan** window, do the following:
   a. For security purposes, retype your password and click **Next**.
   b. Select the **Run a quick check after reset is complete** check box to run a readiness check after the reset.

c. Review configuration information and click **Reset**.

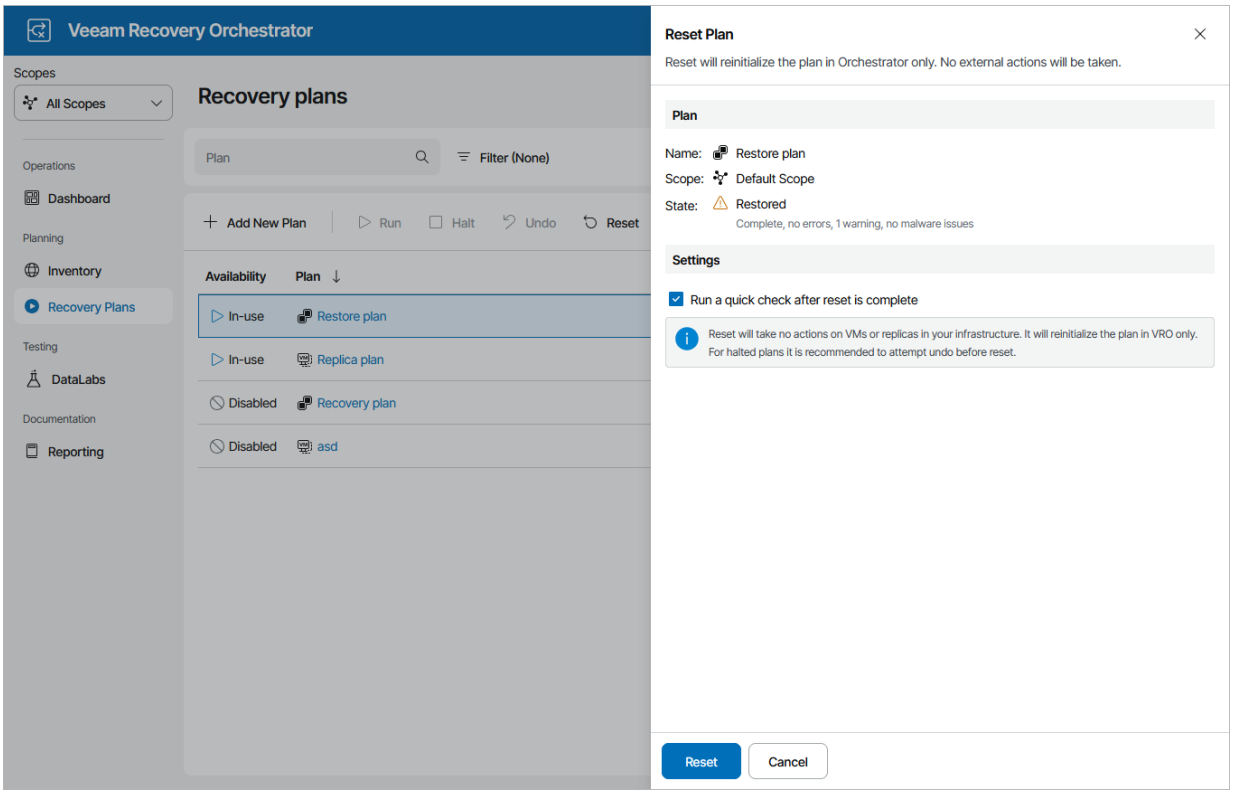

Halting Plan Testing

The **Halt** action interrupts plan testing. You may need to halt plan testing, for example, if you need to fix some environment-related issues and then proceed with testing later (in this case, recovered VMs will still continue to run). Or you may need to stop the testing process completely, for example, if you no longer need to test the selected restore plan (in this case, recovered VMs will be deleted).

To halt testing of a restore plan:

1. Navigate to **Recovery Plans**.
2. Click the plan name to switch to the **Plan Details** page.
3. On the **Plan Details** page, select the plan and click **Halt**.
4. In the **Halt Plan Test** window, click **Halt Plan Test** to confirm the action.

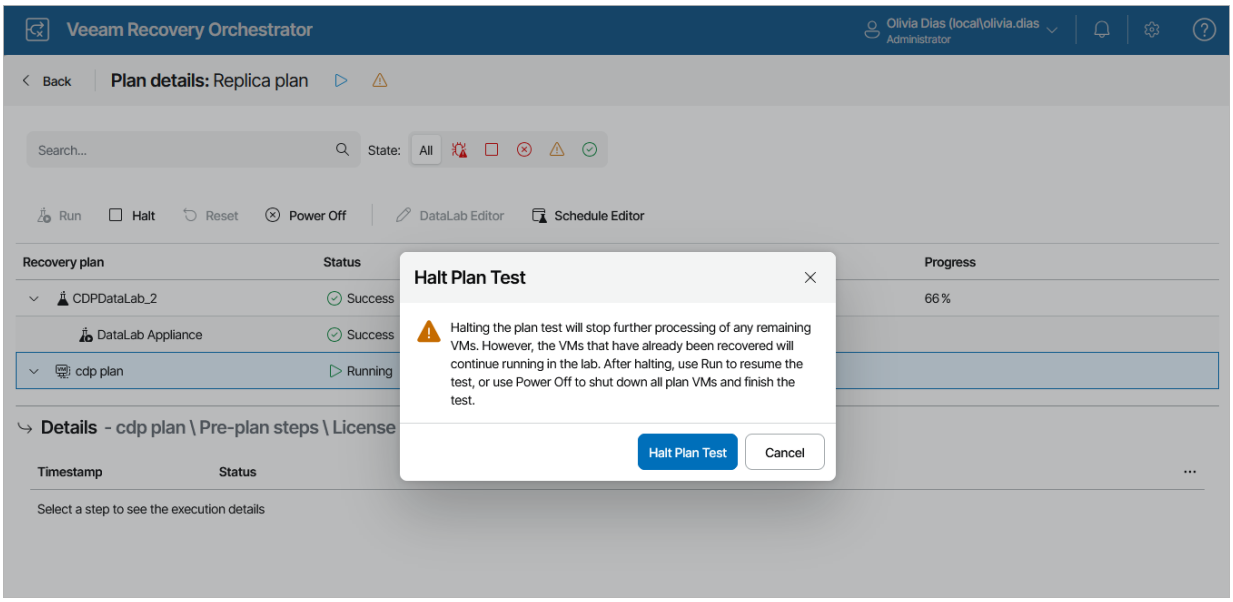

To cancel testing of a restore plan:

1. Navigate to **Recovery Plans**.
2. Click the plan name to switch to the **Plan Details** page.
3. On the **Plan Details** page, select the plan and click **Power Off**.
4. In the **Power Off Plan Testing** window, click **Power Off** to confirm the action.

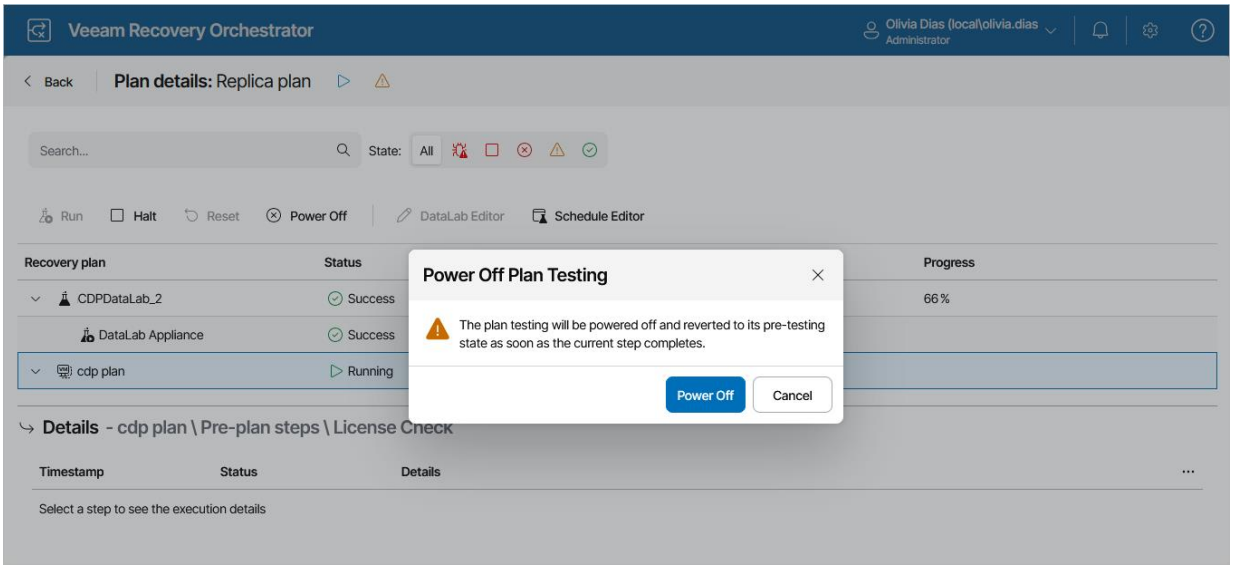

Managing Halted Restore Plans

If a critical step fails for a machine from a critical inventory group, the plan may enter the *HALTED* state. To troubleshoot reasons why a plan failed, use the **Plan Execution Report** generated as soon as the currently performed action completes. For more information on how to track plan performance history, see Viewing Plan Execution History.

After you eliminate the problem that caused the plan to become *HALTED*, you have the following options to resume the plan:

- Repeat the last failed step.
- Proceed to the next step.

Running Halted Restore Plans

To run a *HALTED* restore plan:

1. Navigate to **Recovery Plans**.
2. Select the halted plan and click **Run**.
3. In the **Continue Execution** window, do the following:
   a. For security purposes, retype your password and click **Next**.
   b. In the **Retry settings** step, select an option to resume plan execution.

      Choose whether you want to proceed with plan execution from the next plan step or to retry the failed step.

   c. Review configuration information and click **Continue**. The restore process will be started.

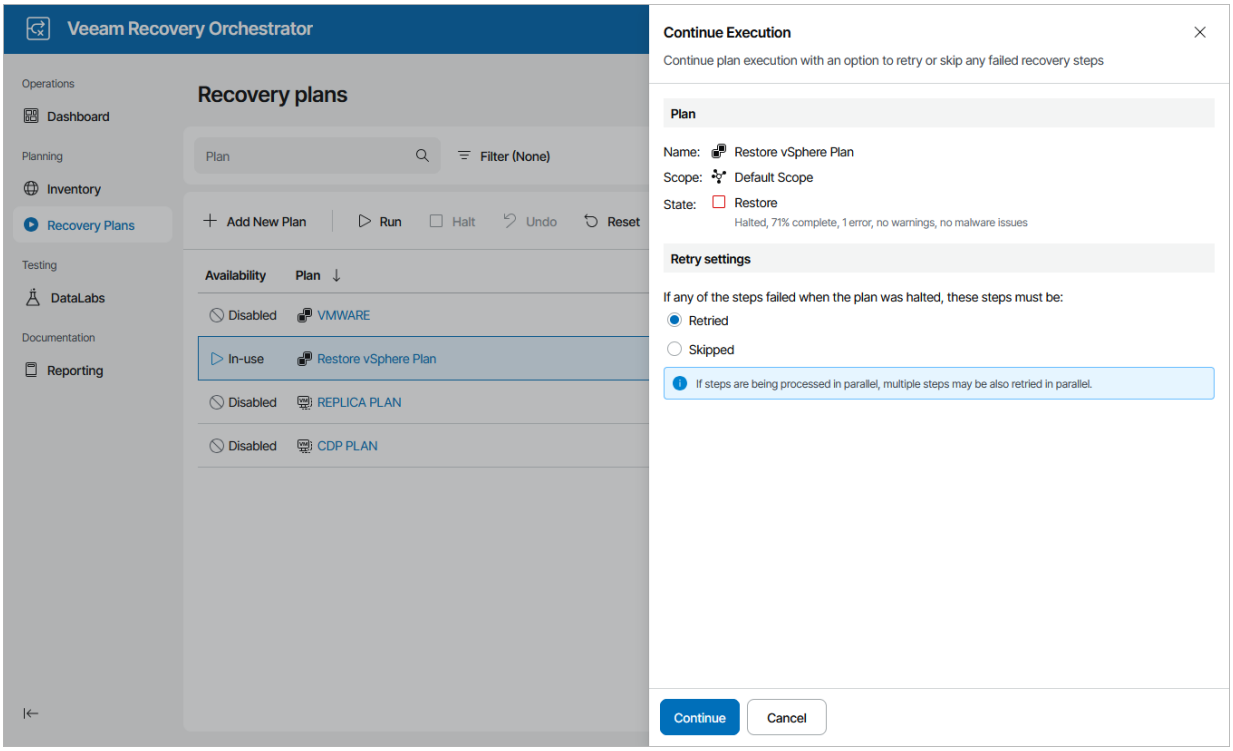

Resuming Plan Testing

To start the *HALTED* restore plan testing process:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Test**.

3. In the **Resume DataLab Test** window, choose whether you want to proceed with test execution from the next plan step or to retry the failed step, and then click **Continue**. The testing process will be started.

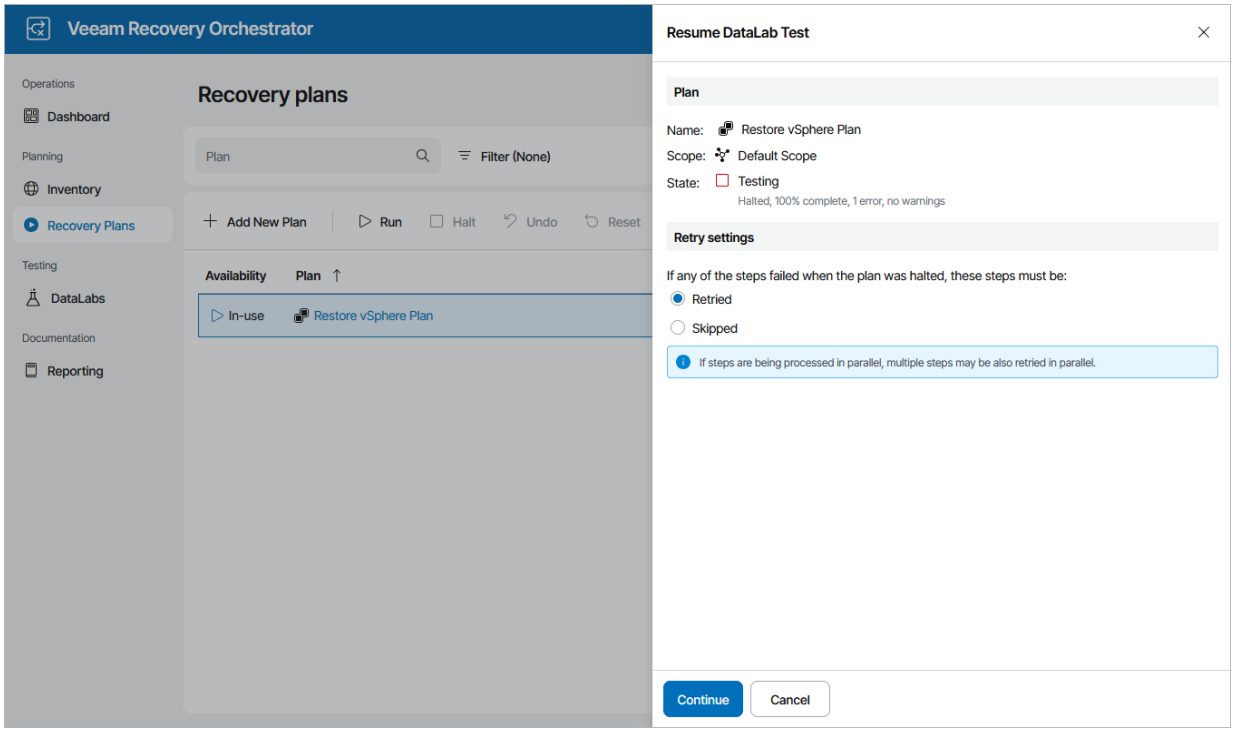

Resetting Halted Restore Plans

To reset a *HALTED* restore plan, follow the instructions provided in section Resetting Restore Plans.

> **NOTE**
>
> When you reset a restore plan, Orchestrator returns it to the *DISABLED* state without making any changes to the external virtual infrastructure. You may need to deal with any infrastructure reconfiguration manually.

Working with Storage Plans

The type of a recovery plan you create depends on whether you intend to use Orchestrator to switch to VM replicas, to restore machines from backups or backup copies, or to serve data from a destination (NetApp) or secondary (HPE) volume in case a disaster strikes.

If you want to recover volumes protected by storage replication, create a storage plan.

> **NOTE**
>
> It is recommended that you drive the creation and transfer of storage snapshots using Veeam Backup & Replication, as described in the Veeam Backup & Replication User Guide, section Integration with Storage Systems.

Creating Storage Plans

To create a storage plan:

1. Navigate to **Recovery Plans**.
2. Click **Add New Plan**.
3. Complete the **New Recovery Plan** wizard:
    a. Choose a type of the plan.
    b. Choose a storage vendor for the plan.
    c. Choose a scope for the plan.
    d. Specify a plan name and description.
    e. Specify the target RTO and RPO.
    f. Select a template for plan reports.
    g. Finish working with the wizard.

Considerations and Limitations

When you create a storage plan, keep in mind the following:

- Since Orchestrator orchestrates storage failover at the volume level, all VMs that belong to a specific datastore must be processed as part of the same storage plan.
- If a VM that you want to fail over stores its disk files on multiple datastores, make sure to include in the plan all inventory groups related to these datastores. Also, make sure to include the group with the .VMX file into the plan first, before the groups with .VMDK files.
- If the VMs that you want to fail over belong to a datastore in a VMware Storage DRS cluster, make sure to include in the plan all inventory groups related to this cluster.
- Failover to the same VMware vSphere datacenter where the source VMs reside is not supported.
- Failover of VMs that store disks on volumes protected using SnapVault is not supported.
- Failover of VMs with RDM disks is not supported.
- For datastores connected through the NFSv4.1 protocol, Orchestrator supports failover to a recovery location only in the case that target hosts included in the location have the NFSv3 export policy enabled (since the recovered datastores will be mounted to the hosts through NFSv3). For datastores connected through other protocols, no limitations apply.
- If the LUN ID of a datastore where the selected inventory groups belong is greater than 256, Orchestrator may not be able to orchestrate storage failover properly. If the LUN ID is greater than 256, make sure that your equipment supports this ID.
- For Orchestrator to be able to recover a VM correctly, the VM must have VMware Tools installed. The presence of VMware Tools is checked automatically on the vCenter Server side — for both Windows-based and Linux-based VMs. To know how to install and upgrade VMware Tools in vSphere, see this VMware KB article.

Step 1. Choose Plan Type

At the **Plan Type** step of the wizard, select the **Storage Failover** option.

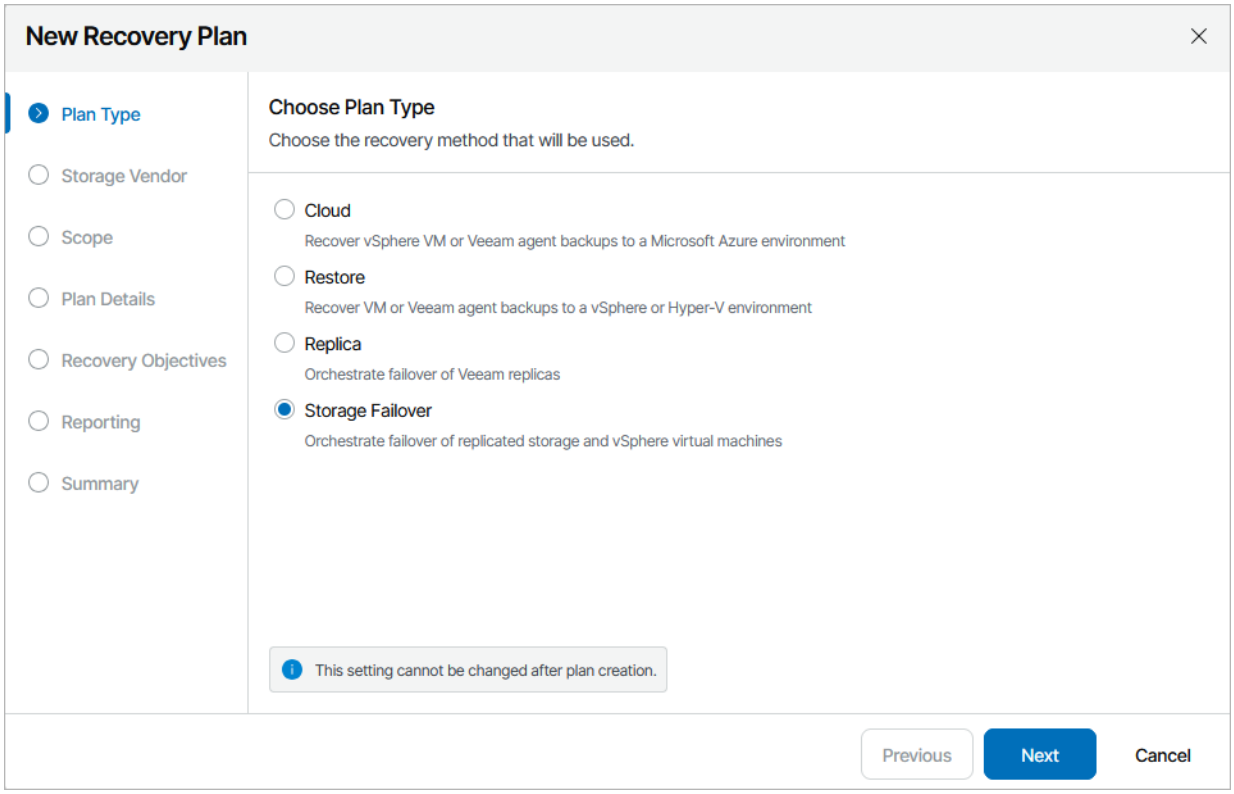

Step 2. Choose Storage Vendor

At the **Storage Vendor** step of the wizard, choose whether VMs that you plan to recover are located on datastores backed by NetApp or HPE storage systems.

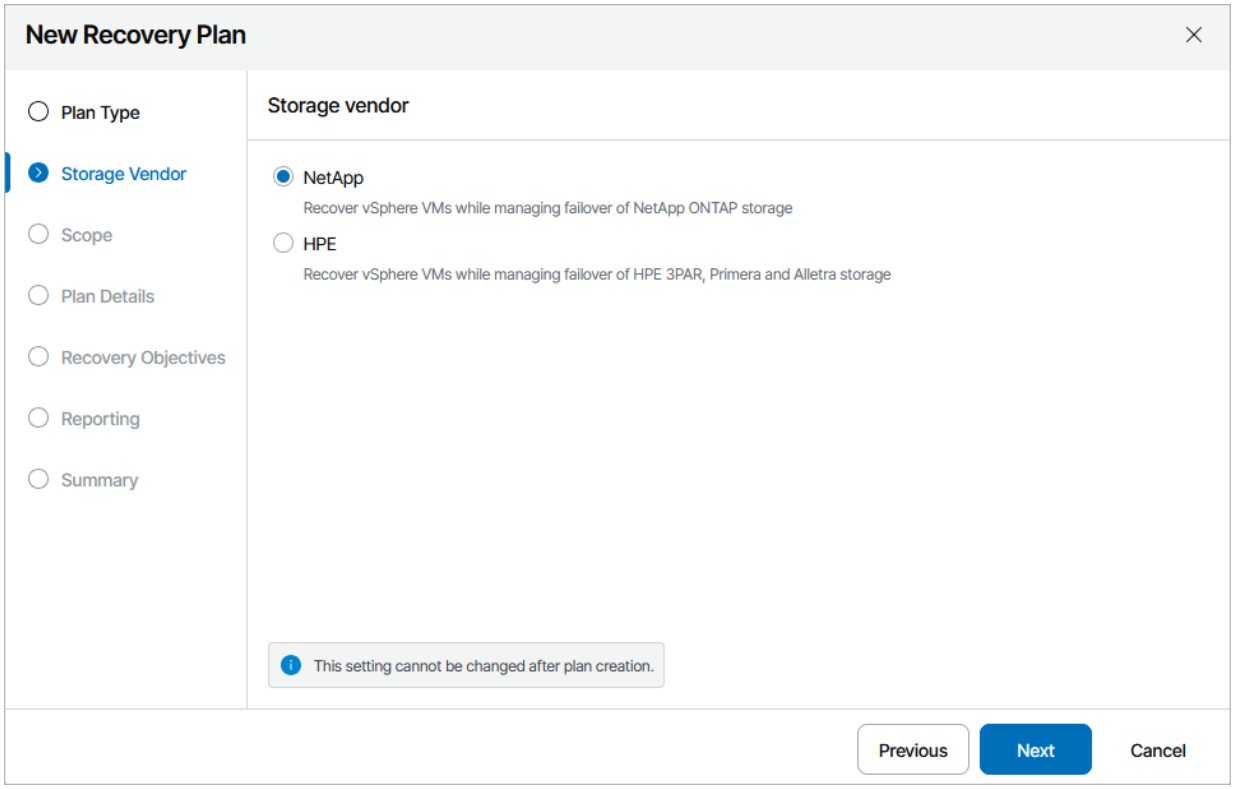

Step 3. Choose Plan Scope

At the **Scope** step of the wizard, select a scope for which you want to create the plan.

For a scope to be displayed in the list, it must be created and customized as described in section Managing Scopes.

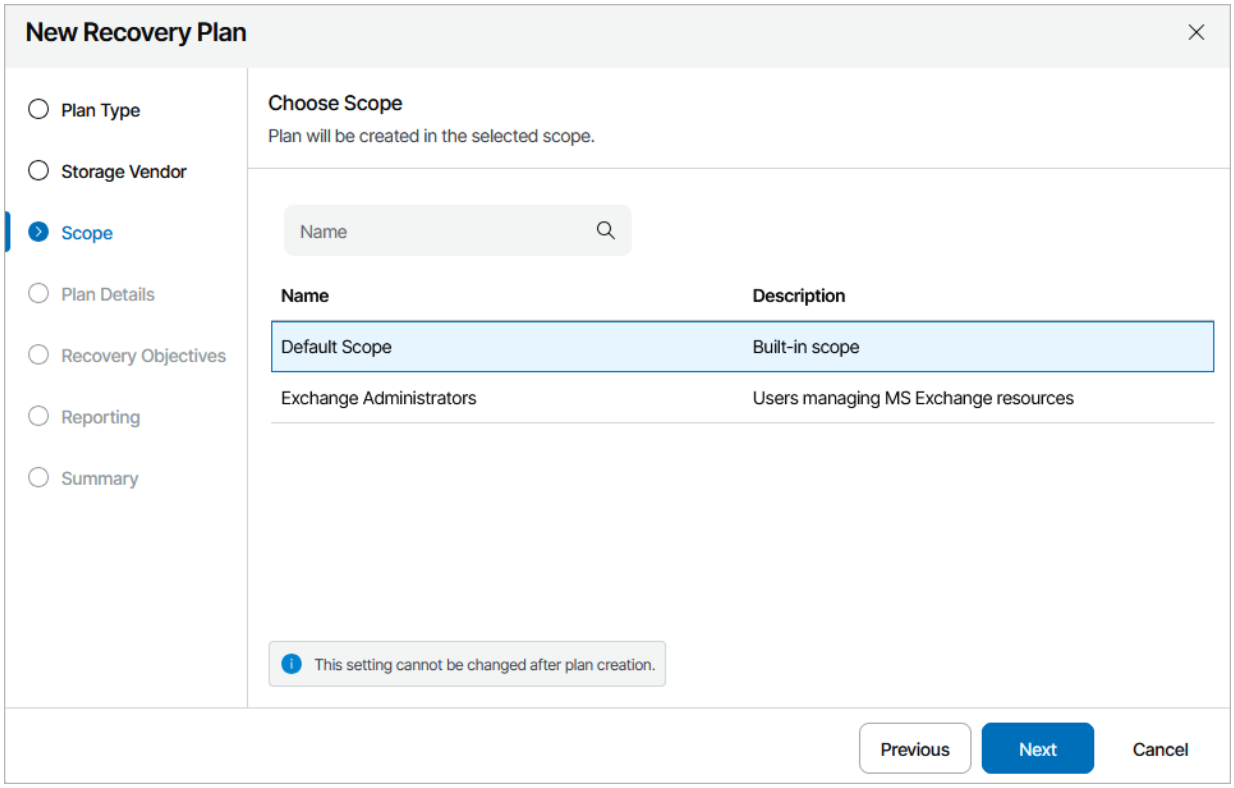

Step 4. Specify Plan Name and Description

At the **Plan Details** step of the wizard, use the **Plan name** and **Description** fields to enter a name for the new plan and to provide a description for future reference. The maximum length of the plan name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

You can also provide a contact name, email and telephone number of a person responsible for the plan.

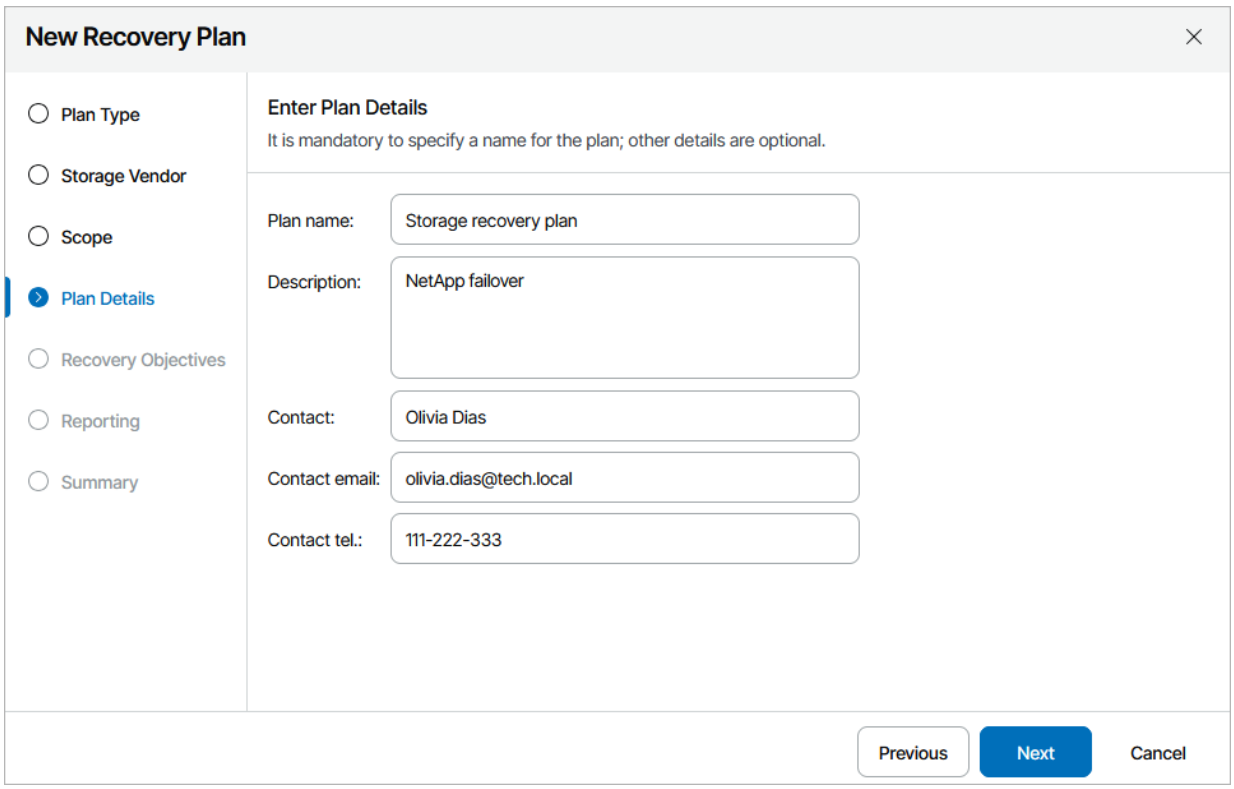

Step 5. Specify Target RTO and RPO

At the **Recovery Objectives** step of the wizard, define your Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for the plan:

- The **RPO** defines the maximum acceptable period of data loss.
- The **RTO** represents the amount of time it should take to recover from an incident.

RTO and RPO performance will be recorded in the Plan Readiness Check, Plan Execution, Plan Audit and DataLab Test reports, and you will be able to track the achieved RTO and RPO objectives for each plan on the Home Page Dashboard.

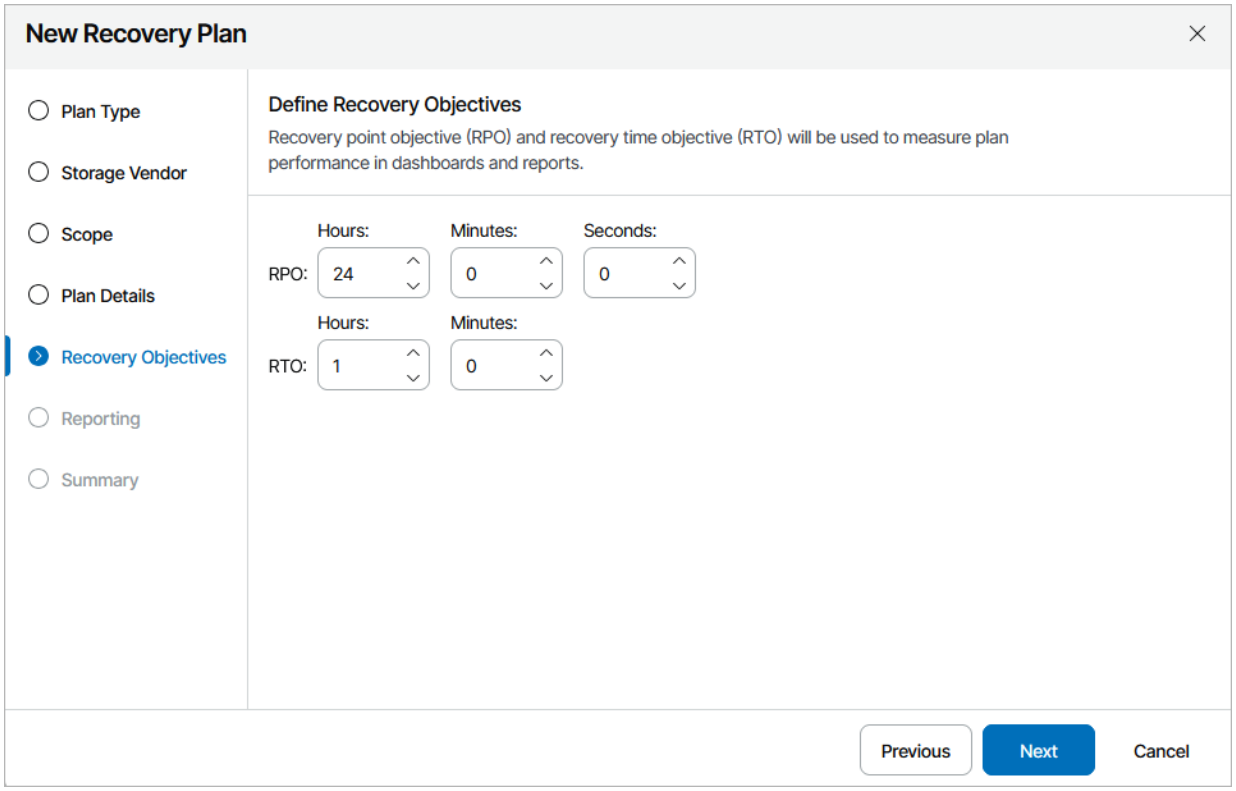

Step 6. Select Report Template

At the **Reporting** step of the wizard, select a document template that will be used as the cover page for all Orchestrator reports.

For a custom document template to be displayed in the list, it must be created and customized as described in section Generating Reports.

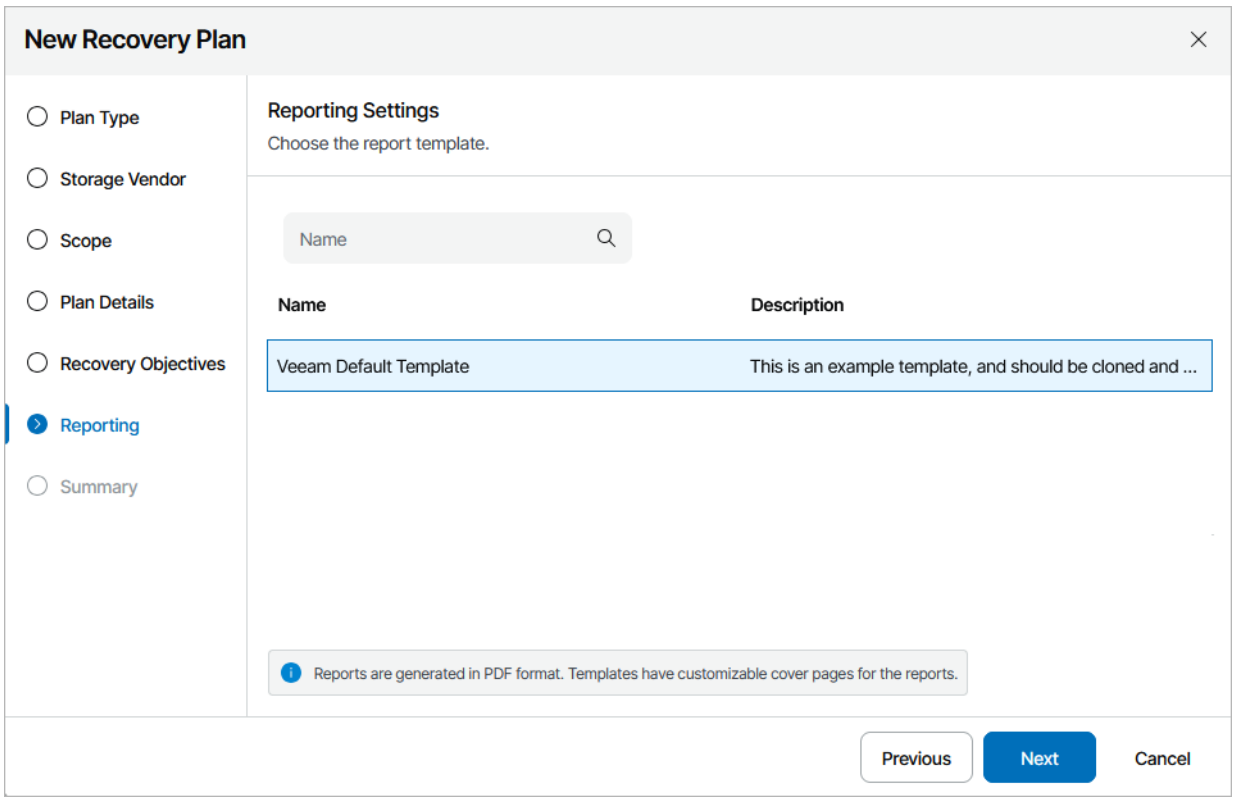

Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.

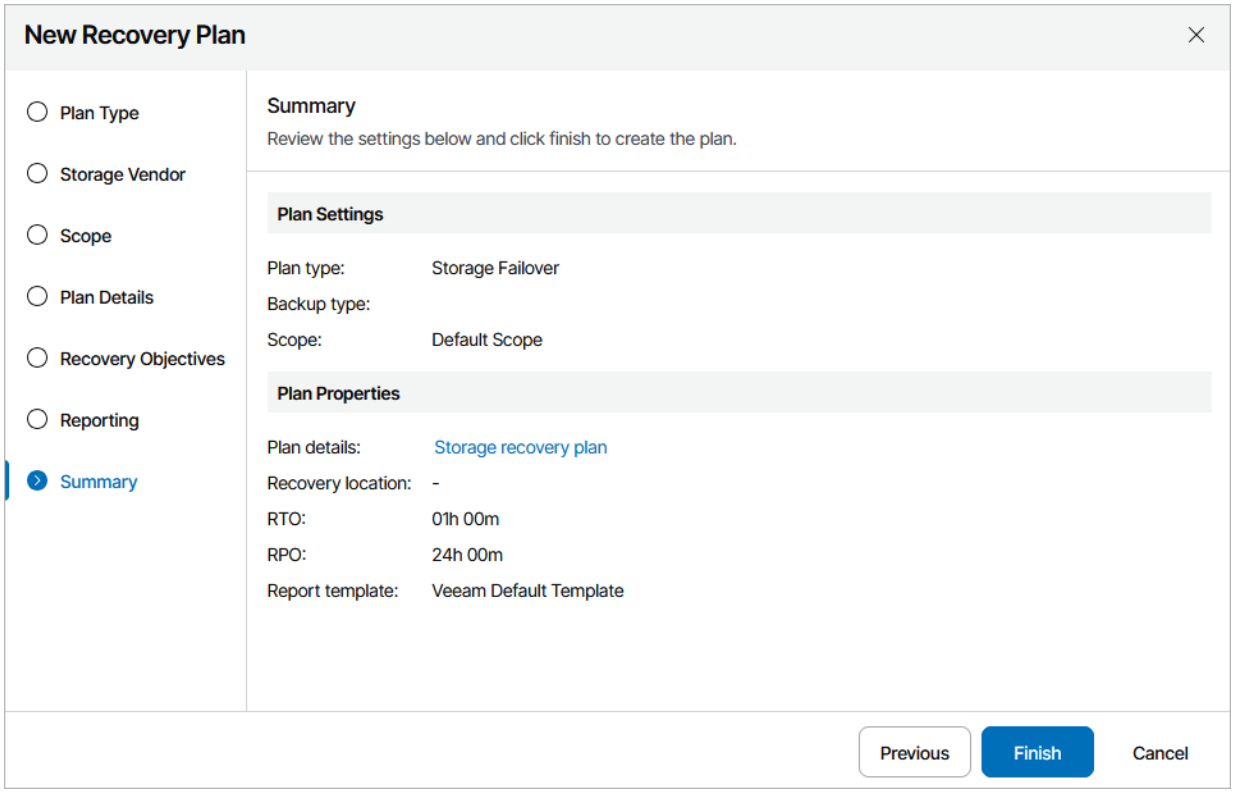

Editing Storage Plans

If you want to specify granular settings not provided in the New Recovery Plan wizard, the Orchestrator UI allows you to customize storage plans and configure the settings for groups, recovered VMs, plan steps and step parameters.

The procedures to edit replica, CDP replica, restore, storage and cloud plans are almost identical. For more information, see Editing Recovery Plans.

Testing Storage Plans

You can start on-demand plan testing and configure test scheduling for any storage plan. There is almost no difference between the procedures performed for replica, CDP replica, restore and storage plans. For more information, see Testing Recovery Plans.

Running and Scheduling Storage Plans

To run a storage plan, it must be *ENABLED*. To enable a plan:

1. Navigate to **Recovery Plans**.
2. Select the plan.
3. From the **Manage** menu, select **Properties**.

   OR-

   Right-click the plan name and select **Manage** > **Properties**.
4. Set the **Availability** toggle to *Enabled*.
5. Click **Save**.

If you do not enable a plan before you run it, the Run Plan wizard will force you to do that as soon as you try running the plan.

> **NOTES**
>
> 1. An Orchestrator Administrator or Plan Operator can force-enable a plan in the **Run Plan** wizard. However, a Plan Operator will not be able to run a disabled storage plan.
>
>   For more information on roles that can be assigned to users and user groups working with the Orchestrator UI, see Managing User Accounts.
> 2. For security purposes, all 'real-world' actions associated with storage plans require password confirmation.

Scheduling Storage Failover

You can schedule a time for a storage plan to execute. Only the failover process can be scheduled — all other operations must be performed manually in the Orchestrator UI.

> **NOTE**
>
> If you configure a schedule for a storage plan, Orchestrator will not be able to trigger reverse replication to reprotect volumes included in the plan — this option is available only when you run the storage failover process manually.

To schedule a storage plan:

1. Navigate to **Recovery Plans**.
2. Select the plan. From the **Manage** menu, select **Schedule**.

   -OR-

   Right-click the plan name and select **Manage** > **Schedule**.
3. In the **Scheduled Tasks** window, do the following:
   a. Set the **Schedule plan execution** toggle to *Enabled*.
   b. Click the **Configure schedule** link and choose whether you want to run the plan on schedule or after any other plan:

- If you want to run the plan at a specific time, click the **Schedule** icon in the **Run on** field, set the desired date and time, and click **Apply**.
- If you want to run the plan after another plan, select the **Schedule after plan** check box and click **Choose plan**. Then, in the **Select Plan** window, select the necessary plan and click **Apply**.

  For a plan to be displayed in the list of available plans, it must be *ENABLED* as described in section Running and Scheduling Storage Plans.

c. Review the configuration information and click **Save**.

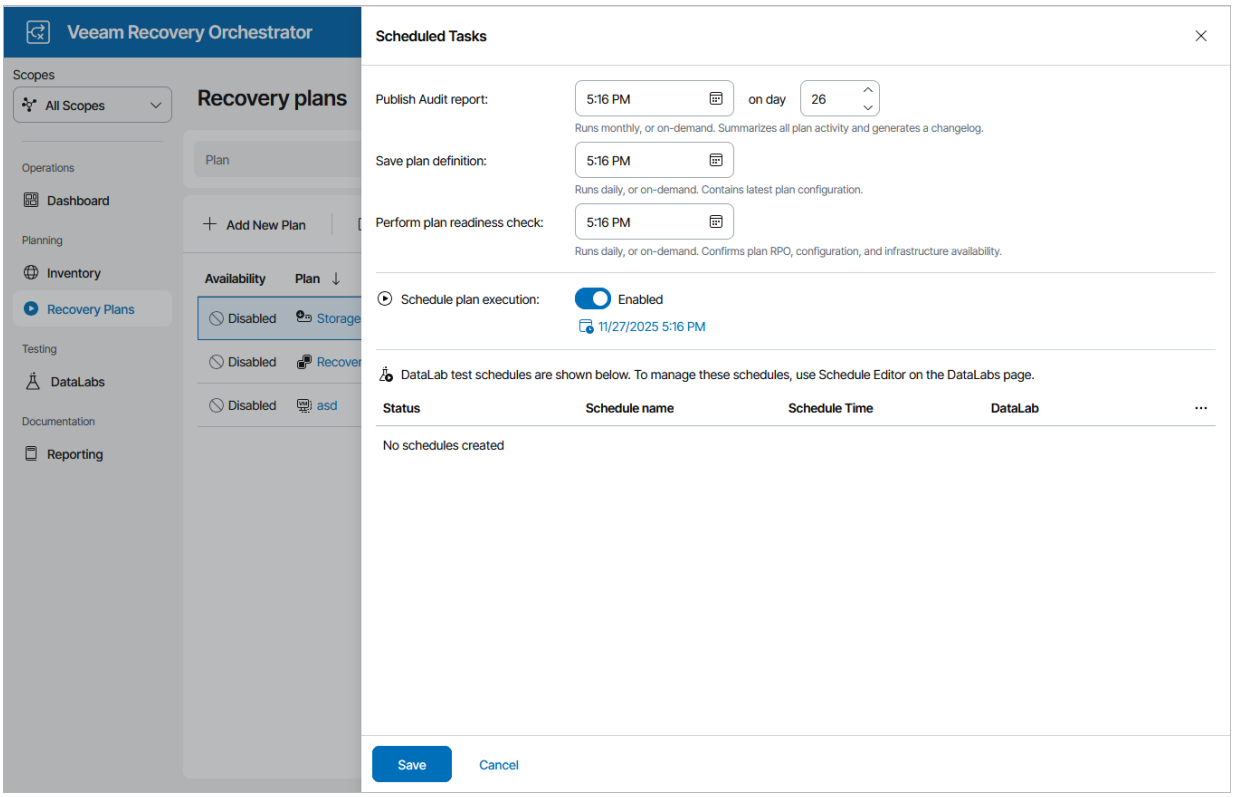

TIP

You can disable a configured schedule if you no longer need it. To do that, set the **Schedule plan execution** toggle to *Disabled* in the **Scheduled Tasks** window.

Running Storage Failover

The **Run** action causes VMs in a plan to fail over to destination (NetApp) or secondary (HPE) storage volumes. For more information on the data recovery process, see the NetApp ONTAP Documentation Center and Hewlett Packard Enterprise Support Center.

NOTE

To allow Orchestrator to complete a storage failover successfully, make sure that all the prerequisites provided in the checklist in this KB article are met.

Running NetApp Storage Failover

To run a NetApp storage plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Run**.
3. In the **Run Plan** window, do the following:
   a. For security purposes, retype your password and click **Next**.

      You must also select the **Force-enable the plan** check box if you have not enabled the plan yet.

   b. In the **Timestamp** field, select a snapshot that will be used to recover VMs.

      To choose target storage systems to be used to recover VMs, Orchestrator will analyze settings specified during the configuration of storage recovery locations. For more information, see the Veeam Recovery Orchestrator User Guide, section How Orchestrator Places VMs During Storage Failover.

> **NOTE**
>
> This setting applies only to volumes protected by asynchronous replication. If a volume is protected by synchronous replication, Orchestrator will always use the most recent replicated data. This is a limitation of the synchronous SnapMirror technology.

   c. In the **Reverse replication** field, choose whether you want Orchestrator to trigger reverse replication to reprotect volumes included in the plan. This option can be useful if you plan to fail back to the production location.

      If you select the to trigger reverse replication to reprotect the failed-over volumes, Orchestrator will add the **Protect Storage Volumes** step to the list of plan steps. This step will resynchronize the data protection relationship in the reverse direction as soon as the storage failover process completes.

      Note that when running the **Protect Storage Volumes** step, Orchestrator only triggers the reprotect operation and reports whether the step itself started successfully — Orchestrator does not check whether the operation of resynchronizing the relationship in the reverse direction completes successfully.

      For more information on ONTAP data protection, see the NetApp ONTAP Documentation Center.

> **IMPORTANT**
>
> The **Protect Storage Volumes** step will interfere with the existing jobs that use storage snapshots in Veeam Backup & Replication. Since the step reverses the source and destination roles of SnapMirror relationships, these jobs will no longer function properly after the storage failover process completes.
>
> To work around the issue, add the **Veeam Job Actions** step to the list of pre-plan steps before running the plan. To disable a job, specify its name when configuring the step parameters.

d. Review configuration information and click **Run**.

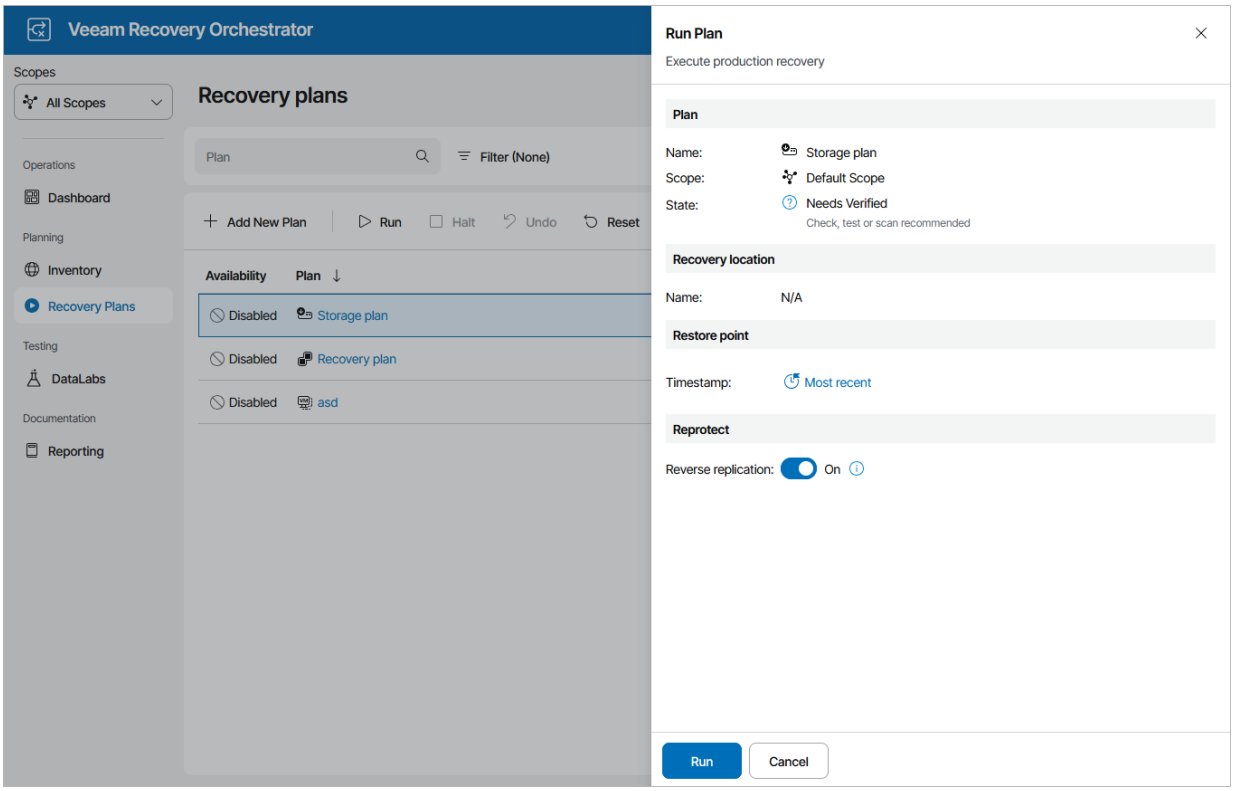

The plan goal is to reach the *FAILOVER* state. If any critical error is encountered, the plan will stop with the *HALTED* state. To learn how to work with *HALTED* storage plans, see Managing Halted Plans.

> **IMPORTANT**
>
> After the storage failover process completes, Orchestrator will leave the plan in the *IN-USE* mode. By design, this makes the results of the storage failover process accessible in the Orchestrator UI as long as required, and also prevents the plan from being modified by any automatic updates related to infrastructure changes.
>
> If you want to perform any further actions with the plan (for example, to test the plan, to run readiness checks or to execute the plan again), reset the plan as described in section Resetting Storage Plans.

Running HPE Storage Failover

> **IMPORTANT**
>
> It is recommended that you do not enable the auto synchronize option for the remote copy group as it may cause performance issues during the failover process. For more information on the auto synchronize option, see Hewlett Packard Enterprise Support Center.

To run an HPE storage plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Run**.

3. In the **Run Plan** window, do the following:

   a. For security purposes, retype your password and click **Next**.

      You must also select the **Force-enable the plan** check box if you have not enabled the plan yet.

   b. Review configuration information and click **Finish**.

> **NOTE**
>
> For HPE storage plans, the **Run Plan** wizard does not offer you to choose a restore point that will be used to recover VMs. By design, Orchestrator will always use the most recent replicated data. This is a limitation of HPE storage systems.

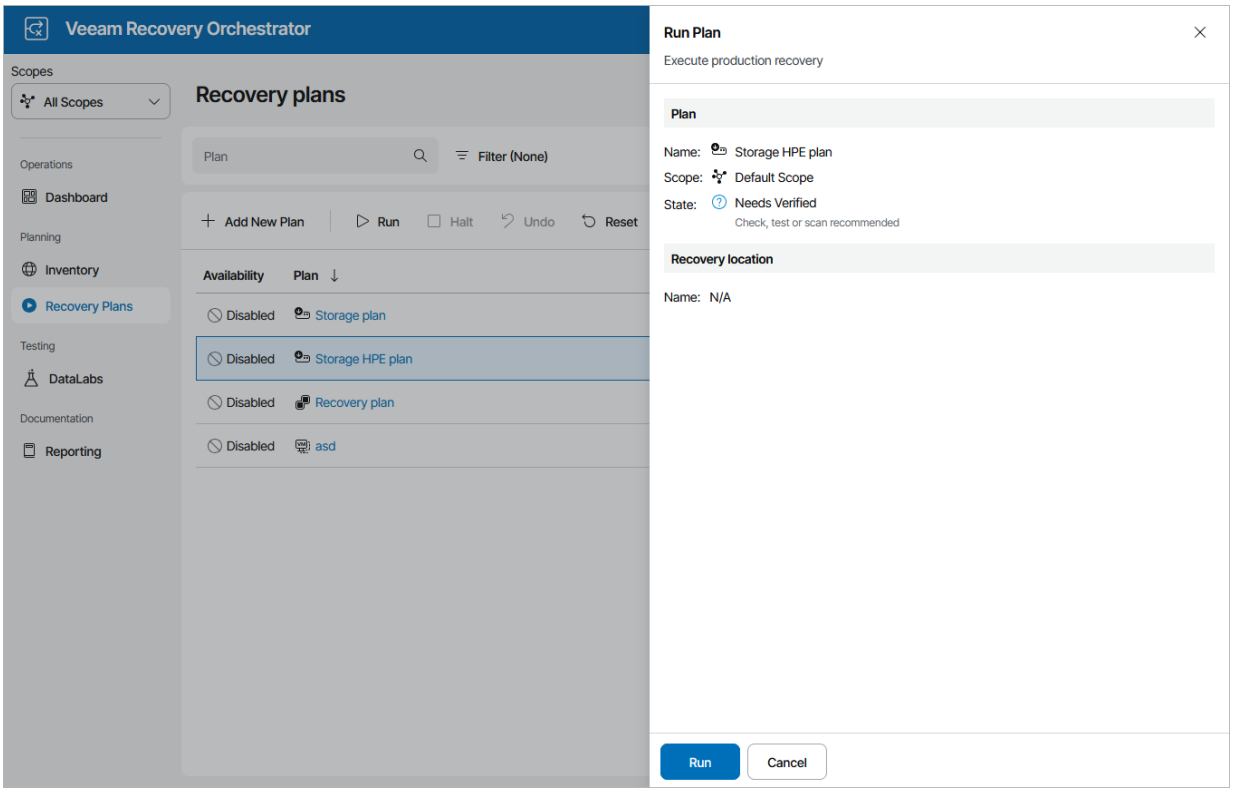

The plan goal is to reach the *FAILOVER* state. If any critical error is encountered, the plan will stop with the *HALTED* state. To learn how to work with *HALTED* storage plans, see Managing Halted Plans.

> **IMPORTANT**
>
> After the storage failover process completes, Orchestrator will leave the plan in the *IN-USE* mode. By design, this makes the results of the storage failover process accessible in the Orchestrator UI as long as required, and also prevents the plan from being modified by any automatic updates related to infrastructure changes.
>
> If you want to perform any further actions with the plan (for example, to test the plan, to run readiness checks or to execute the plan again), reset the plan as described in section Resetting Storage Plans.

Halting Storage Failover

The **Halt** action interrupts plan execution. Any steps currently executing will be completed, then the plan will enter the *HALTED* state. To learn how to work with *HALTED* storage plans, see Managing Halted Plans.

To stop a running storage plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Halt**.
3. In the **Halt Plan Execution** window, do the following:
   a. For security purposes, retype your password and click **Next**.
   b. Review configuration information and click **Halt**.

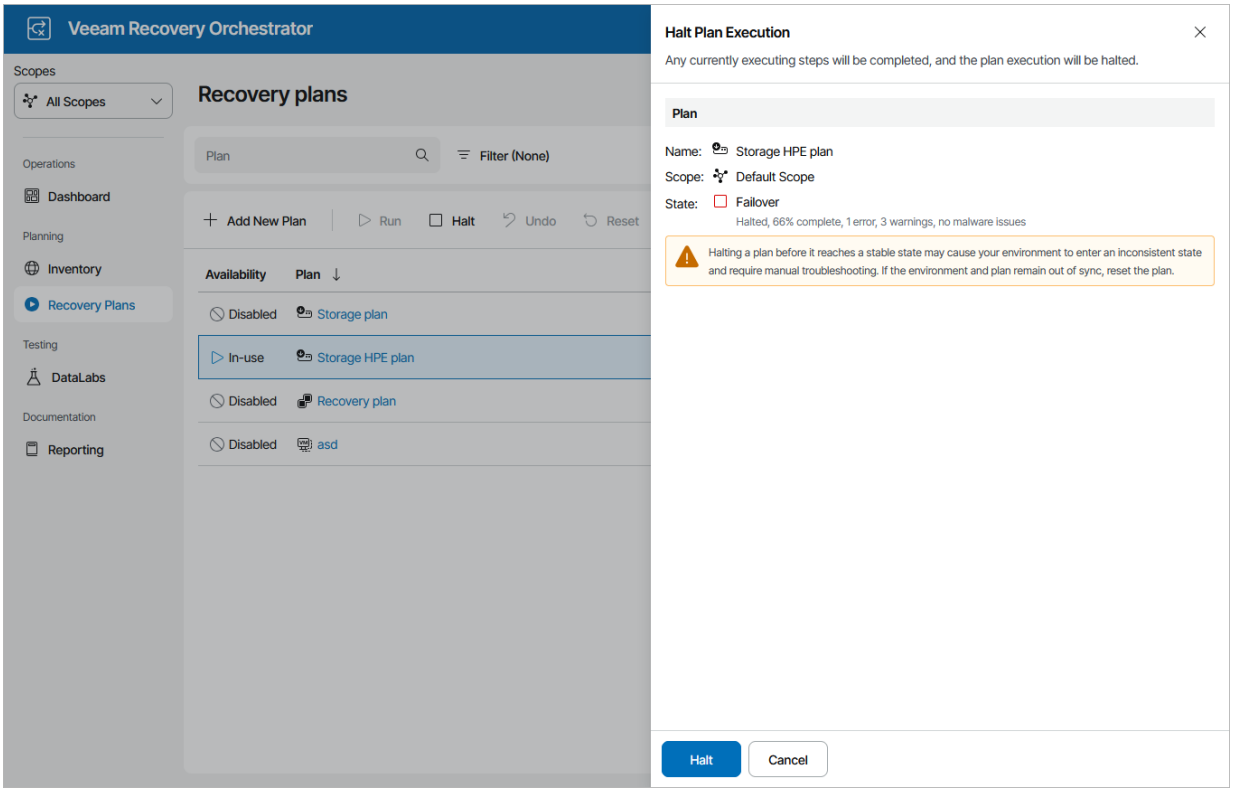

Resetting Storage Plans

After you run a storage plan and it acquires the *FAILOVER* state, you must reset the plan if you wish to run it again (for example, to perform failback). The **Reset** action returns the plan to the *DISABLED* state and updates the Orchestrator database to reflect the changes made to the location of VMs included in the plan. The configuration of plan steps and their parameter settings in this case remain the same.

You may also require to reset a storage plan if the plan becomes inconsistent with the virtual environment. This will return the plan to the *DISABLED* state, without making any changes to the external virtual infrastructure.

To reset a storage plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Reset**.
3. In the **Reset Plan** window, do the following:
   a. For security purposes, retype your password and click **Next**.
   b. Select the **Run a quick check after reset is complete** check box to run a readiness check after the reset.

c. Review configuration information and click **Reset**.

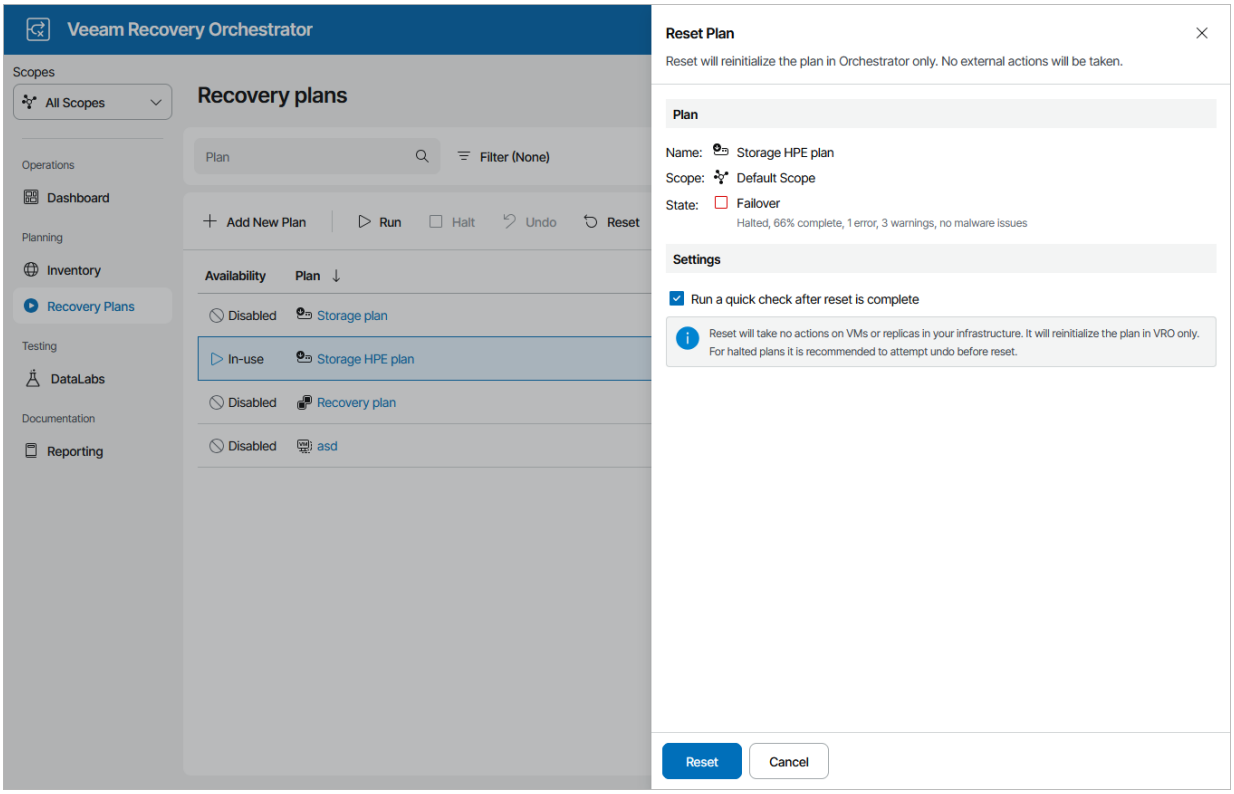

Halting Plan Testing

The **Halt** action interrupts plan testing. You may need to halt plan testing, for example, if you need to fix some environment-related issues and then proceed with testing later (in this case, recovered VMs will still continue to run). Or you may need to stop the testing process completely, for example, if you no longer need to test the selected storage plan (in this case, recovered VMs will be deleted).

To halt testing of a storage plan:

1. Navigate to **Recovery Plans**.
2. Click the plan name to switch to the **Plan Details** page.
3. On the **Plan Details** page, select the plan and click **Halt**.
4. In the **Halt Plan Test** window, click **Halt Plan Test** to confirm the action.

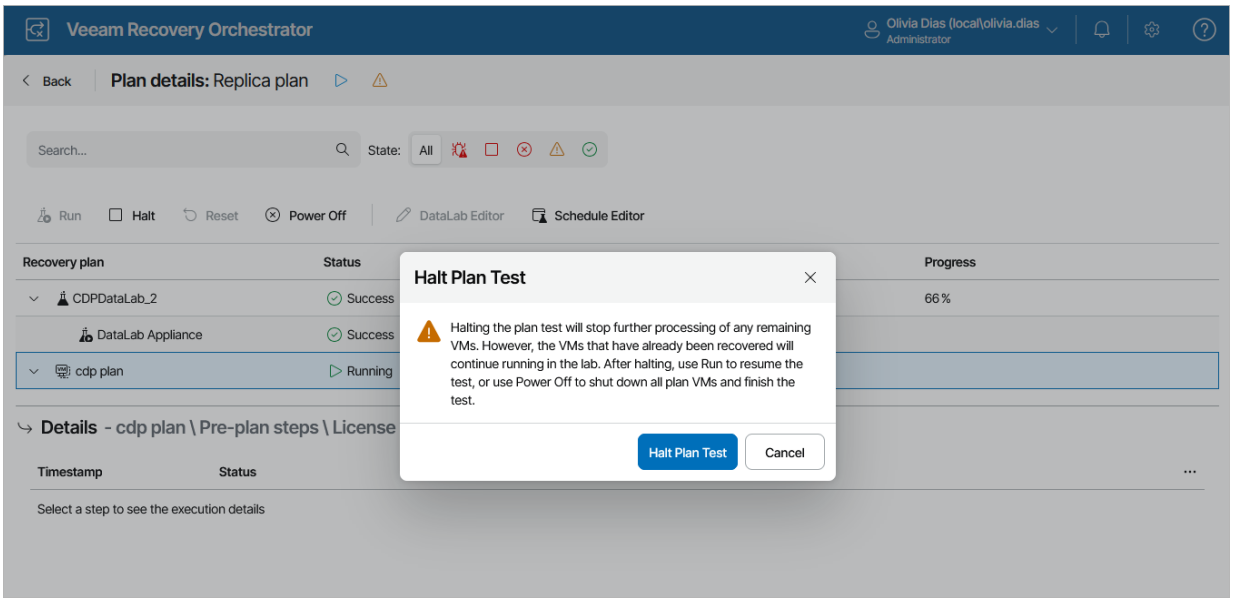

To cancel testing of a storage plan:

1. Navigate to **Recovery Plans**.
2. Click the plan name to switch to the **Plan Details** page.
3. On the **Plan Details** page, select the plan and click **Power Off**.
4. In the **Power Off Plan Testing** window, click **Power Off** to confirm the action.

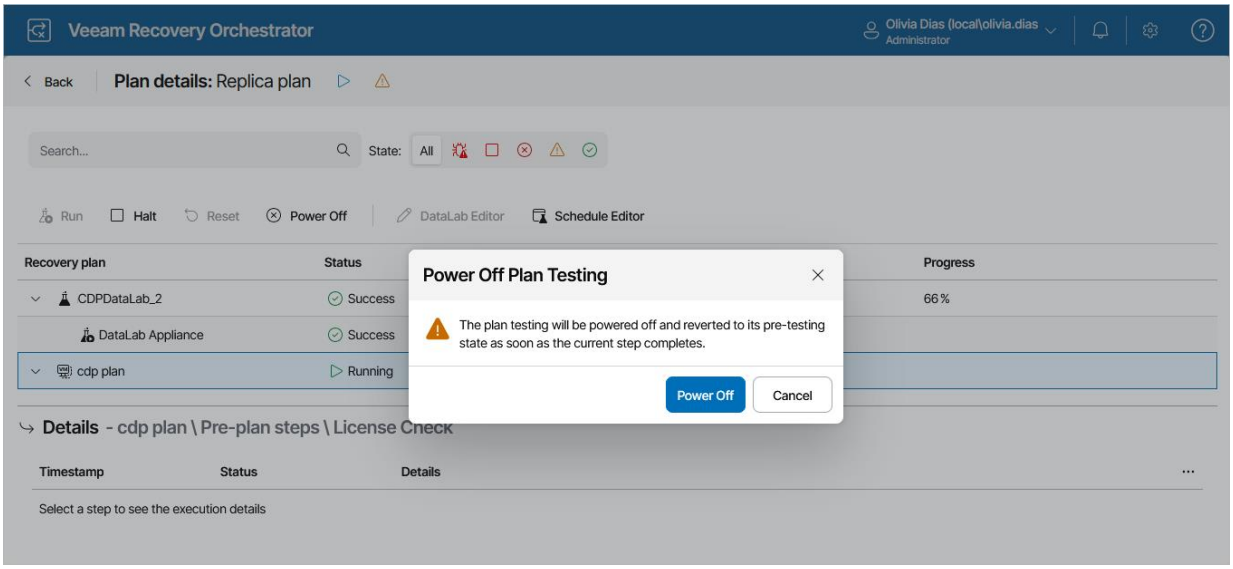

Managing Halted Storage Plans

If a critical step fails for a VM from a critical inventory group, the plan may enter the *HALTED* state. To troubleshoot reasons why a plan failed, use the Plan Execution History Report generated as soon as the currently performed action completes. For more information on how to track plan performance history, see Viewing Plan Execution History.

After you eliminate the problem that caused the plan to become *HALTED*, you have the following options to resume the plan:

- Repeat the last failed step.
- Proceed to the next step.

Running Halted Storage Plans

To run a *HALTED* restore plan:

1. Navigate to **Recovery Plans**.
2. Select the halted plan and click **Run**.
3. In the **Continue Execution** window, do the following:
   a. For security purposes, at the **Credentials** step, retype your password.
   b. In the **Retry settings** step, select an option to resume plan execution.

      Choose whether you want to proceed with plan execution from the next plan step or to retry the failed step.

> **NOTE**
>
> If you select the **Retried** option, Orchestrator will execute the **Storage Failover** step again only in case the plan halts when trying to execute the **Register Replica VM (Storage)** step. For more information on steps performed by Orchestrator, see Appendix A. Recovery Plan Steps.

c. Review configuration information and click **Continue**. The failover process will be started.

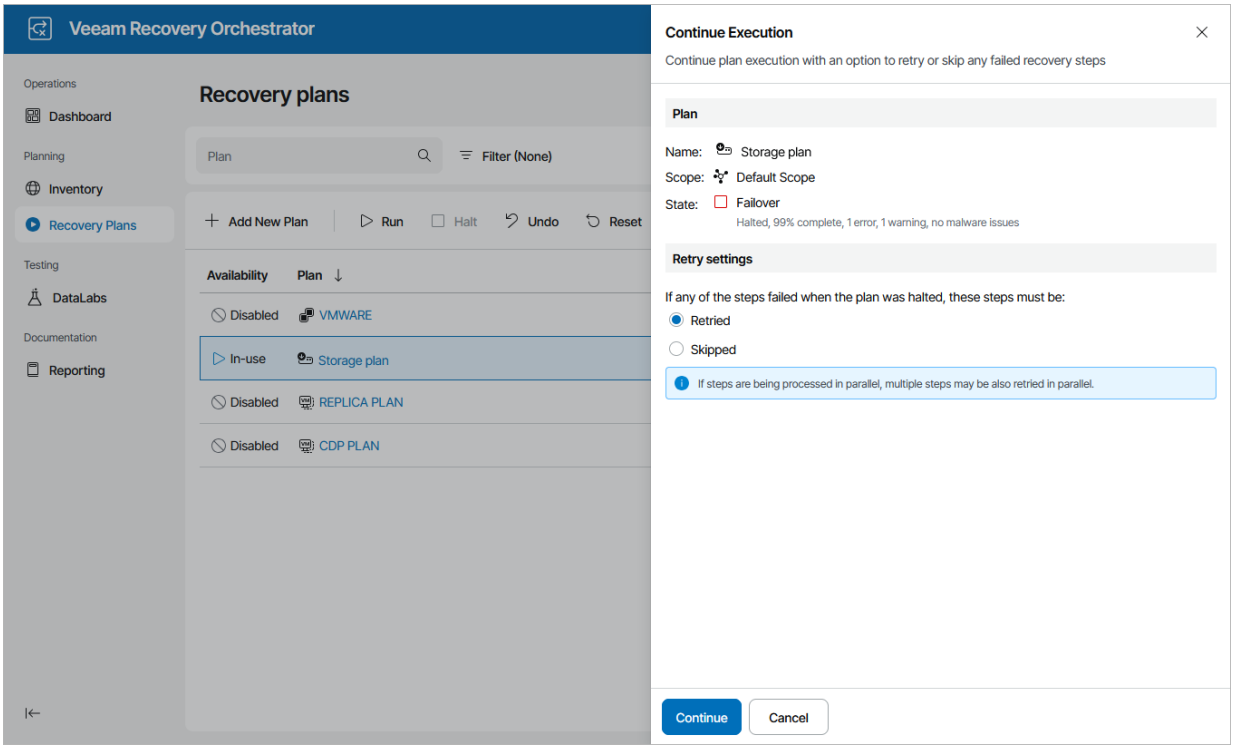

Resuming Plan Testing

To start the *HALTED* storage plan testing process:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Test**.

3. In the **Resume DataLab Test** window, choose whether you want to proceed with test execution from the next plan step or to retry the failed step, and then click **Continue**. The testing process will be started.

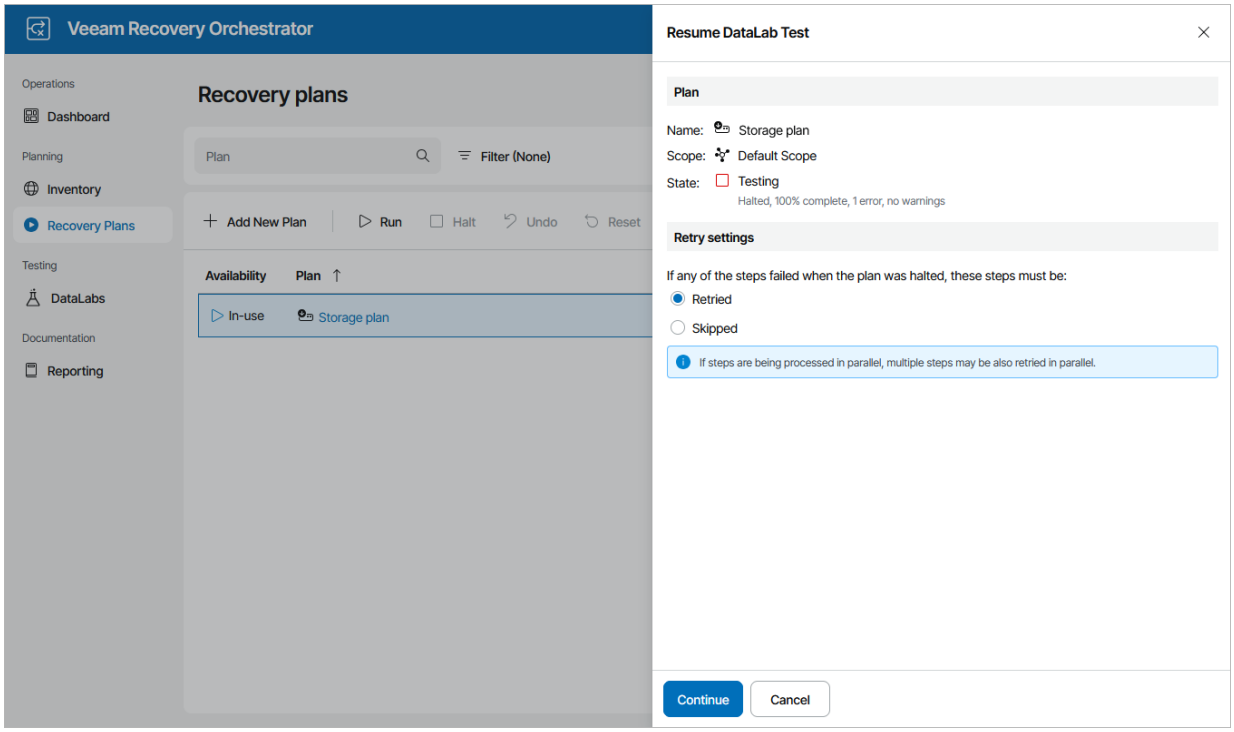

Undoing Halted Storage Plans

> **NOTE**
>
> This action is currently not supported for HPE storage systems.

To perform an undo operation for a *HALTED* storage plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Undo**.
3. In the **Undo** window, do the following:
    a. For security purposes, retype your password and click **Next**.

b. Review configuration information and click **Undo**. The failover process will be started.

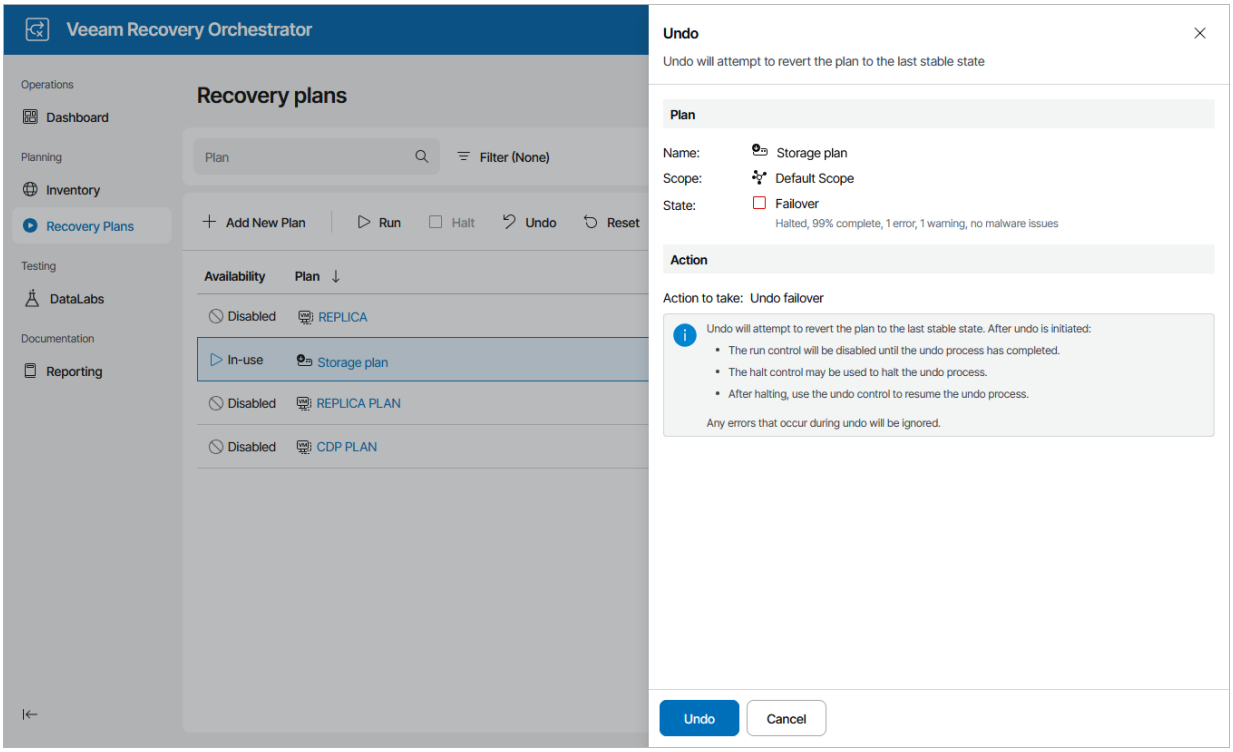

If a plan repeatedly enters the *HALTED* state due to misconfiguration or changes in the external environment, the only option left may be to RESET the plan.

Resetting Halted Storage Plans

To reset a *HALTED* storage plan, follow the instructions provided in section Resetting Storage Plans.

> **NOTE**
>
> When you reset a storage plan, Orchestrator returns it to the *DISABLED* state without making any changes to the external virtual infrastructure. You may need to deal with any infrastructure reconfiguration manually.

Working with Cloud Plans

The type of a recovery plan you create depends on whether you intend to use Orchestrator to switch to VM replicas, to restore machines from backups or backup copies, or to serve data from a destination (NetApp) or secondary (HPE) volume in case a disaster strikes.

If you want to recover machines from vSphere VM and Veeam agent backups to a cloud environment, create a cloud plan.

Creating Cloud Plans

To create a cloud plan:

1. Navigate to **Recovery Plans**.
2. Click **Add New Plan**.
3. Complete the **New Recovery Plan** wizard:
    a. Choose a type of the plan.
    b. Choose a type of backup files.
    c. Choose a scope for the plan.
    d. Specify a plan name and description.
    e. Specify the target RTO and RPO.
    f. Choose a recovery location.
    g. Select a template for plan reports.
    h. Finish working with the wizard.

Step 1. Choose Plan Type

At the **Plan Type** step of the wizard, select the **Cloud** option.

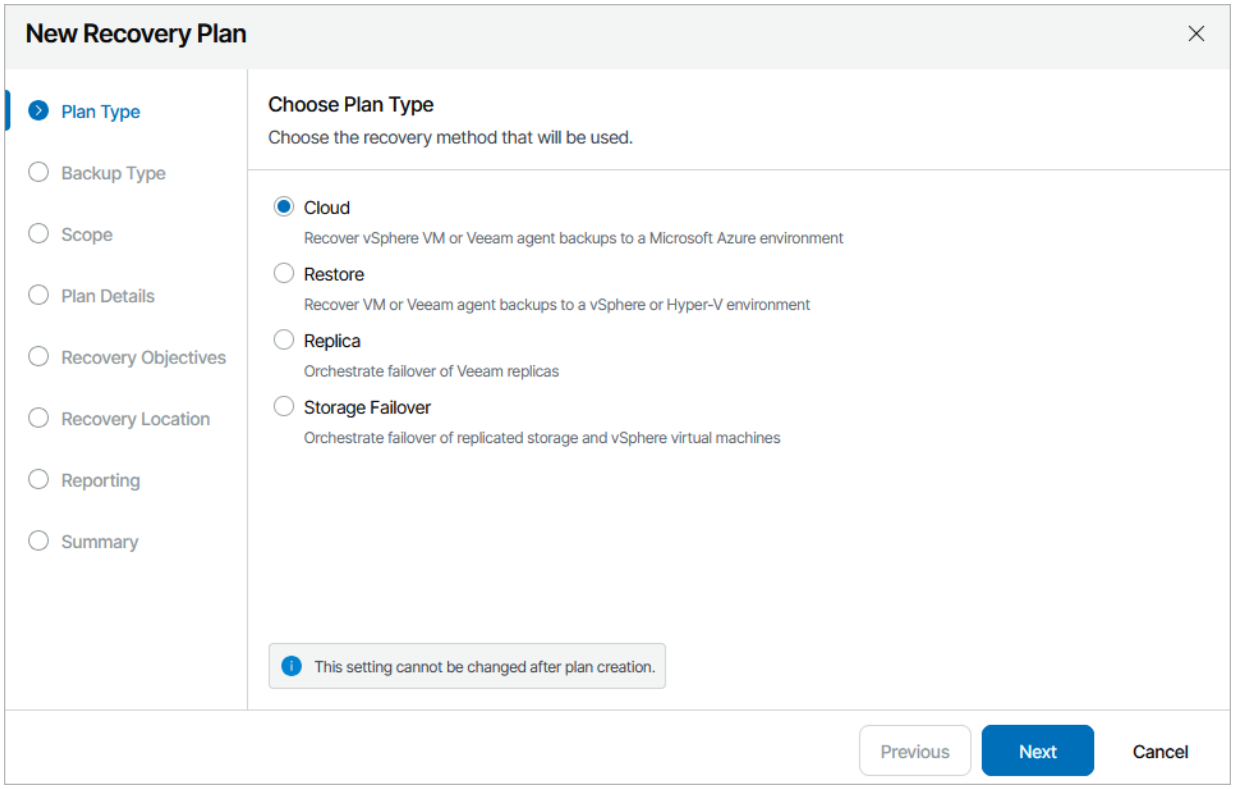

Step 2. Choose Backup Type

At the **Backup Type** step, specify whether you want to recover machines from VM backups or Veeam agent backups.

> **NOTE**
>
> Orchestrator only supports agent backups created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux.

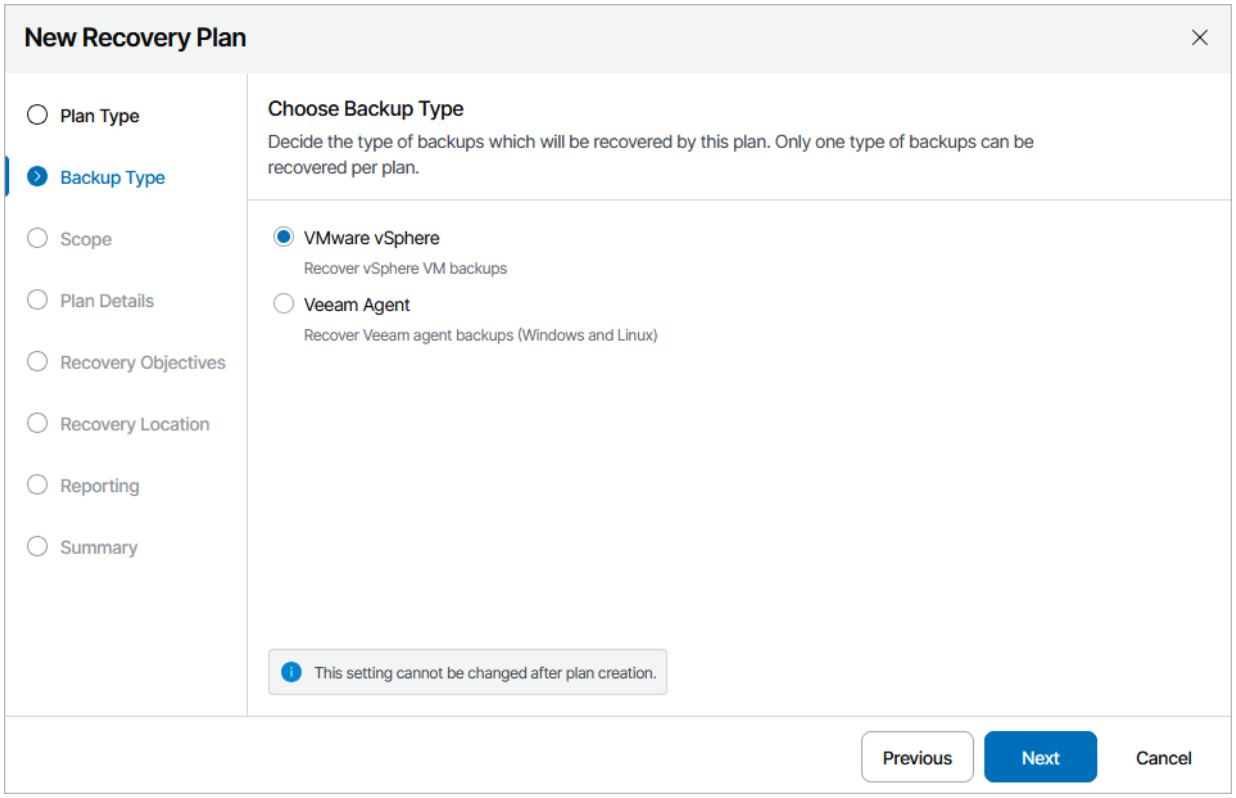

Step 3. Choose Plan Scope

At the **Scope** step of the wizard, select a scope for which you want to create the plan.

For a scope to be displayed in the list, it must be created and customized as described in section Managing Scopes.

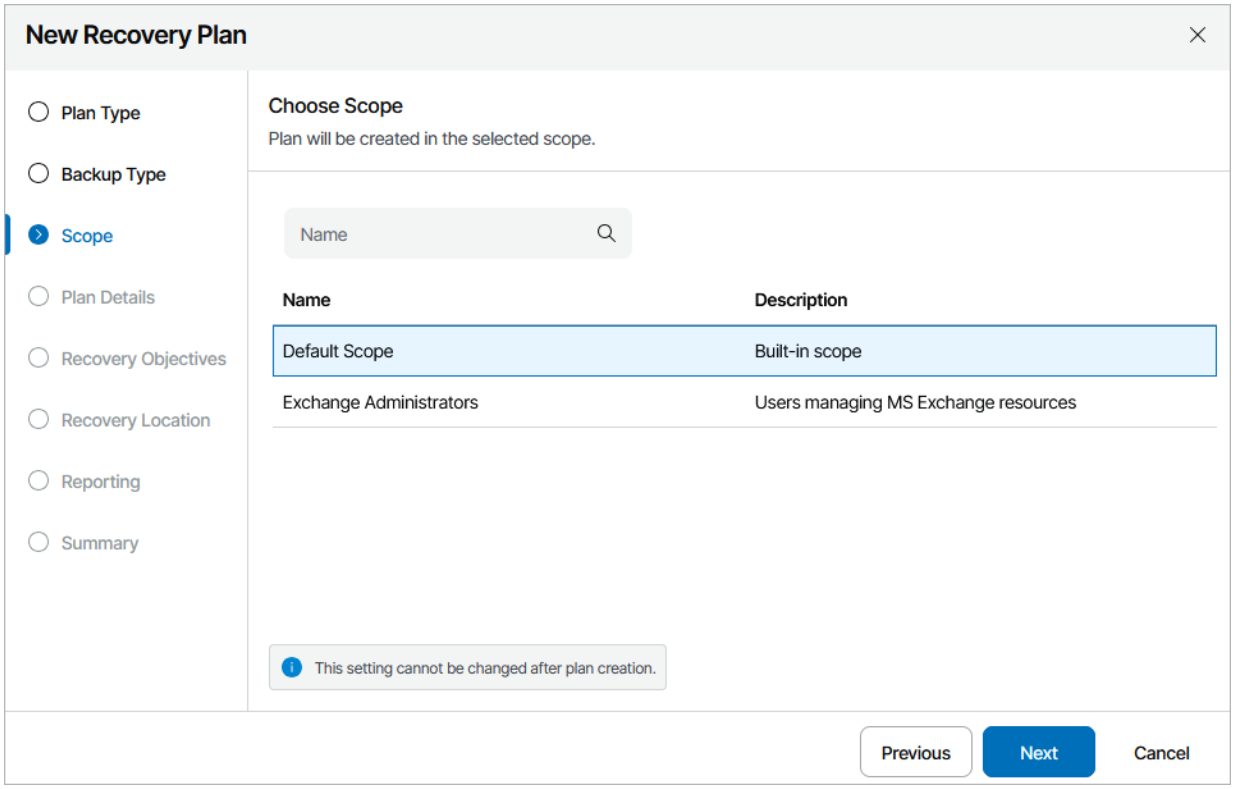

Step 4. Specify Plan Name and Description

At the **Plan Details** step of the wizard, use the **Plan name** and **Description** fields to enter a name for the new plan and to provide a description for future reference. The maximum length of the plan name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

You can also provide a contact name, email and telephone number of a person responsible for the plan.

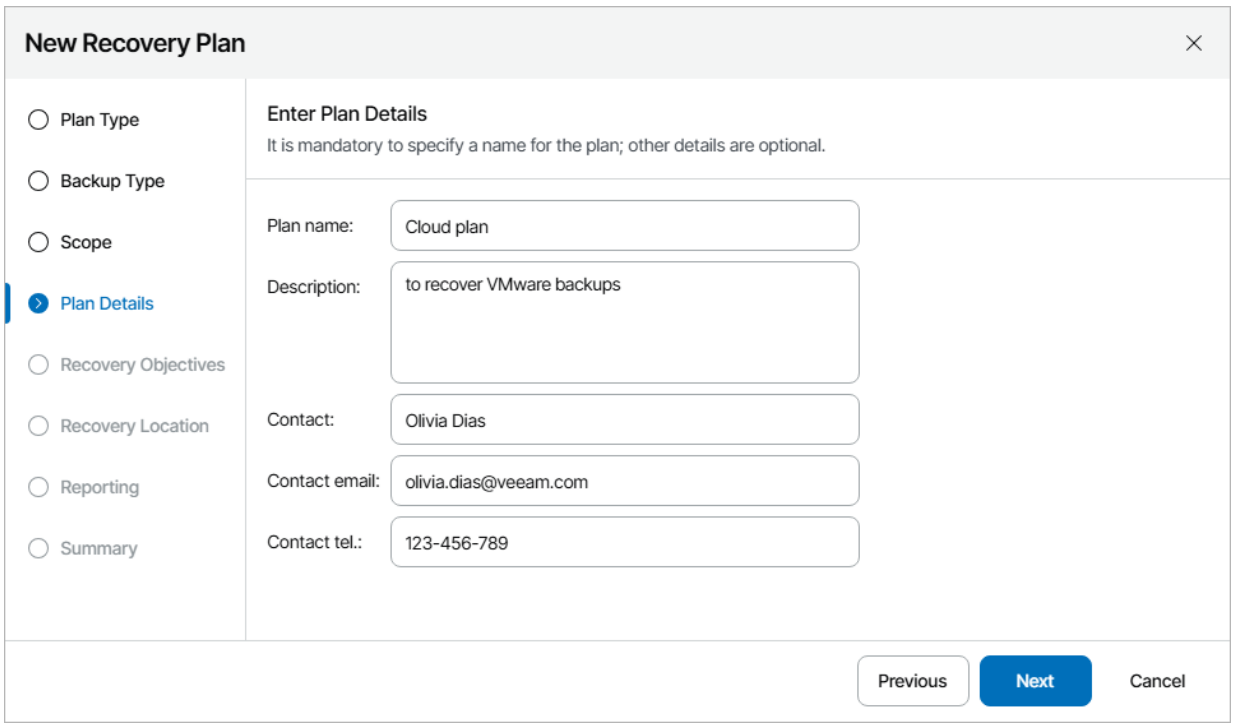

Step 5. Specify Target RTO and RPO

At the **Recovery Objectives** step of the wizard, define your Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for the plan:

- The **RPO** defines the maximum acceptable period of data loss.
- The **RTO** represents the amount of time it should take to recover from an incident.

> **NOTE**
>
> If you choose to perform malware scan while running the plan, Orchestrator will scan one disk per mount server at a time. This process may take a while, affecting the plan RTO.

RTO and RPO performance will be recorded in the Plan Readiness Check, Plan Execution and Plan Audit reports, and you will be able to track the achieved RTO and RPO objectives for each plan on the Home Page Dashboard.

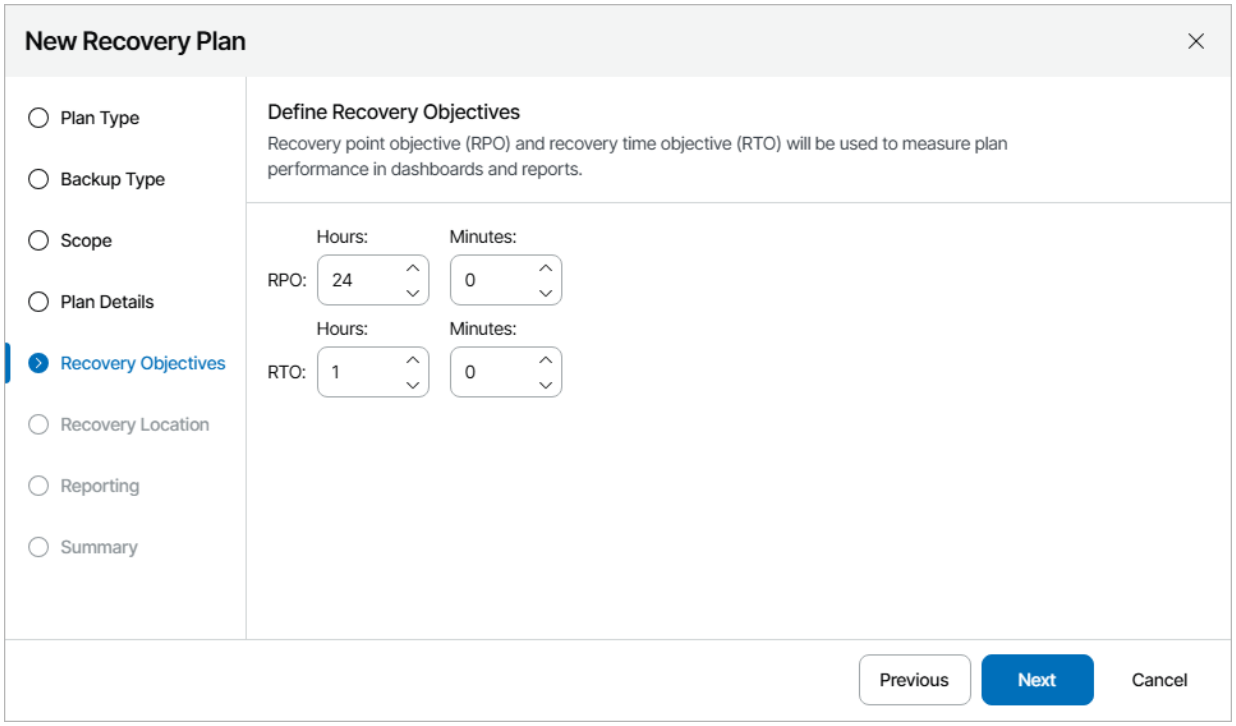

Step 6. Select Recovery Location

At the **Recovery Location** step of the wizard, select a location to which inventory groups included in the plan will be restored.

For a recovery location to be displayed in the list of available locations, it must be created and added to the list of available inventory items available for the scope, as described in section Managing Recovery Locations.

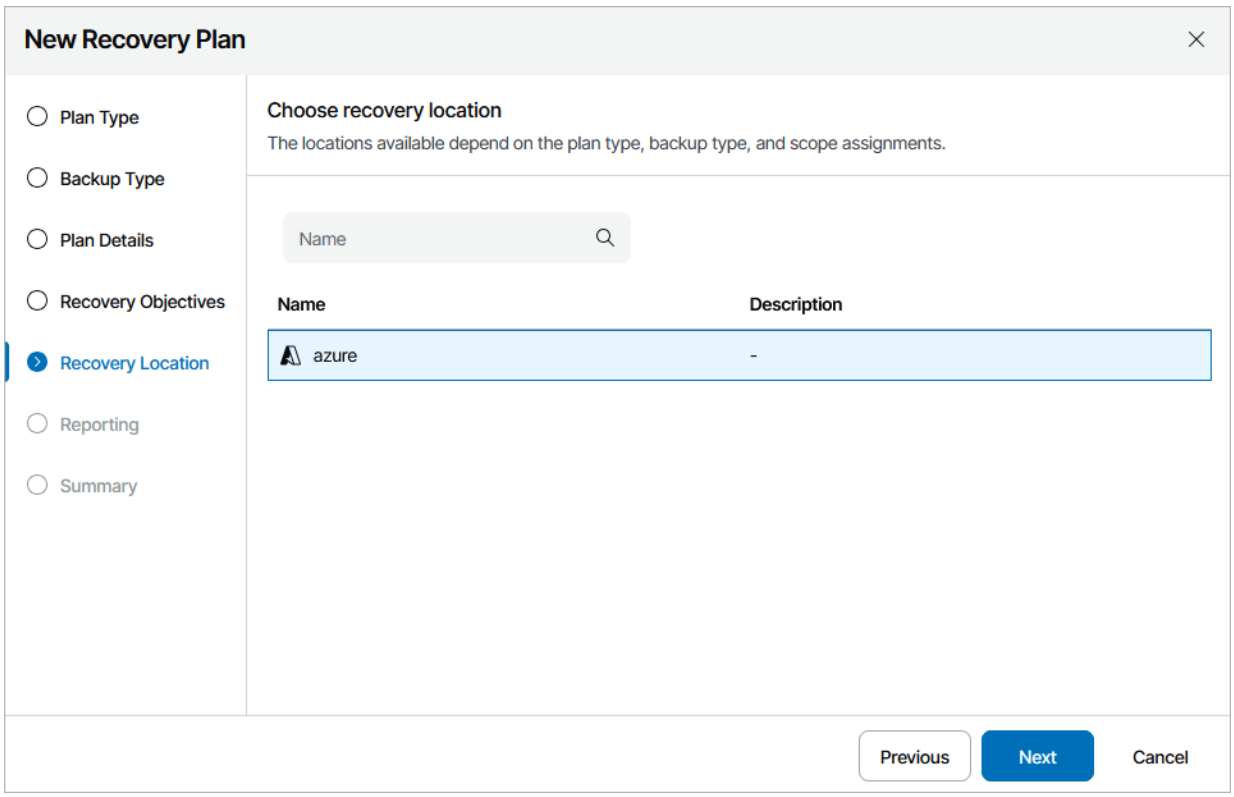

Step 7. Select Report Template

At the **Reporting** step of the wizard, select a document template that will be used as the cover page for all Orchestrator reports.

For a custom document template to be displayed in the list, it must be created and customized as described in section Managing Templates.

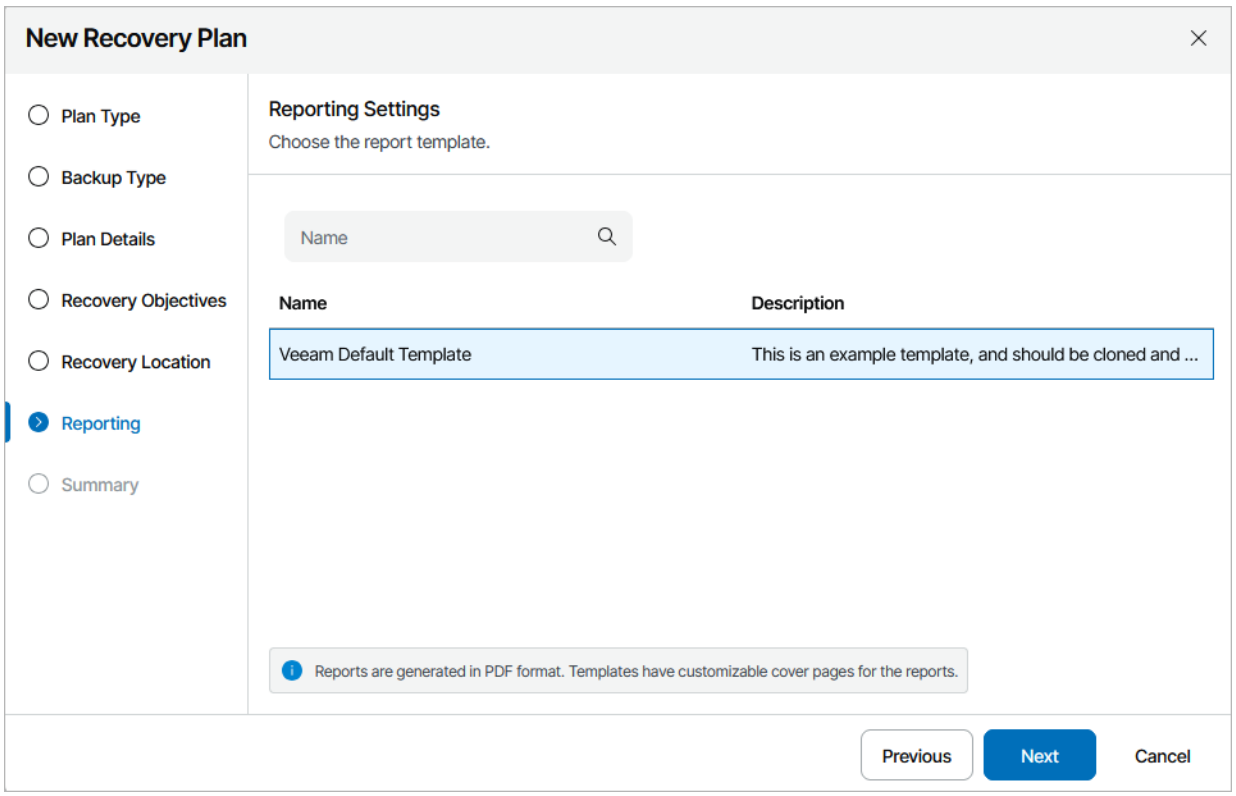

Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.

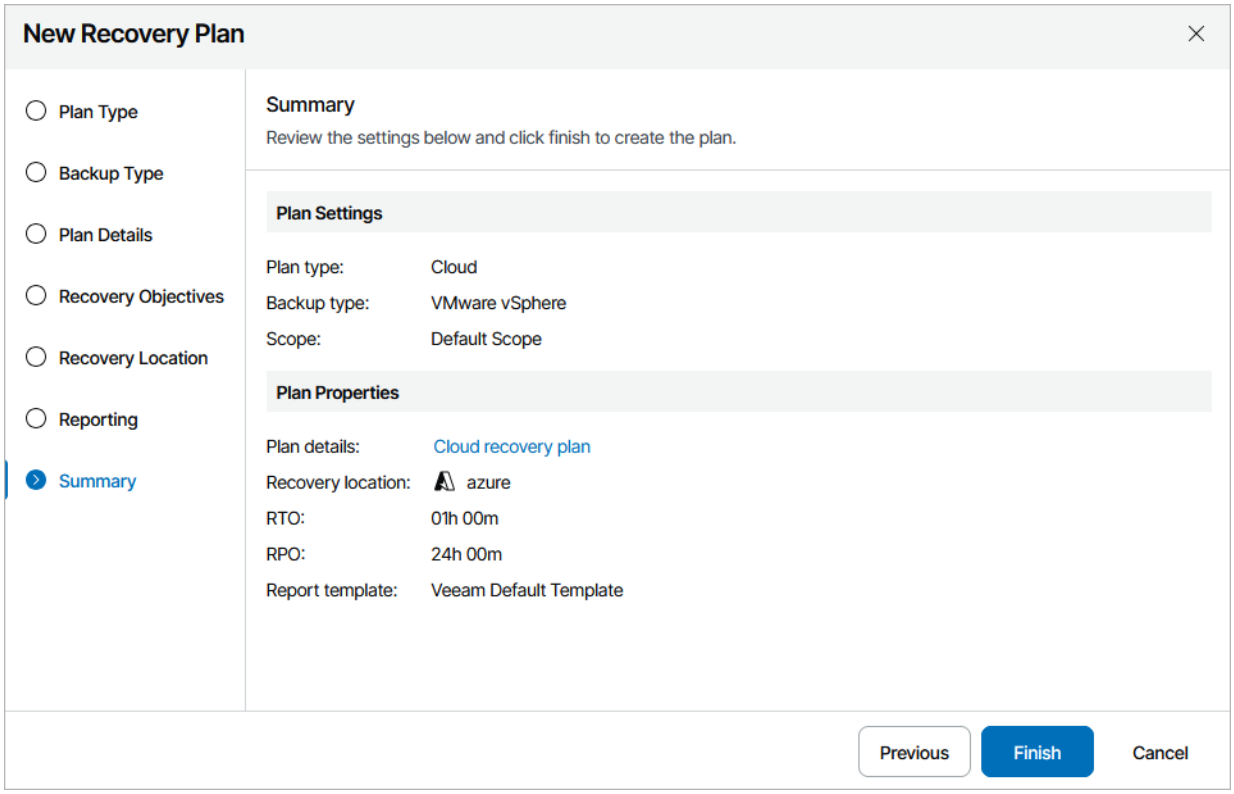

Editing Cloud Plans

If you want to specify granular settings not provided in the New Recovery Plan wizard, the Orchestrator UI allows you to customize cloud plans and configure the settings for groups, recovered VMs, plan steps and step parameters.

The procedures to edit replica, CDP replica, restore, storage and cloud plans are almost identical. For more information, see Editing Recovery Plans.

Scanning Cloud Plans

You can start on-demand plan scanning for malware and configure scan scheduling for any cloud plan. There is almost no difference between the procedures performed for replica, CDP replica, restore and cloud plans. For more information, see Scanning Recovery Plans.

Running and Scheduling Cloud Plans

To run a cloud plan, it must be *ENABLED*. To enable a plan:

1. Navigate to **Recovery Plans**.
2. Select the plan.
3. From the **Manage** menu, select **Properties**.

   OR-

   Right-click the plan name and select **Manage** > **Properties**.
4. Set the **Availability** toggle to *Enabled*.
5. Click **Save**.

If you do not enable a plan before you run it, the Run Plan wizard will force you to do that as soon as you try running the plan.

> **NOTES**
>
> 1. An Orchestrator Administrator can force-enable a plan in the **Run Plan** wizard. However, a Plan Operator will not be able to run a disabled replica plan.
>
>   For more information on roles that can be assigned to users and user groups working with the Orchestrator UI, see Managing User Accounts.
> 2. For security purposes, all 'real-world' actions associated with restore plans require password confirmation.

Scheduling Cloud Restore

You can schedule a time for a cloud plan to execute. Only the restore process can be scheduled — all other operations must be performed manually in the Orchestrator UI.

To schedule a cloud plan:

1. Navigate to **Recovery Plans**.
2. Select the plan. From the **Manage** menu, select **Schedule**.

   -OR-

   Right-click the plan name and select **Manage** > **Schedule**.
3. In the **Scheduled Tasks** window, do the following:

   a. Set the **Schedule plan execution** toggle to *Enabled*.

   b. Click the **Configure schedule** link and choose whether you want to run the plan on schedule or after any other plan:

   - If you want to run the plan at a specific time, click the **Schedule** icon in the **Run on** field, set the desired date and time, and click **Apply**.
   - If you want to run the plan after another plan, select the **Schedule after plan** check box and click **Choose plan**. Then, in the **Select Plan** window, select the necessary plan and click **Apply**.

     For a plan to be displayed in the list of available plans, it must be *ENABLED* as described in section Running and Scheduling Cloud Plans.

c. Set the **Malware actions** toggle to *Enabled* if you want to check restore points created for machines included in the plan for malware flags. You can also decide whether you want to scan these restore points with antivirus software, YARA rules or both.

   By default, Orchestrator checks the most recent restore point on each machine. If no clean restore point is found, Orchestrator performs the following actions depending on whether you have specified a quarantine network when configuring the cloud recovery location:

   - In case you have specified a quarantine network, Orchestrator connects the machine to the network.
   - In case you have not specified a quarantine network, Orchestrator halts the plan and cancels the restore operation.

   For more information on how Orchestrator performs malware scan, see the Veeam Recovery Orchestrator User Guide, section Overview.

d. Review configuration information and click **Save**.

> **TIP**
>
> You can also scan a recovery plan for possible malware without scheduling the plan execution. To do that, follow the instructions provided in section Scanning Recovery Plans.

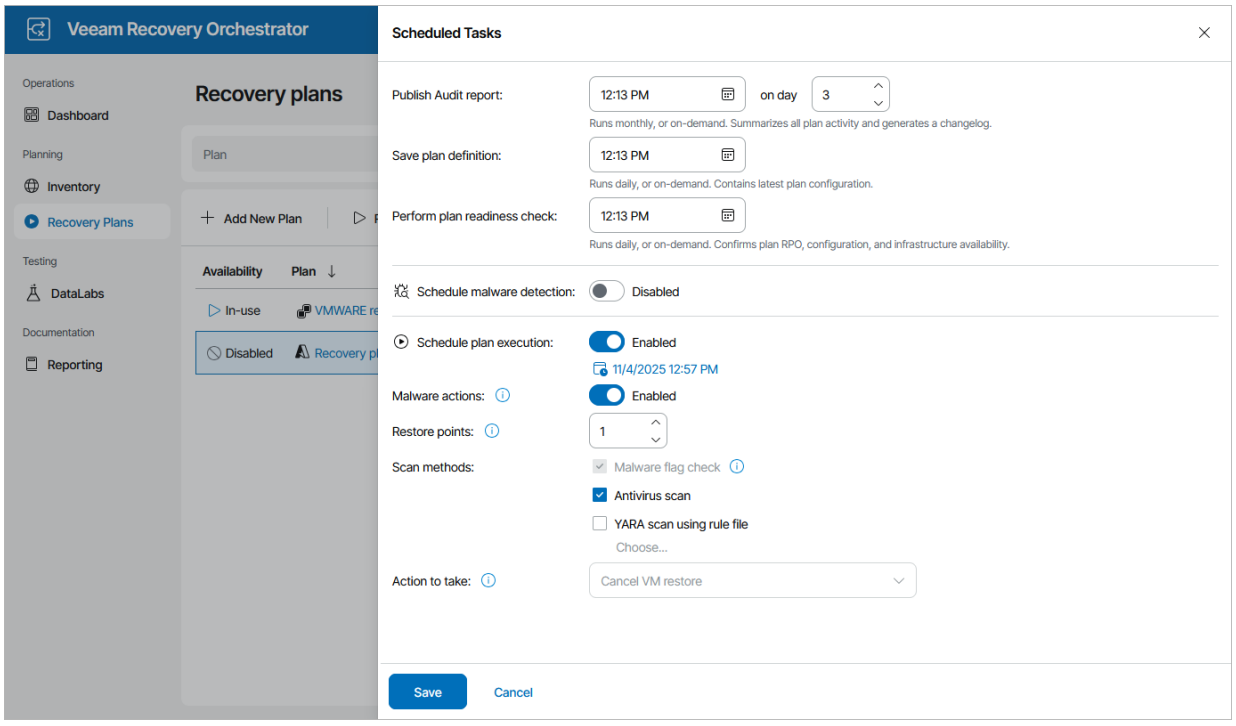

> **TIP**
>
> You can disable a configured schedule if you no longer need it. To do that, set the **Schedule plan execution** toggle to *Disabled* in the **Scheduled Tasks** window.

Running Cloud Restore

The **Run** action causes machines in a plan to recover from their backup files. For more information on the data recovery process, see the Veeam Backup & Replication User Guide, section Data Recovery.

To run a cloud plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Run**.
3. In the **Run Plan** window, do the following:
   a. For security purposes, retype your password and click **Next**.

      You must also select the **Force-enable the plan** check box if you have not enabled the plan yet.

   b. In the **Recovery location** section, select a location to which inventory groups included in the plan will be restored. For a recovery location to be displayed in the list of available locations, it must be created and added to the list of inventory items available for the scope, as described in section Managing Recovery Locations.

      If the selected recovery location includes multiple proxies, Orchestrator will use the round-robin algorithm to restore machines added to the plan. For more information, see the Veeam Recovery Orchestrator User Guide, section How Orchestrator Places VMs During Cloud Restore.

   c. In the **Restore point** section, choose a restore point that will be used to recover machines.

      Keep in mind that recovering data from the archive tier is not supported. If you select the **Most recent** option, make sure to choose a restore point that is stored in either the capacity or the performance tier. For more information on Veeam Backup & Replication tiering options, see the Veeam Backup & Replication User Guide, section Scale-Out Backup Repository.

   d. In the **Malware detection** section, choose whether you want to check restore points created for machines included in the plan for malware flags. You can also decide whether you want to scan these restore points with antivirus software, YARA rules or both.

      By default, Orchestrator checks the most recent restore point on each machine. If no clean restore point is found, Orchestrator performs the following actions depending on whether you have specified a quarantine network when configuring the cloud recovery location:

      - In case you have specified a quarantine network, Orchestrator connects the machine to the network.
      - In case you have not specified a quarantine network, Orchestrator halts the plan and cancels the restore operation.

      For more information on how Orchestrator performs malware scan, see the Veeam Recovery Orchestrator User Guide, section Overview.

   e. At the **Summary** step, review configuration information and click **Finish**.

> **TIP**
>
> You can also scan a restore plan for possible malware without running the plan. To do that, follow the instructions provided in section Scanning Recovery Plans.

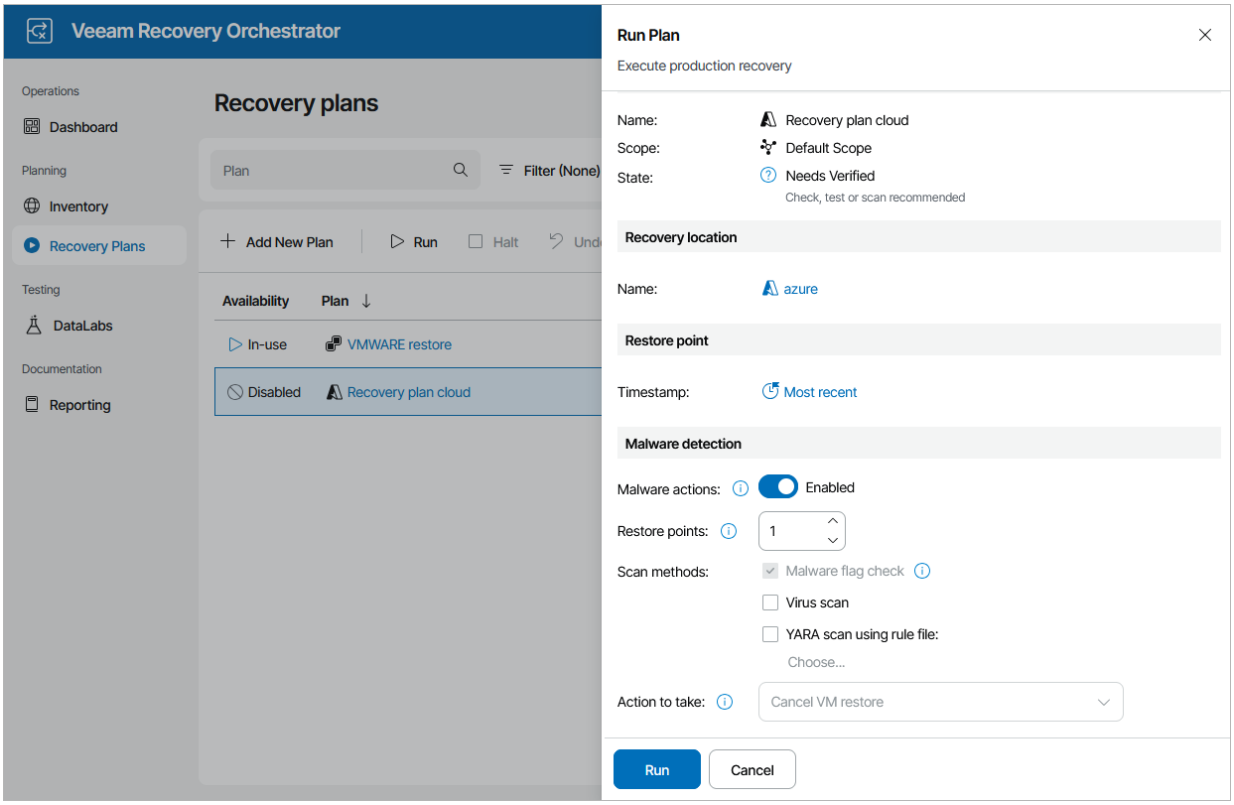

The plan goal is to reach the *RESTORED* state. If any critical error is encountered, the plan will stop with the *HALTED* state. To learn how to work with *HALTED* cloud plans, see Managing Halted Plans.

Halting Cloud Restore

The **Halt** action interrupts plan execution. Any steps currently executing will be completed, then the plan will enter the *HALTED* state. To learn how to work with *HALTED* cloud plans, see Managing Halted Plans.

To stop a running cloud plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Halt**.
3. In the **Halt Plan Execution** window, do the following:
    a. For security purposes, retype your password and click **Next**.

b. Review configuration information and click **Halt**.

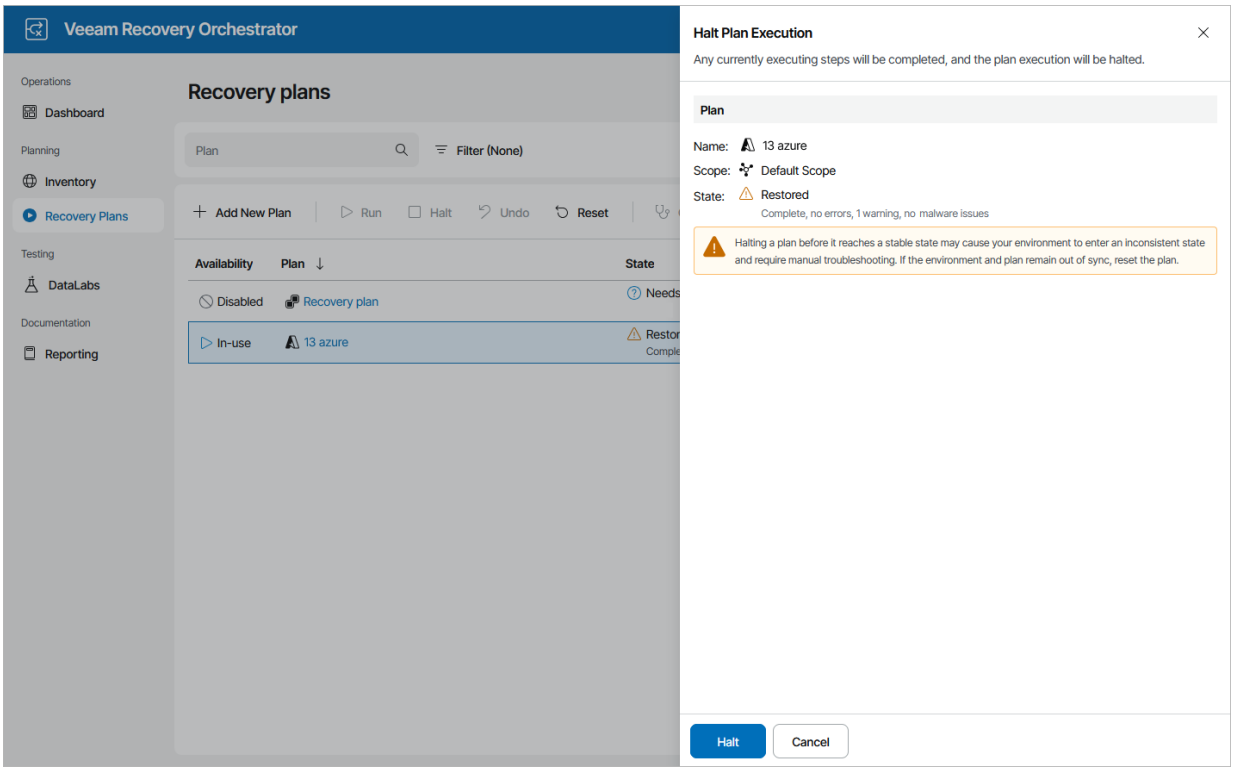

Resetting Cloud Plans

If a cloud plan becomes inconsistent with the virtual environment, you can reset the plan. This will return the plan to the *DISABLED* state, without making any changes to the external virtual infrastructure.

To reset a cloud plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Reset**.
3. In the **Reset Plan** window, do the following:
   a. For security purposes, retype your password and click **Next**.
   b. Select the **Run a quick check after reset is complete** check box to run a readiness check after the reset.

c. Review configuration information and click **Reset**.

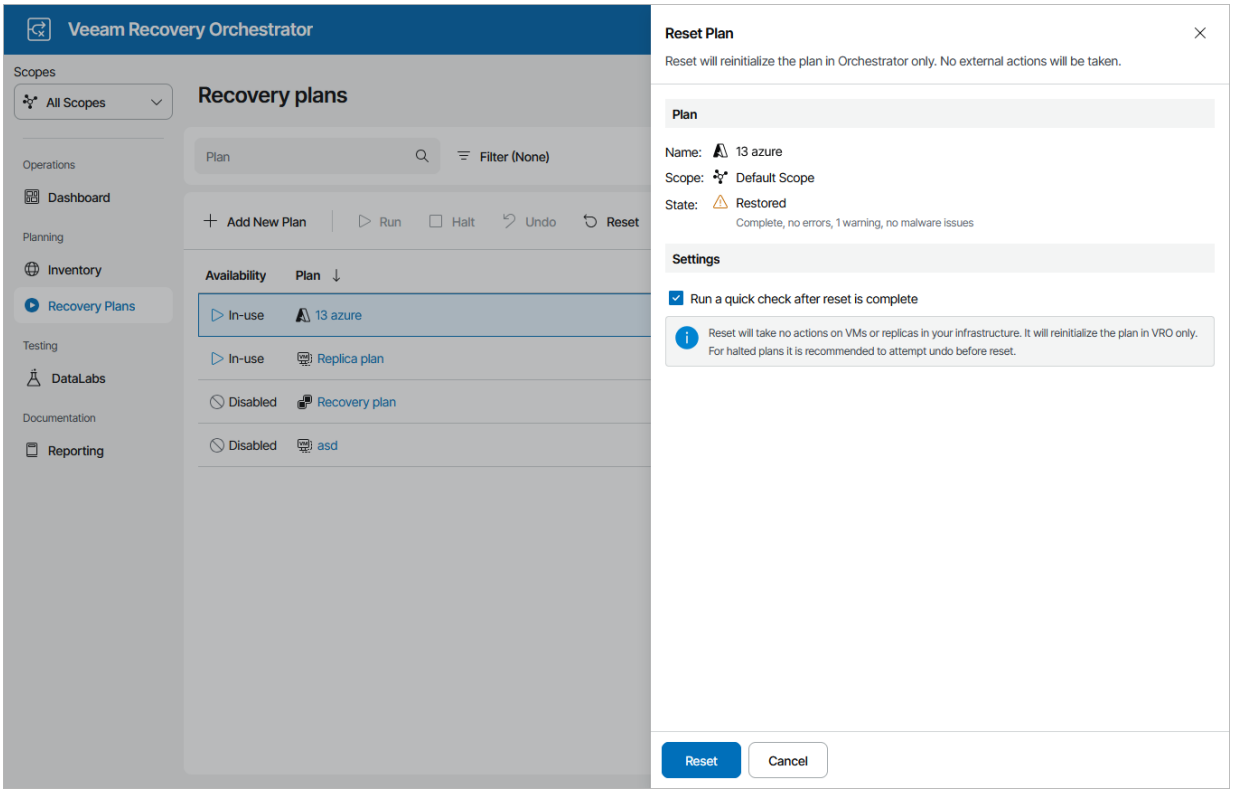

Managing Halted Cloud Plans

If a critical step fails for a machine from a critical inventory group, the plan may enter the *HALTED* state. To troubleshoot reasons why a plan failed, use the **Plan Execution Report** generated as soon as the currently performed action completes. For more information on how to track plan performance history, see Viewing Plan Execution History.

After you eliminate the problem that caused the plan to become *HALTED*, you have the following options to resume the plan:

- Repeat the last failed step.
- Proceed to the next step.

Running Halted Cloud Plans

To run a *HALTED* cloud plan:

1. Navigate to **Recovery Plans**.
2. Select the halted plan and click **Run**.
3. In the **Continue Execution** window, do the following:
   a. For security purposes, at the **Credentials** step, retype your password.
   b. In the **Retry settings** step, select an option to resume plan execution.

      Choose whether you want to proceed with plan execution from the next plan step or to retry the failed step.

c. Review configuration information and click **Finish**. The restore process will be started.

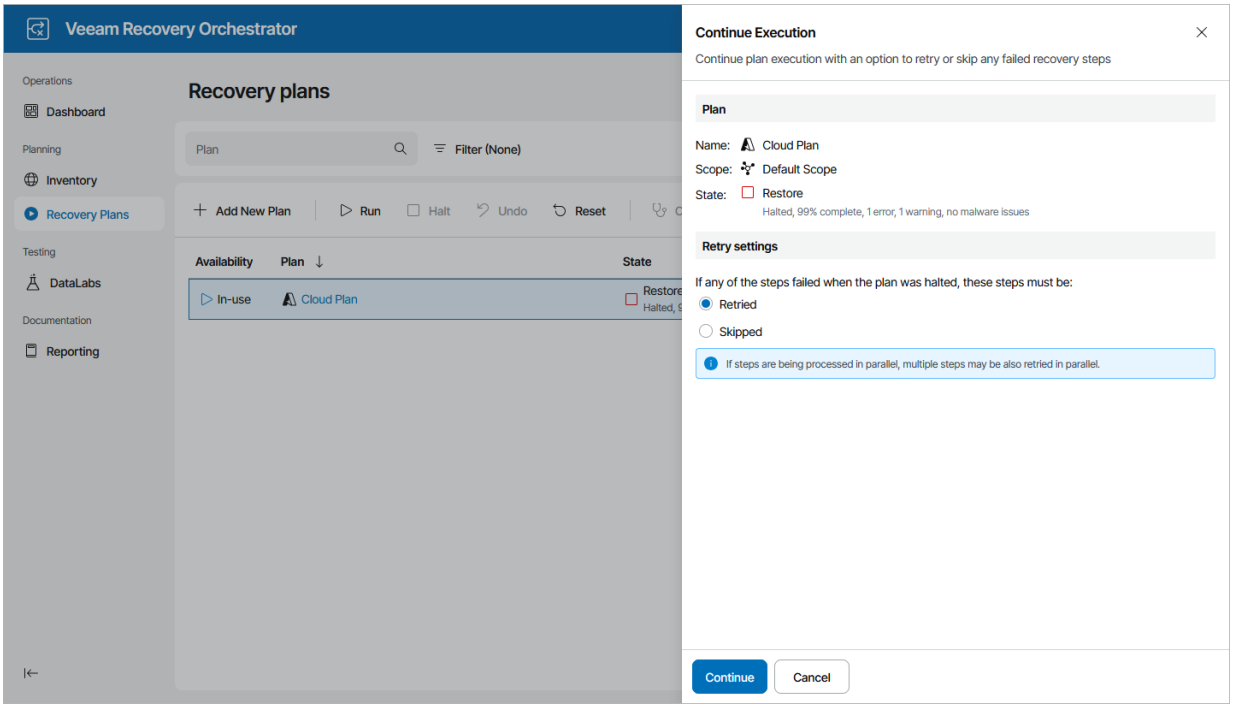

Resetting Halted Cloud Plans

To reset a *HALTED* cloud plan, follow the instructions provided in section Resetting Cloud Plans.

> **NOTE**
>
> When you reset a cloud plan, Orchestrator returns it to the *DISABLED* state without making any changes to the external virtual infrastructure. You may need to deal with any infrastructure reconfiguration manually.

Editing Recovery Plans

In addition to default recovery plan settings that you specify when creating a plan, you can also specify granular settings to customize the plan. The Orchestrator UI allows you to adjust the following:

- Plan properties
- Group settings
- Plan step settings
- Step parameter settings

> **NOTE**
>
> You cannot edit a recovery plan if the plan is in the *IN-USE* mode. If a recovery plan is in the *ENABLED* mode, you can edit only plan properties. To allow editing plan properties and other plan settings, you must disable the plan.
>
> For the list of modes that different types of recovery plans can acquire, see Replica Plans, CDP Replica Plans, Restore Plans, Storage Plans and Cloud Plans.

Configuring Plan Properties

For each recovery plan, you can configure settings specified while creating the plan:

1. Navigate to **Recovery Plans**.
2. Select the plan.
3. From the **Manage** menu, select **Properties**.

   OR-

   Right-click the plan name and select **Manage** > **Properties**.
4. In the **Plan Properties** window, do the following:

   a. To provide a new name, description, contact name, email or telephone number of a person responsible for the plan, follow the instructions provided in section Creating Replica Plans (step 4), Creating CDP Replica Plans (step 4), Creating Restore Plans (step 4), Creating Storage Plans (step 4) or Creating Cloud Plans (step 4).

   b. [Applies only to cloud and restore plans] To select a new location to which inventory groups included in the plan will be restored, follow the instructions provided in section Creating Restore Plans (step 6) or Creating Cloud Plans (step 6).

   c. To modify the configured Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for the plan, follow the instructions provided in section Creating Replica Plans (step 5), Creating CDP Replica Plans (step 5), Creating Restore Plans (step 5), Creating Storage Plans (step 5) or Creating Cloud Plans (step 5).

   d. To select a new document template that will be used to create documents for the plan, follow the instructions provided in section Creating Replica Plans (step 6), Creating CDP Replica Plans (step 6), Creating Restore Plans (step 7), Creating Storage Plans (step 6) or Creating Cloud Plans (step 7).

   e. Review configuration information and click **Save**.

If you want to run the plan immediately, set the **Availability** toggle to *Enabled* and follow the instructions provided in section Running and Scheduling Replica Plans, Running and Scheduling CDP Replica Plans, Running and Scheduling Restore Plans, Running and Scheduling Storage Plans or Running and Scheduling Cloud Plans.

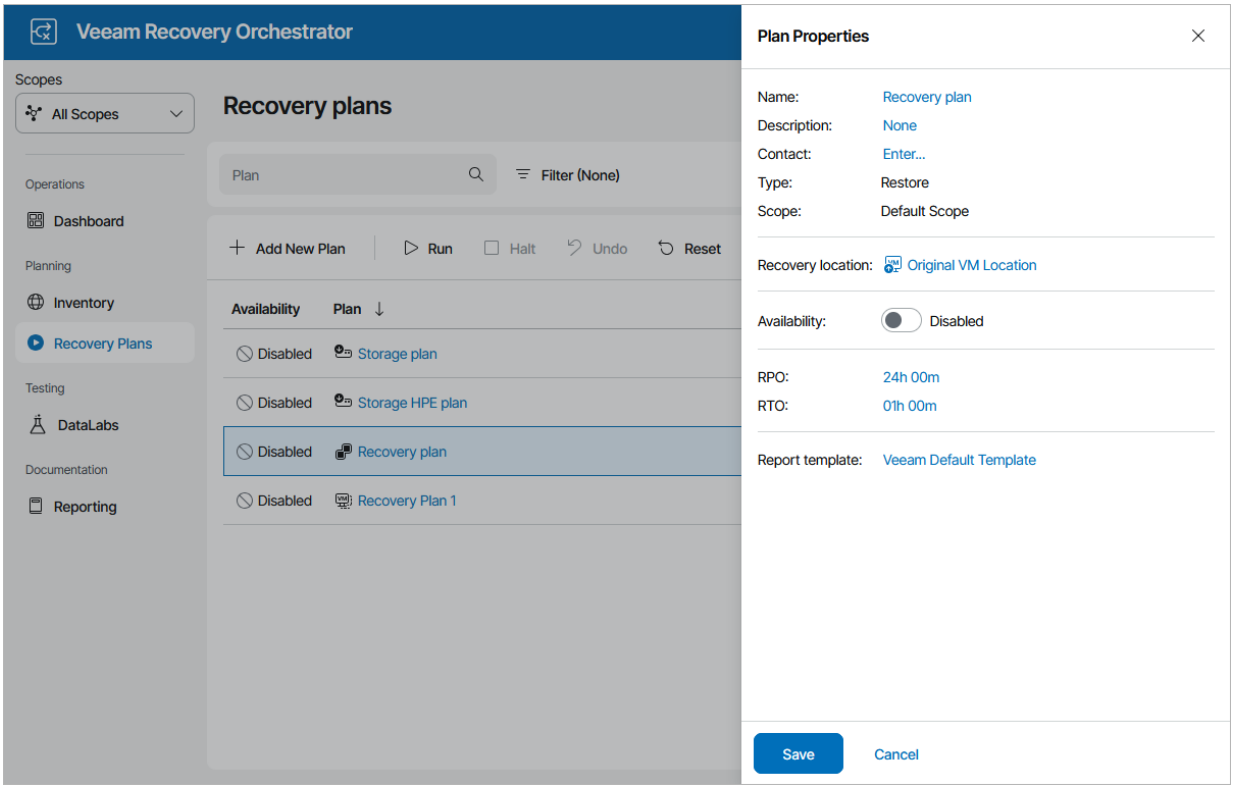
Veeam Recovery Orchestrator
Scopes
All Scopes
Operations
Dashboard
Planning
Inventory
Recovery Plans
Testing
DataLabs
Documentation
Reporting
Recovery plans
Plan
Filter (None)
Add New Plan
Run
Halt
Undo
Reset
Availability
Plan
Disabled
Storage plan
Disabled
Storage HPE plan
Disabled
Recovery plan
Disabled
Recovery Plan 1
Plan Properties
Name:
Recovery plan
Description:
None
Contact:
Enter...
Type:
Restore
Scope:
Default Scope
Recovery location:
Original VM Location
Availability:
Disabled
RPO:
24h 00m
RTO:
01h 00m
Report template:
Veeam Default Template
Save
Cancel

Configuring Groups

Options on the **Edit Plan** page allow you to:

- Add inventory groups to a plan
- Change the processing order for groups in a plan
- Turn on and turn off post-recovery protection of machines in a group
- Customize recovery settings for machines in a group

Adding Inventory Groups

After you create a recovery plan, you must add to this plan inventory groups that contain VMs that you plan to restore:

1. Navigate to **Recovery Plans**.
2. Select the plan to which you want to add inventory groups and click **Manage** > **Edit**.
3. On the **Edit Plan** page, click **Add**.
4. In the **Add Inventory Group** window, choose groups that you want to add to the plan and click **Next**. For an inventory group to be displayed in the **Group** list, it must be added to the list of inventory items available for the scope, as described in section Managing Inventory Items.

   Note that one restore plan can contain inventory groups of one type only (either vSphere, Veeam Agent or Hyper-V). To recover workloads added to inventory groups of different types, create separate restore plans.
5. In the **Group Options** window, configure the group settings as described in section Configuring Group Settings.
6. To save changes made to the plan settings, click **Apply**.

> **IMPORTANT**
>
> For Orchestrator to be able to recover a machine to a VMware vSphere environment, it is recommended that the machine has VMware Tools installed:
>
> - For VMs recovered from Veeam agent backups, Orchestrator automatically verifies whether VMware Tools are installed on all machines included in a plan when running a readiness check or a DataLab test for the plan. However, this verification is supported for Windows-based machines only. For Linux-based machines, you must perform the verification manually.
> - For VMs recovered from vSphere backups, the verification is performed automatically on the vCenter Server side — for both Windows-based and Linux-based VMs. To know how to install and upgrade VMware Tools in vSphere, see this VMware KB article.

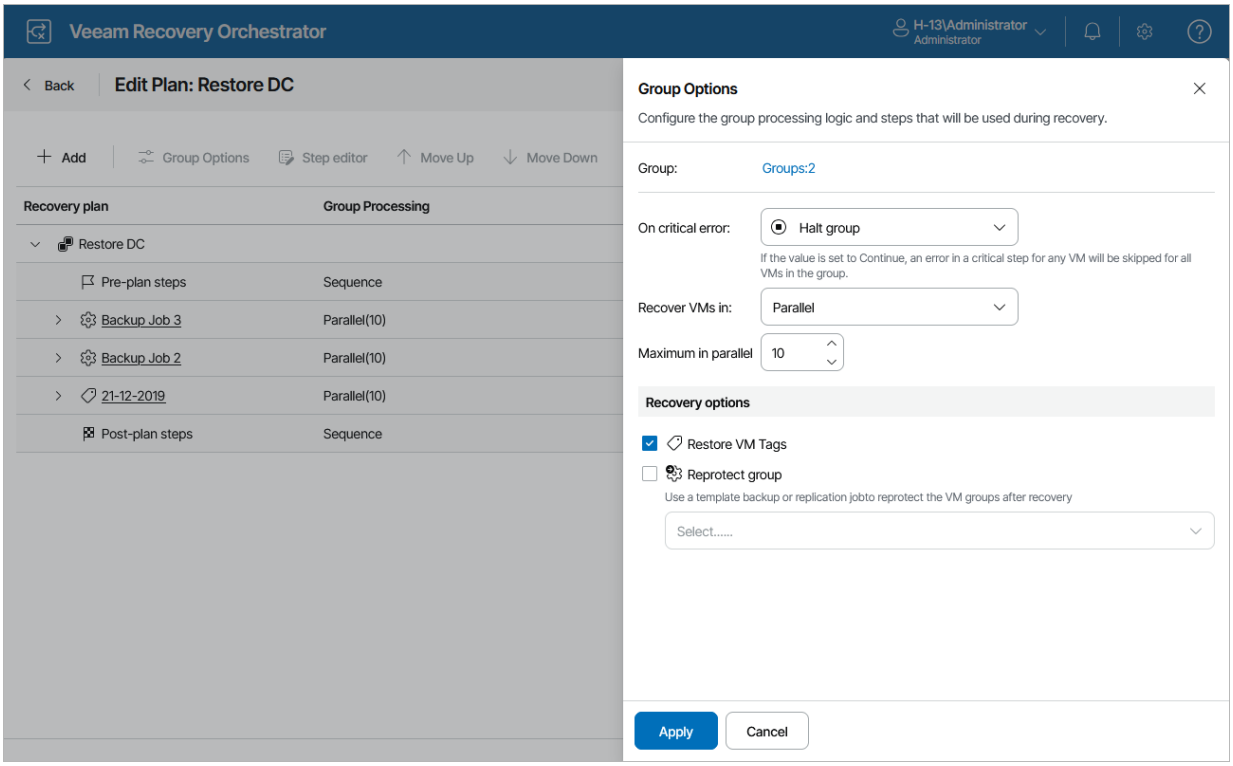

Configuring Group Settings

You can configure the following group settings:

1. Navigate to **Recovery Plans**.
2. Select the plan that contains an inventory group you want to edit and click **Manage** > **Edit**.
3. On the **Edit Plan** page, in the **Recovery plan** column, expand the plan to see all its inventory groups. Then, select the necessary inventory group and click **Group Options**.
4. In the **Group Options** window, do the following:
   a. Use the options in the **On critical error** drop-down list to choose whether you want to halt plan execution if machine recovery fails.
   b. Use the options in the **Recover VMs in** drop-down list to choose whether you want to recover machines in sequence or in parallel. If you select to process machines simultaneously, use the **Maximum in parallel** field to specify the maximum number of VMs processed at the same time.
   c. Select the **Restore VM Tags** check box if you want the recovered VMs to have the same tags as the source machines.
   d. Use the **Reprotect group** check box to choose whether you want to protect VMs in the plan post-recovery with a backup or replication job. Keep in mind that you cannot reprotect VMs recovered to a Microsoft Hyper-V environment.

      If you select the **Reprotect group** check box, you must specify a backup or replication job to be used as a template for a new job that will reprotect recovered VMs. To do that, select the required job from the drop-down list.

      For a template backup or replication job to be displayed in the list of available jobs, it must be created and added to the list of inventory items for the scope, as described in section Editing Template Jobs.

> **IMPORTANT**
>
> The new job will consume Veeam Backup & Replication licenses to protect the machines. That is why you must take into account the number of licenses installed on the Veeam Backup & Replication server, so that the number of managed objects does not exceed the license limit.

e. Click **Apply** to save changes made to the group settings.

5. Repeat the procedure for each inventory group that you want to edit and click **Save Plan**.

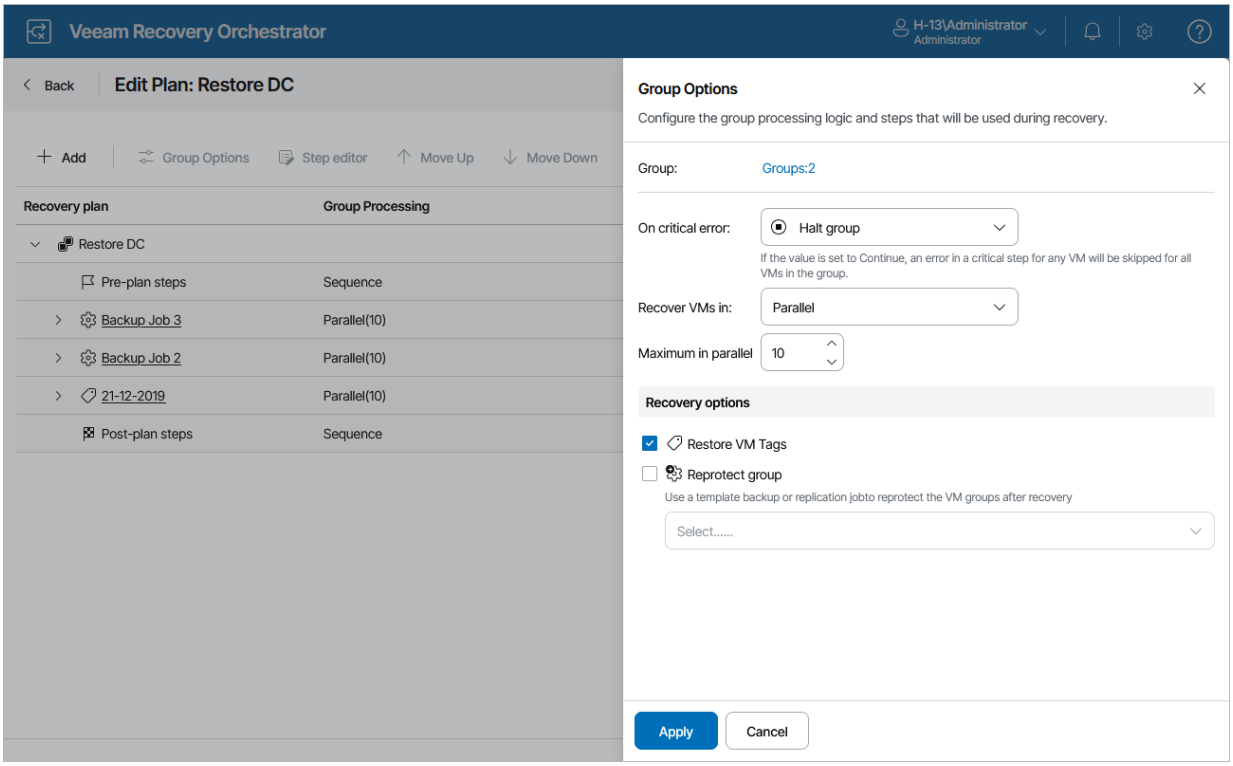

Setting Group Processing Order

Inventory groups in a recovery plan are processed in the order they appear in the **Recovery plan** list. If some machines in a group are dependent upon machines in other groups, make sure that the required group is recovered first.

To change the processing order for inventory groups included in a recovery plan:

1. Navigate to **Recovery Plans**.
2. Select the plan for which you want to change the group processing order and click **Manage** > **Edit**.
3. On the **Edit Plan** page, do the following:
   a. Expand the plan to see all its inventory groups and select an inventory group whose processing order you want to change.
   b. To move the group up or down the list, use the **Move Up** and **Move Down** arrows.
   c. To save changes made to the plan settings, click **Save Plan**.

NOTE

By design, each recovery plan contains 2 default groups — *Pre-plan steps* and *Post-plan steps*. These groups include plan steps that run before and after the recovery process. You cannot change the processing order for the *Pre-plan steps* and *Post-plan steps* groups, but you can add and remove steps for these groups. For more information, see Configuring Steps.

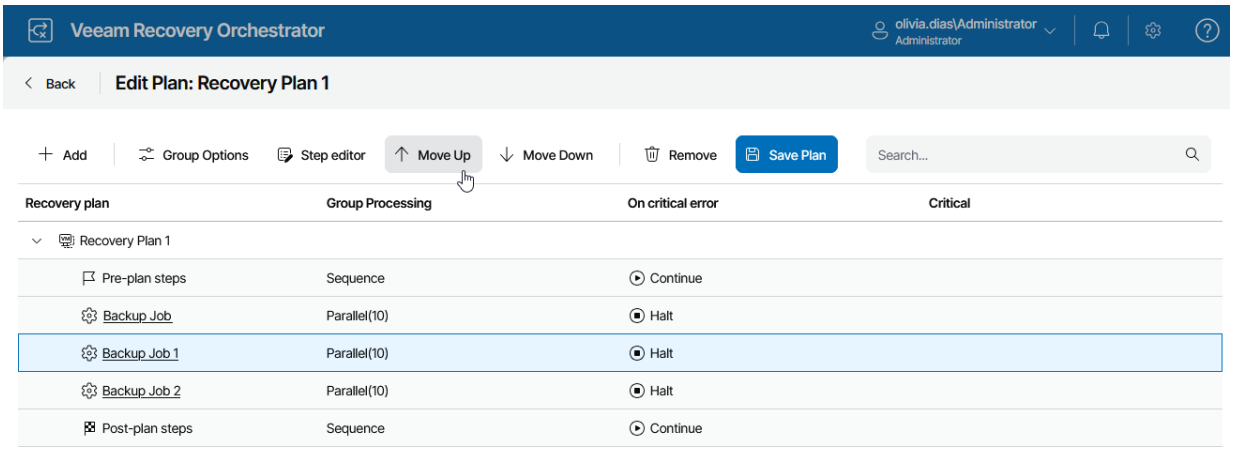

Configuring Machines

The order in which machines in a recovery plan are processed depends on the **Group Options** configured when editing the plan:

- If the **Recover VMs in** > **Sequence** option is selected for an inventory group, machines in the group will be processed in the order they appear in the **Recovery plan** list.
- If the **Recover VMs in** > **Parallel** option is selected for an inventory group, a limited number of machines in the group will be processed at the same time.
- If new machines are added to an inventory group, the entire machine list will be resorted and then processed in the alphabetical order.

If some machines are dependent on other machines, ensure the required machines are started first. To define the recovery order for machines included in an inventory group:

1. Navigate to **Recovery Plans**.
2. Select the plan for which you want to change the group processing order and click **Manage** > **Edit**.
3. On the **Edit Plan** page, do the following:
   a. Expand the plan to see all its inventory groups and select an inventory group in which you want to change the machine processing order.
   b. To move the machine up or down the list, use the **Up** and **Down** arrows.
   c. To save changes made to the plan settings, click **Save Plan**.

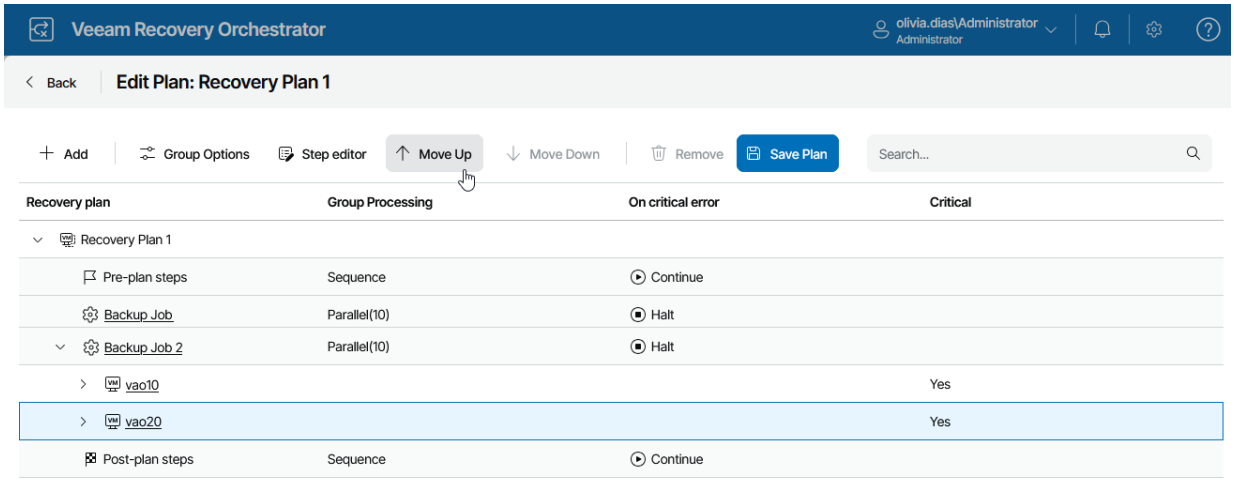

Configuring Steps

For each machine included in a recovery plan, you can add and remove steps performed when processing the machine:

1. Navigate to **Recovery Plans**.
2. Select the plan that contains a machine whose steps you want to edit.
3. From the **Manage** menu, select **Edit**.

   OR-

   Right-click the plan name and select **Manage** > **Edit**.
4. On the **Edit Plan** page, expand the plan, select the necessary machine and click **Step editor**. The the **Step Editor** window will open.
5. In the **Step Editor** window, do the following:
   - To change the step execution order, use the **Move Up** and **Move Down** arrows to move steps up and down the list.
   - To remove a step, select the step and click **Remove**.
   - To add a step, click **Add**. In the **Add Recovery Steps** window, click **Add** and select steps to be performed for each machine during restore.

     For a step to be displayed in the **Step name** list, it must be added to the list of inventory items available for the scope, as described in section Managing Inventory Items. If a step is displayed as not available, this means that the step has already been added to the plan and cannot be added twice.

> **NOTE**
>
> If a VM is included in multiple inventory groups in the same plan, Orchestrator will only run the **Restore VM** step once. However, other steps for this VM will execute when processing it in each group.

5. To save changes made to the plan settings, click **Save**.

> **TIP**
>
> You can simultaneously add steps for multiple machines in each inventory group. To do that, select an inventory group in the **Recovery plan** column and click **Step editor**. In the **Step Editor** window, choose whether you want to add current or default steps, and click **Add**. In the **Add Recovery Steps** window, select the required steps that you want to add and then click **Save**.

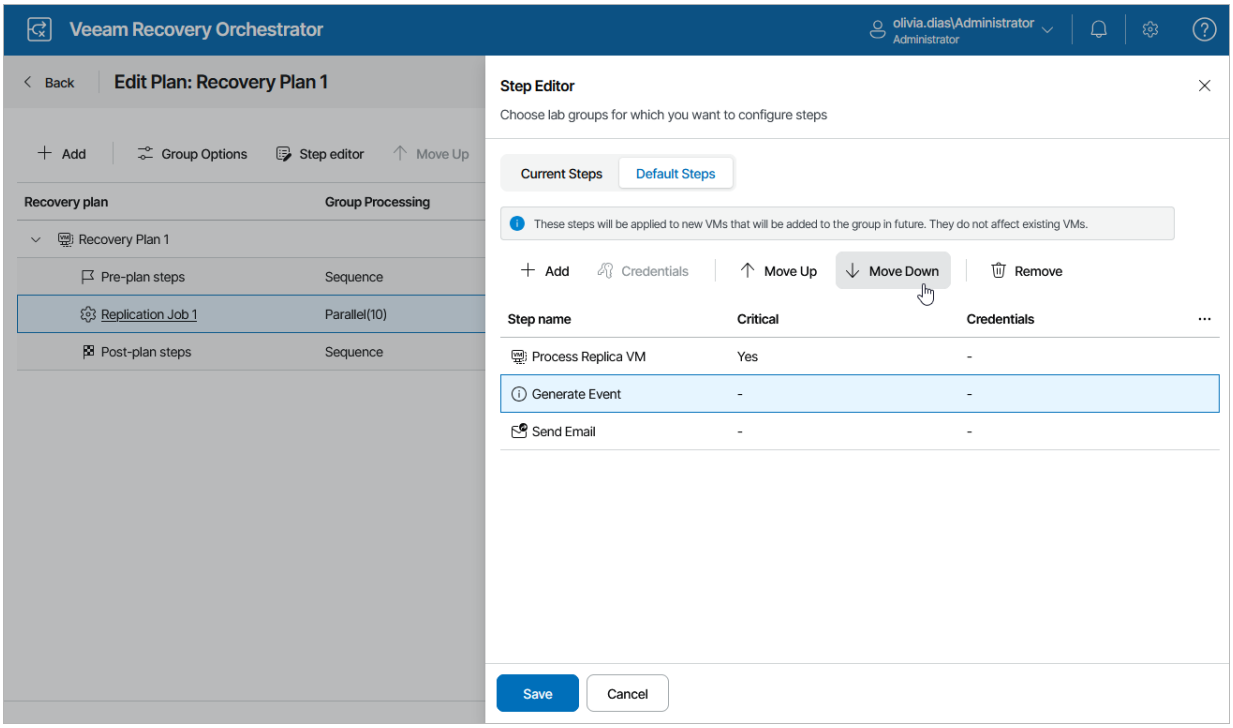
Veeam Recovery Orchestrator
olivia.dias\Administrator
Administrator
Back
Edit Plan: Recovery Plan 1
Add
Group Options
Step editor
Move Up
Recovery plan
Group Processing
Recovery Plan 1
Pre-plan steps
Sequence
Replication Job 1
Parallel(10)
Post-plan steps
Sequence
Step Editor
Choose lab groups for which you want to configure steps
Current Steps
Default Steps
These steps will be applied to new VMs that will be added to the group in future. They do not affect existing VMs.
Add
Credentials
Move Up
Move Down
Remove
Step name
Critical
Credentials
Process Replica VM
Yes
-
Generate Event
-
-
Send Email
-
-
Save
Cancel

Configuring Step Parameters

For each plan step performed during recovery, you can customize parameter settings:

1. Navigate to **Recovery Plans**.
2. Select the plan that contains a VM whose step you want to edit and click **Manage** > **Edit**. The **Edit Plan** page will open.
3. On the **Edit Plan** page, in the **Recovery plan** column, expand the plan to see all its inventory groups. Then, expand the necessary inventory group, select a VM whose step parameters you want to edit, and click **Step editor**.
4. In the **Step Editor** window, select the necessary step and click **Step parameters**.
5. In **Step parameters** window, set the desired step parameter values and click **Save**.

   For detailed description of step parameters that you can configure for recovery plan steps, see Appendix A. Recovery Plan Steps.
6. To save changes made to the plan settings, click **Save Plan**.

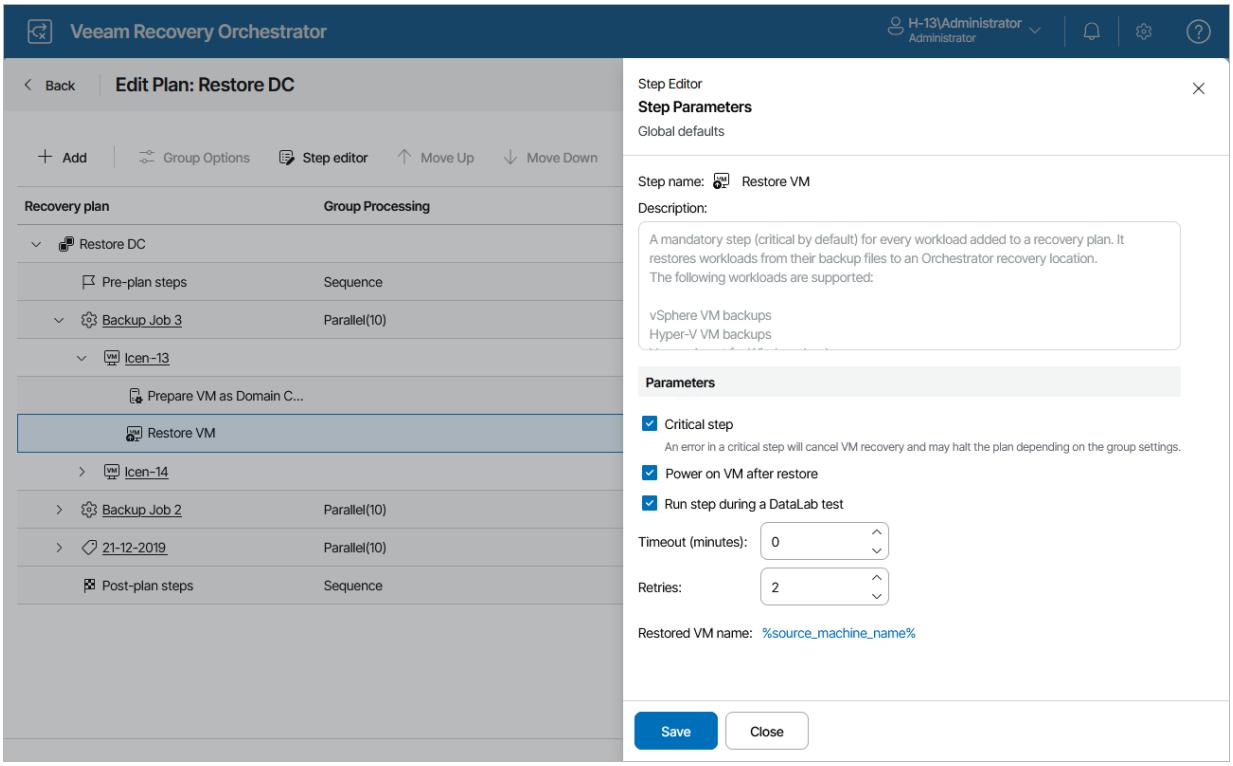

Testing Recovery Plans

Before you run a recovery plan, you can use an isolated Orchestrator DataLab to test the entire plan, including the verification of vSphere and agent backups, replicas and storage snapshots. All changes made to machines during a lab session will be discarded as soon as the testing process is over.

> **NOTE**
>
> DataLab testing is currently not supported for Microsoft Azure and Microsoft Hyper-V recovery locations.

To test a recovery plan in a DataLab, perform a number of steps:

1. Create a virtual lab in Veeam Backup & Replication and configure a connection to this lab.
2. Assign the DataLab to a scope.
3. Associate the DataLab with a recovery location.
4. [Optional] Create a lab group.
5. Start on-demand testing or configure test scheduling settings.

Connecting DataLabs

To validate your disaster recovery plans without impacting the production infrastructure, you can configure automatic scheduled testing for the verification of vSphere and agent backups, VM replicas, applications and storage snapshots. For this purpose, Orchestrator uses virtual labs created in the Veeam Backup & Replication console. These virtual labs provide an isolated environment in which Orchestrator performs verification tests.

After you create a virtual lab on a Veeam Backup & Replication server connected to your Orchestrator server, you must configure a connection to the VMware Server used to manage the lab, as described in section Connecting VMware vSphere Servers. Otherwise, Orchestrator will not be able to discover the virtual lab and make it available in the Orchestrator UI. Note that the data synchronization process between Orchestrator and the Veeam Backup & Replication server may take several minutes to complete.

> **NOTE**
>
> Virtual labs created in Veeam Backup & Replication are referred to as DataLabs in Orchestrator.

Associating DataLabs

If you want to verify machines in a DataLab that contains a lab group, you must associate the DataLab with a recovery location whose settings will be applied to the machines being verified to connect them to the correct network, to reconfigure machine IP addresses and to set the backup copy preference. For more information on these settings, see Adding VMware vSphere Recovery Locations.

To associate a DataLab with a recovery location, do the following:

1. Navigate to **DataLabs**.
2. In the **DataLab** column, select a DataLab to which you want to assign the recovery location and click **DataLab Editor**.

   For a DataLab to be displayed in the **DataLab** list, it must be added to the scope as described in section Managing Inventory Items.
3. On the **Edit DataLab** page, click **DataLab options**. The **Edit DataLab** window will open.
4. In the **Edit DataLab** window, click the link in the **Associated recovery location** field, select a recovery location with which you want to associate this DataLab, and click **Apply**.

   For a recovery location to be displayed in the list of available recovery locations, its compute resources must contain the host where the DataLab is deployed, and the location must be added in the selected scope as described in section Managing Inventory Items.

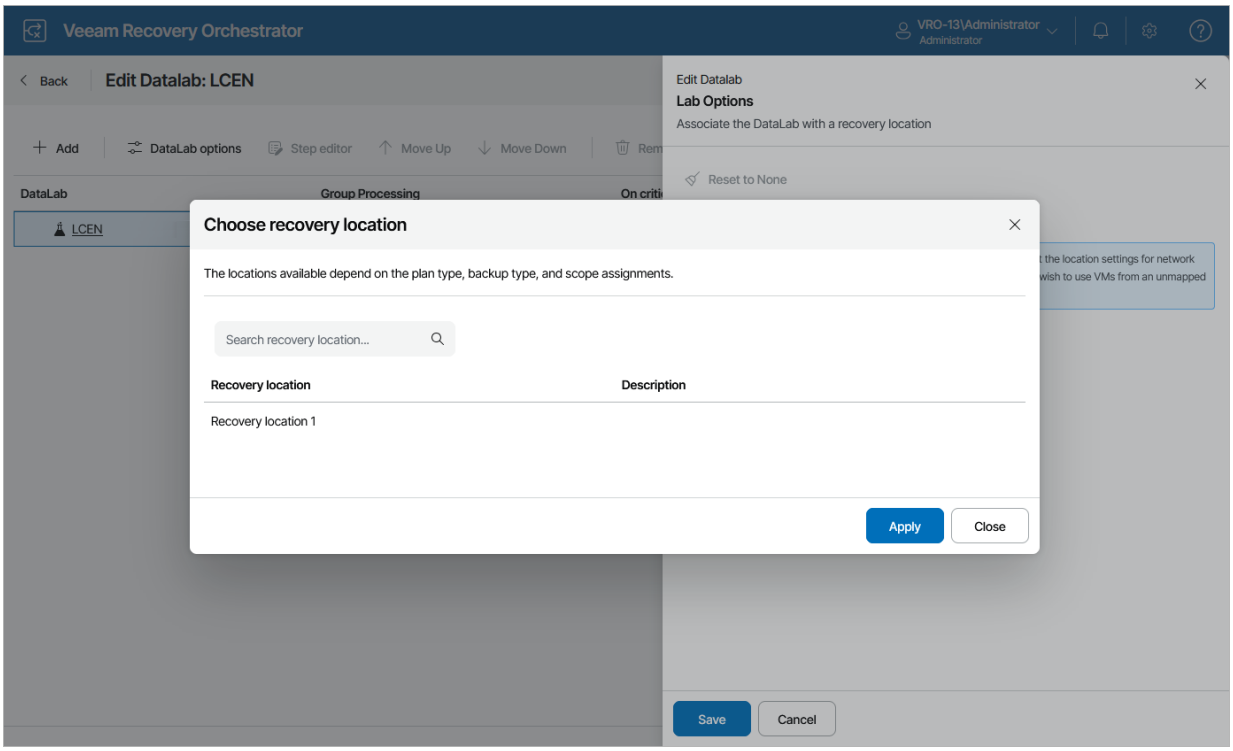

Creating Lab Groups

In most cases, a machine does not work in isolation but has dependencies on other services and components, such as Active Directory or DNS. To verify such a machine, the DataLab will have to supply all services on which this machine is dependent. For this purpose, Orchestrator uses lab groups.

> **NOTE**
>
> If a recovery plan contains a particular inventory group or machine, it is recommended that you do not test this plan in a DataLab that includes a lab group with the same machine or inventory group.

To create a lab group:

1. Navigate to **DataLabs**.
2. In the **DataLab** column, select a DataLab for which you want to create the lab group, and click **DataLab Editor**.

   For a DataLab to be displayed in the **DataLab** list, it must be added to the scope as described in section Managing Inventory Items.
3. On the **Edit DataLab** page, click **Add**.
4. In the **Add Inventory Group** window, select inventory groups that you want to include in the DataLab and click **Next**.

   For an inventory group to be displayed in the **Group** list, it must be added to the list of inventory items available for the scope, as described in section Managing Inventory Items.
5. In the **Group Options** window, do the following:

   a. In the **Group settings** section, choose whether the lab group will contain VMs recovered from backups, replicas or CDP replicas.

   b. In the **Processing logic** section, choose how Orchestrator will process VMs in this lab group. For more information on the group settings, see Configuring Groups.

   c. Review configuration information and click **Apply**.

> **NOTE**
>
> When you remove a DataLab from a scope, all lab groups in the DataLab are automatically deleted from the DataLab.

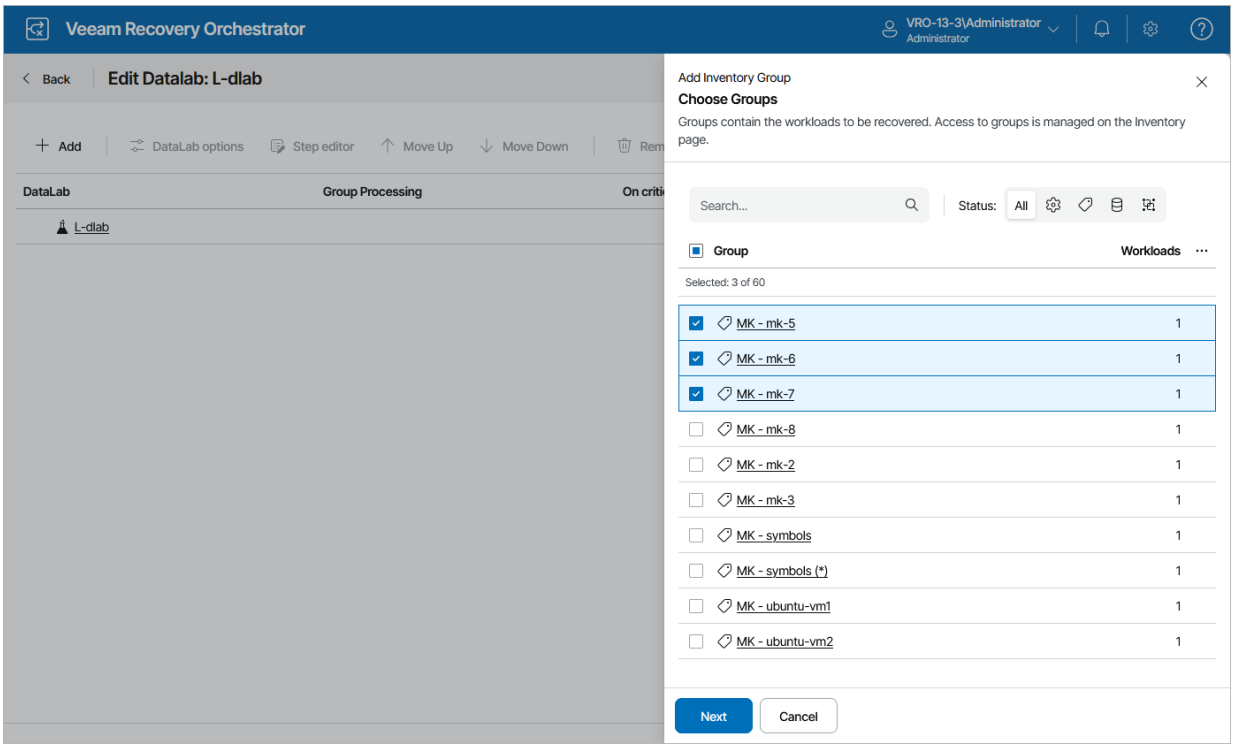

Configuring Lab Groups

If required, you can customize lab group settings in much the same way as editing a recovery plan.

1. Navigate to **DataLabs**.
2. In the **DataLabs** column, click the name of a DataLab whose groups you want to configure.

   For a DataLab to be displayed in the **DataLabs** list, it must be added to the scope as described in section Managing Inventory Items.
3. On the **DataLab Details** page, in the **Recovery plan** column, select the necessary DataLab and click **DataLab editor**.
4. On the **Edit DataLab** page, select the newly created lab group and do the following:
   - To customize the configured VM recovery options, click **DataLab options**. In the **Group Options** window, specify the required settings following the instructions provided in section Configuring Group Settings, and click **Apply**.
   - To define the order in which machines will be started, use the **Move Up** and **Move Down** arrows.
   - To select steps performed when processing each machine, expand the lab group, choose the necessary machine, click **Step editor**, and follow the instructions provided in section Configuring Steps.
   - To modify parameter settings for each step, expand the lab group, choose the necessary machine, select the step whose parameters you want to modify, click **Step editor**, and follow the instructions provided in section Configuring Step Parameters.
   - To delete the lab group, click **Remove**.
5. To save changes made to the lab group settings, click **Save Lab**.

By default, Orchestrator skips a number of steps during the plan testing process — **Generate Event**, **Send Email**, **Shutdown Source VM** and **VM Power Actions**. That is why when you create or edit a DataLab, you cannot add these steps. If you still want to add these steps, set the **During DataLab Tests** parameter value to *Execute* for each step as described in section Configuring Step Parameters.

> **IMPORTANT**
>
> A common use case for lab groups is to provide domain controllers for the test environment. If there are domain controllers in a lab group, it is essential to add the **Prepare VM as Domain Controller** step. By design, it will automatically become the first step in the step execution order.
>
> You may also optionally add domain controller-specific checks, such as **Verify Domain Controller Port** and **Verify Global Catalog Port**. These steps must be performed after the **Check Networks** step.

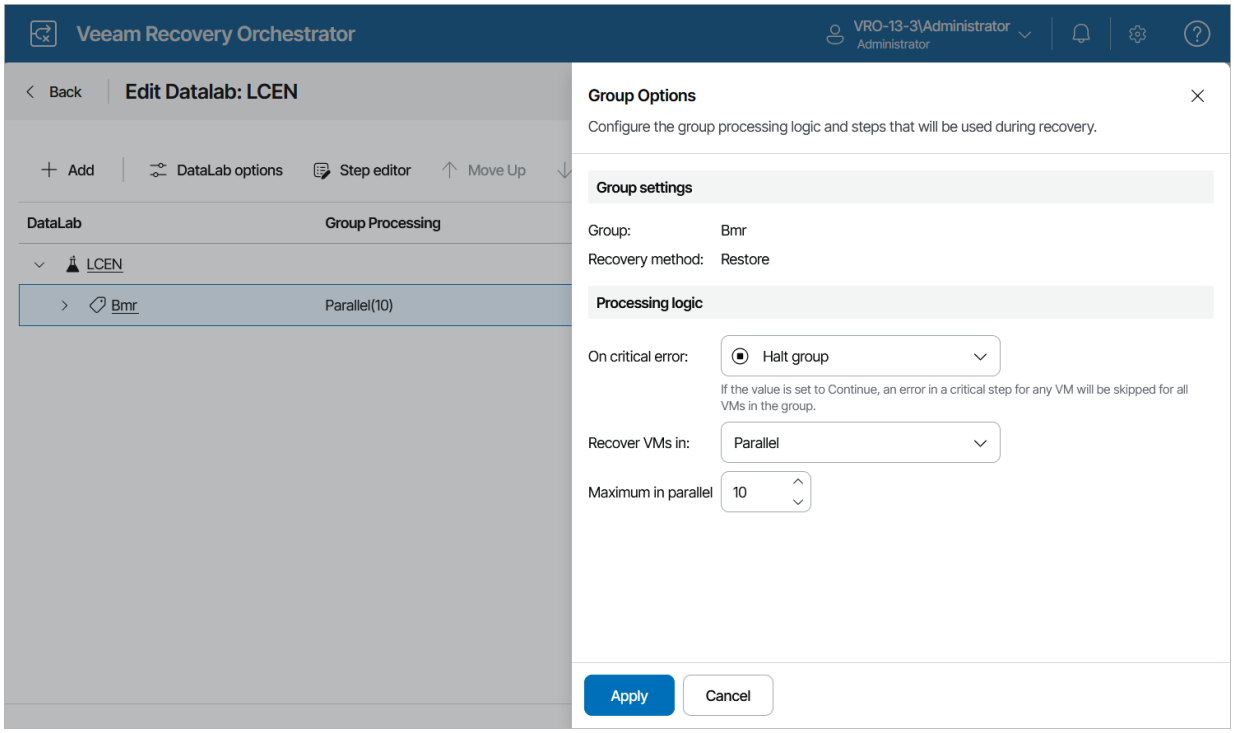

> **NOTE**
>
> There is no clear use case for replicating a domain controller. Failing over to a domain controller that contains an old version of the Active Directory database is not recommended by Microsoft. The only real use case for replicating a domain controller is to use it in an isolated lab group, and you may need to create a replication job specifically for that purpose.
>
> To learn how to restore a domain controller from an image-aware backup, see this Veeam KB article. To learn how to back up a domain controller, see this Veeam KB article.

Working with Default Lab Groups

Default lab groups are lab groups created by an Administrator but available for plan testing by Plan Authors and Plan Operators for any scope.

To allow Plan Authors and Plan Operators to use a lab group as a default group when testing plans for their scope:

1. Assign a DataLab to the scope as described in section Managing Inventory Items.
2. Add the lab group to the DataLab as described in section Creating Lab Groups.

The added lab group will now be considered a default lab group. The DataLab with the lab group will become available for the scope, and Plan Authors and Plan Operators will be able to use the group for plan testing. The lab group will be preselected in the Run Lab Tests and Create Test Schedule wizards, and it will start before all other lab groups in the DataLab every time Orchestrator powers on the lab to test a recovery plan.

> **NOTE**
>
> Plan Authors and Plan Operators cannot edit default lab groups or delete them from DataLabs because these groups can be managed only by Administrators. However, if an Administrator assigns a DataLab to the *Default Scope*, lab groups added to the DataLab will not be treated as default lab groups. This means that Plan Authors will still be able to edit and delete these lab groups as described in section Configuring Lab Groups.

Starting On-Demand Plan Test

DataLab testing can be started on-demand for any recovery plan in the *ENABLED* or *DISABLED* state. To start testing for a plan, perform the following steps:

1. Navigate to **Recovery Plans**.
2. Select the necessary plan and click **Test**.
3. In the **Run DataLab Test** window, do the following:
   a. In the **DataLab** field, select a DataLab in which the plan will be verified.For a DataLab to be displayed in the **DataLab** list, it must be added to the scope as described in section Managing Inventory Items.
   b. In the **Lab groups** field, click **Select** and add the necessary lab groups required to support the test environment. For a lab group to be displayed in the **Group** list, it must be created and configured as described in section Creating Lab Groups.

      Note that all default lab groups previously created by Administrators will automatically become preselected in the **Group** list, and you will not be able to remove them. For more information, see Working with Default Lab Groups.
   c. [Applies only to restore plans] In the **Test method** field, choose whether you want to verify both backups of plan machines and the recovery location used to restore the machines, or backups only.
      - If you select the **Quick test** option, Orchestrator will verify whether machines included in the plan will be able to recover from their backup files.

        In this case, plan machines will be verified in the Instant VM Recovery environment.
      - If you select the **Full restore to target storage** option, Orchestrator will not only verify that vSphere and agent backups are ready-to-use, but also check that the recovery location to which the machines will be restored is available and has enough resources to support the recovery process.

        In this case, you must also specify the location explicitly — to do that, click the link in the in the **Recovery location** section. For a recovery location to be displayed in the list of available recovery locations, it must be created and added to the list of inventory items available for the scope, as described in section Managing Recovery Locations.

      In both cases, Orchestrator will run all verification steps added to the plan to make sure that the plan will be able to complete successfully.

TIPS

- For the test to run successfully, you must map isolated networks of the virtual lab to all target networks present in the network mapping table of the selected recovery location. In case you want any of the recovered VMs to be connected to the same networks as the source machines, you must map isolated networks of the virtual lab to those source networks.

  To do that, configure the **Isolated Networks** settings of the virtual lab in the Veeam Backup & Replication console, as described in the Veeam Backup & Replication User Guide, section Recovery Verification.
- If you have selected the **Quick test** option at the **Test method** step of the wizard, you can only select a location that has Instant VM Recovery enabled.

  If you want to test the plan using a location with Instant VM Recovery enabled (location A) but then to restore to a location with Instant VM Recovery disabled (location B), clone the location B and change the Instant VM Recovery setting for the clone. Then, use the location A for testing and the location B for the recovery.

d. In the **Power options** field, choose an action to perform after the plan testing process is over:

   - To keep all plan VMs and the lab running in case you are willing to perform further tests, select the **Test then power off after** option. This option will book a time slot in the lab schedule and prevent other tests from being scheduled for the same period.
   - To power off all plan VMs and the lab, select the **Test then power off immediately** option.

   Note that if you have selected a starting or running lab at step 3a, the lab will not be powered off after the plan testing process is over — even if the **Test then power off immediately** is selected. In this case, Orchestrator will power off only plan VMs and keep the lab running.

e. In the **Malware detection** section, choose whether you want to check restore points created for machines included in the plan for malware flags. For restore plans, you can also decide whether you want to scan these restore points with antivirus software, a YARA rule or both.

   By design, Orchestrator checks only the most recent restore point for each machine and stops plan testing if the restore point is marked as *Suspicious* or *Infected*. However, if this restore points is created for a machine added to a lab group, Orchestrator issues a warning and continues plan testing. For more information, see the Veeam Recovery Orchestrator User Guide, section Malware Scan.

f. Review configuration information and click **Run Test**.

The lab will power on, start lab groups and begin testing the plan. If the lab halts, the plan will fail to be tested. To learn how to resume plan testing, see Managing Halted Replica Plans, Managing Halted Storage Plans, Managing Halted Restore Plans or Managing Halted CDP Replica Plans.

> **NOTE**
>
> As soon as the test is over, the DataLab Test Report will be generated. The plan and the DataLab will be powered off or will keep running, depending on the power options chosen in the **Run DataLab Test** window. Keep in mind that even if you have enabled the **Test then power off after** option, the test will be considered to be completed when all plan steps have been run, and the DataLab Test Report will be generated at that point.

If you want to receive notifications on errors that occur while powering off the DataLab, you must connect an SMTP server, add recipients and subscribe to **DataLab Test reports** as described in section Configuring Notification Settings.

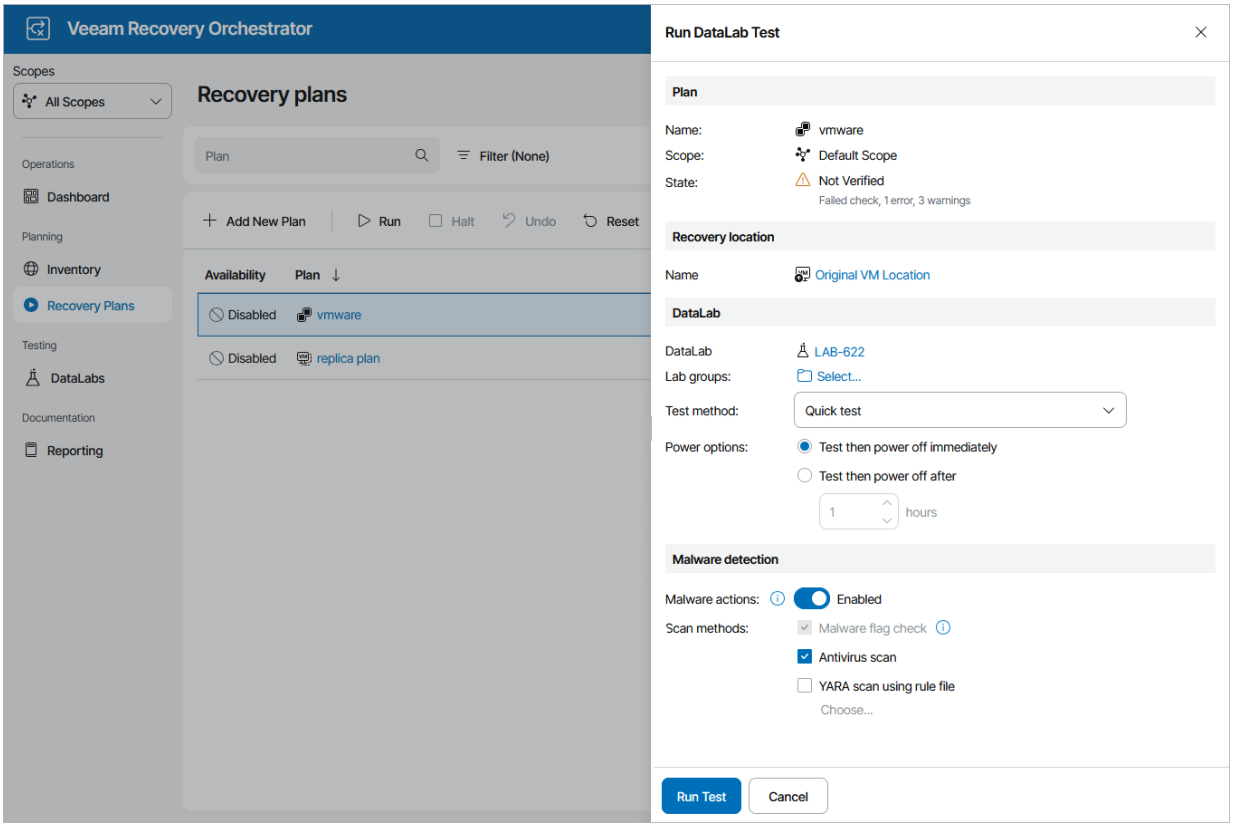

Configuring Test Scheduling

To schedule recovery plan testing:

1. Navigate to **DataLabs**.
2. Choose the DataLab for which you want to create a schedule and click **Schedule editor**.
3. In the **DataLab Schedule Editor** window, click **Add**.
4. Complete the **New Test Schedule** wizard:
    a. Specify a schedule name and description.
    b. Specify scheduling settings.
    c. Add lab groups and configure test options.
    d. Select plans you want to test.
    e. Choose whether you want to check restore points for possible malware.
    f. Finish working with the wizard.

> **NOTE**
>
> When Orchestrator tests a plan according to a specific schedule, the duration of the testing process equals the RTO value configured when creating the plan. If you instruct Orchestrator to test multiple plans at the same time, the duration of the testing process equals the maximum of the configured plan RTO values. Therefore, Orchestrator does not allow you to schedule other tests in the same DataLab until the RTO is over.

Step 1. Specify Schedule Name and Description

At the **Schedule Name** step of the wizard, use the **Name** and **Description** fields to enter a name for the new schedule and to provide a description for future reference. The maximum length of the schedule name is 128 characters; the following characters are not supported: * : / \ ? " < > | .

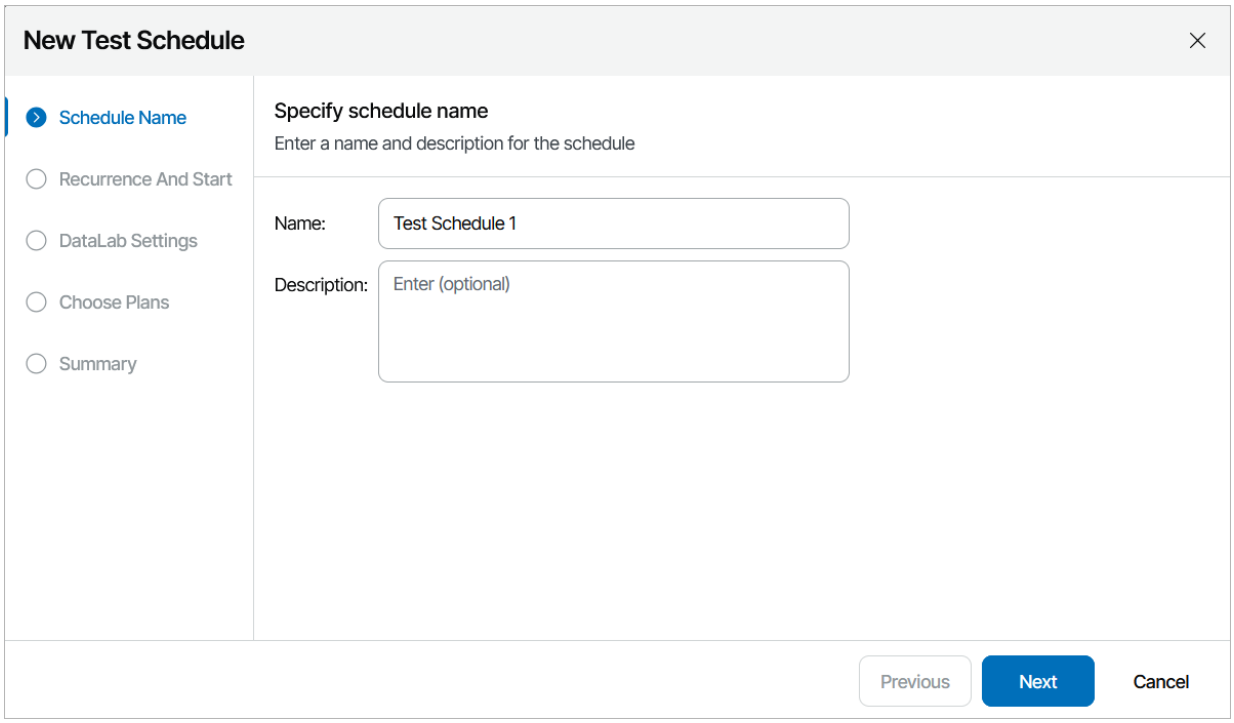

Step 2. Specify Scheduling Settings

At the **Recurrence and Start** step of the wizard, define scheduling settings for the lab:

1. Click the **Schedule** icon in the **Start on** section to configure the necessary schedule, and click **Apply**.
2. In the **Recurrence** section, choose the necessary option:
   - **Once** – to test plans once on the specified day.
   - **Weekly on** – to start testing once a week on the specified day.
   - **Monthly on** – to start testing once a month on the specified day.

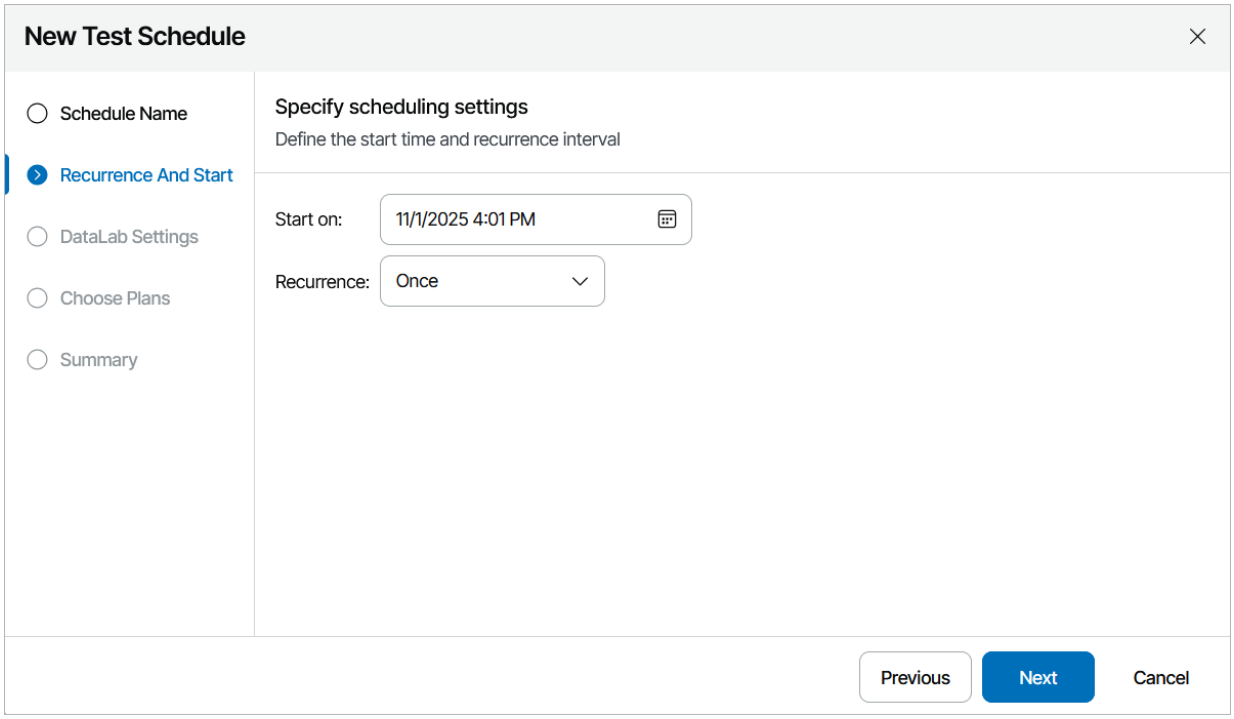

Step 3. Add Lab Groups

At the **DataLab Settings** step of the wizard, add the required lab groups to support the test environment.

For a lab group to be displayed in the **Group** list, it must be created and configured as described in section Creating Lab Groups.

> **NOTE**
>
> All default lab groups previously created by an Administrator will automatically become preselected in the **Lab Groups to use** list, and you will not be able to remove them. For more information, see Working with Default Lab Groups.

Additionally, you can configure the following settings:

- From the **Test method** drop-down list, choose whether you want to verify both backups of plan machines and the recovery location used to restore the machines, or backups only.
  - If you select the **Instant VM recovery** option, Orchestrator will check whether machines included in the plan will be able to recover from their backup files.

    In this case, plan machines will be verified in the Instant VM Recovery environment.
  - If you select the **Full restore to target storage** option, Orchestrator will not only verify that vSphere and agent backups are ready-to-use, but also check that the recovery location to which the machines will be restored is available and has enough resources to support the recovery process.

    In this case, Orchestrator will run all verification steps added to the plan to make sure that the plan will be able to complete successfully.

> **TIP**
>
> For the test to run successfully, you must map isolated networks of the virtual lab to all target networks present in the network mapping table of the selected recovery location. In case you want any of the recovered VMs to be connected to the same networks as the source machines, you must map isolated networks of the virtual lab to those source networks.
>
> To do that, configure the isolated networks settings of the virtual lab in the Veeam Backup & Replication console, as described in the Veeam Backup & Replication User Guide, section Recovery Verification.

- In the **Power Options** field, choose an action to perform after the plan testing process is over:
  - To power off all plan machines and the lab, select the **Test then power off immediately** option.
  - To keep plan all machines and the lab running in case you are willing to perform further tests, select the **Test then power off after** option.

    Use the Test then power off after field to book a time slot in the lab schedule and prevent other tests from being scheduled for the same period.

> **NOTE**
>
> If you have selected a starting or running lab, the lab will not be powered off after the plan testing process is over — even if the **Test then power off immediately** option is selected. In this case, Orchestrator will power off only plan machines and keep the lab running.

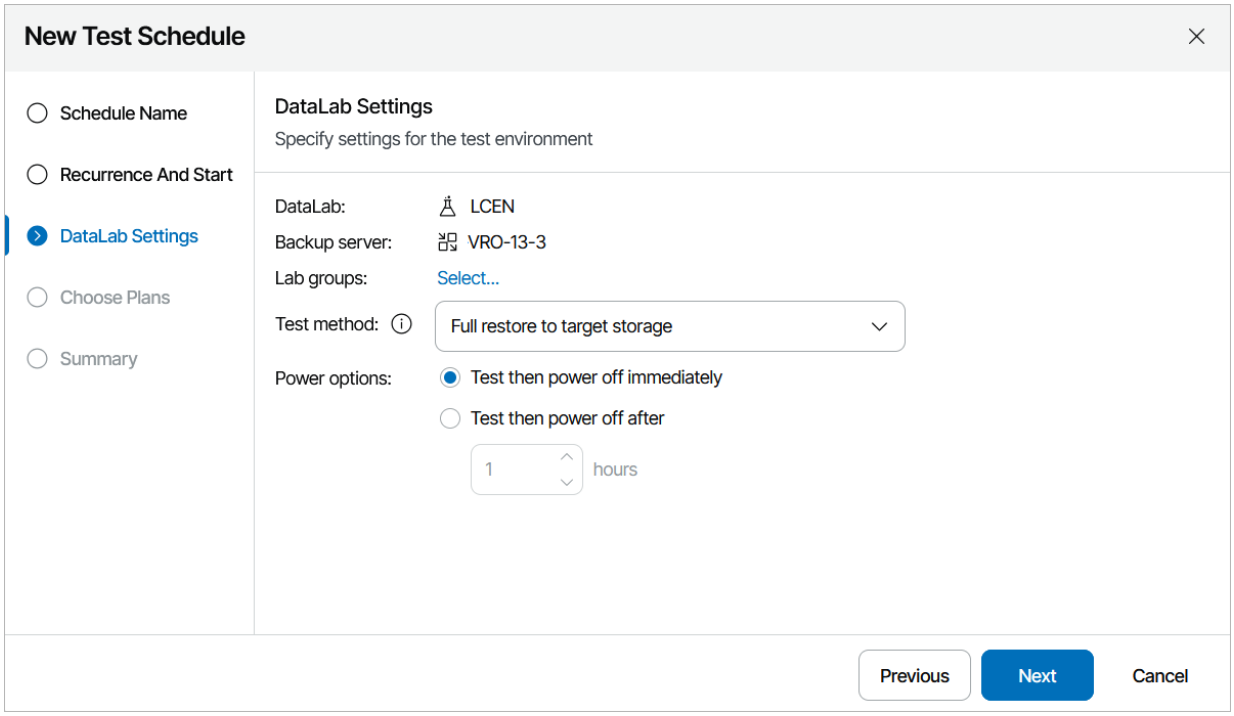
New Test Schedule
Schedule Name
Recurrence And Start
DataLab Settings
Choose Plans
Summary
DataLab Settings
Specify settings for the test environment
DataLab:
LCEN
Backup server:
VRO-13-3
Lab groups:
Select...
Test method:
Full restore to target storage
Power options:
Test then power off immediately
Test then power off after
1
hours
Previous
Next
Cancel

Step 4. Select Plans

At the **Choose Plans** step of the wizard, select recovery plans to be tested in a DataLab.

> **NOTE**
>
> If you select multiple plans, they all will be tested at the same time.

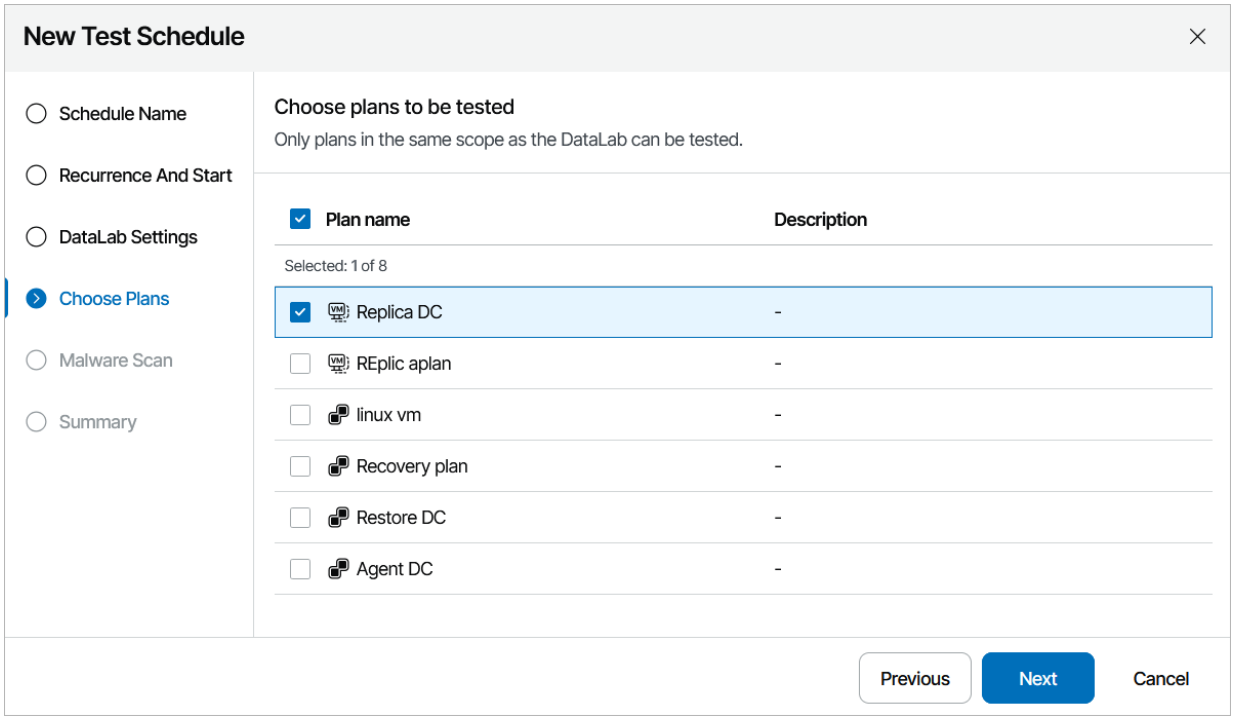

Step 5. Run Malware Scan

[This step applies only if you have included at least one restore, replica or CDP replica plan in the **Plan name** list at the Choose Plans step of the wizard]

At the **Malware Scan** step of the wizard, choose whether you want to check restore points created for machines included in the plan for possible for malware flags. You can also decide whether you want to scan these restore points with antivirus software, a YARA rule or both.

By design, Orchestrator checks only the most recent restore point for each machine and stops plan testing if the restore point is marked as *Suspicious* or *Infected*. For more information, see the Veeam Recovery Orchestrator User Guide, section Malware Scan.

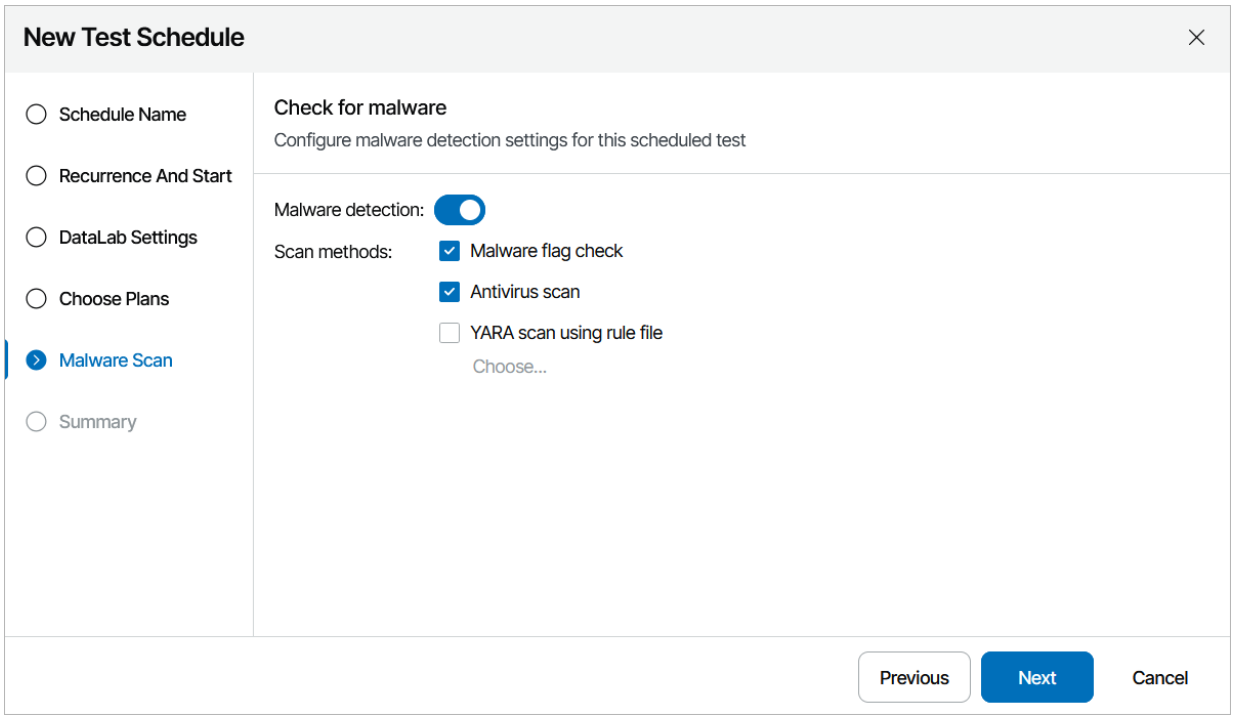

Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.

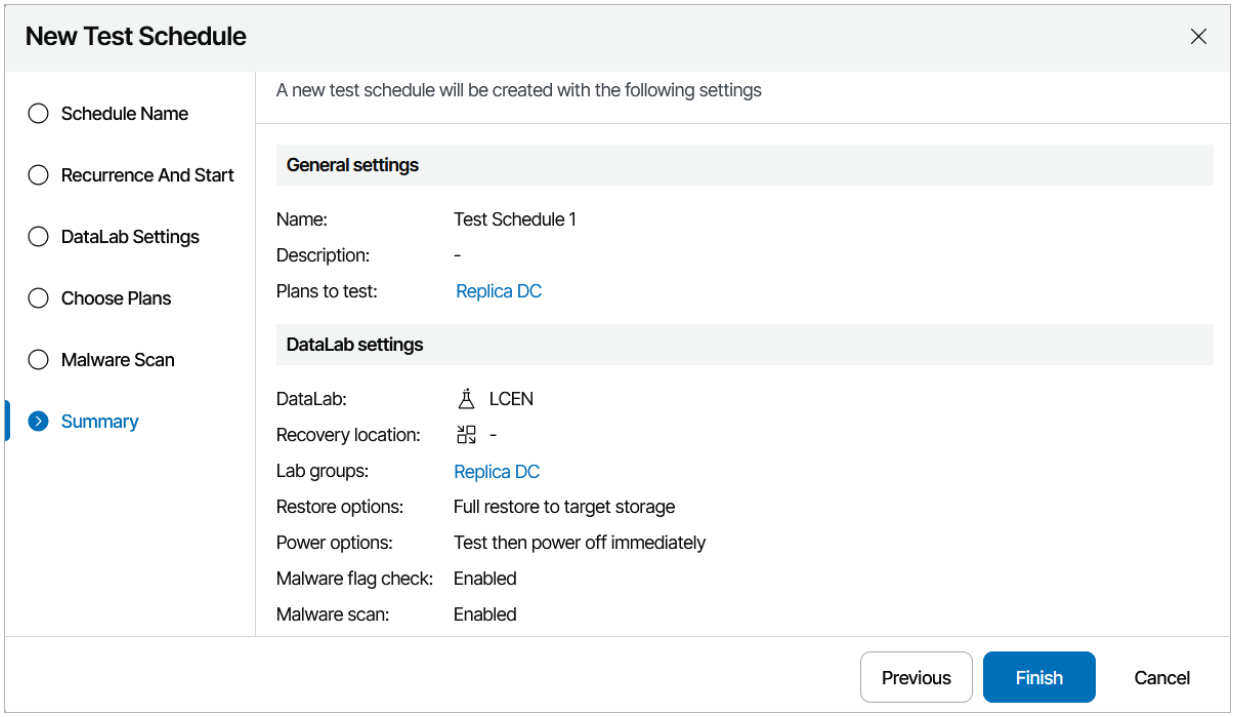

Scanning Recovery Plans

You can check a recovery plan for possible malware before or when running a plan.

IMPORTANT

- To scan CDP replica plans for malware, the Veeam Backup & Replication server that manages replication jobs must run version 13.0.1 or later.
- Scanning storage plans for possible malware is not supported.

Configuring Scan Scheduling

You can schedule a time to perform a malware scan for a plan. To do that:

1. Navigate to **Recovery Plans**.
2. Select the plan that you want to scan for malware.
3. From the **Manage** menu, select **Schedule**.

   OR-

   Right-click the plan name and select **Manage** > **Schedule**.
4. In the **Scheduled Tasks** window, set the **Schedule malware detection** toggle to *Enabled* and do the following:
   a. Click the **Configure schedule** link to set the necessary schedule and specify recurrence settings, and click **Save**.
   b. In the **Restore points** field, specify the number of restore points on each machine that you want to scan for malware.

      By default, Orchestrator checks the most recent restore point. However, if you specify to scan more than 1 restore point, Orchestrator will perform scanning starting from the earliest available restore point to the most recent one. If all of the restore points are infected, the plan will acquire the *NOT VERIFIED* state after the scan process completes.

      For more information on the way Orchestrator chooses restore points for malware scan, the Veeam Recovery Orchestrator User Guide, section How Orchestrator Selects Restore Points During On-Demand Malware Scan.
   c. [Applies only to restore and cloud plans] Choose whether you want to scan the restore points created for machines included in the plan with antivirus software, a YARA rule or both.
   d. [Applies only if you have selected **Virus scan**, **YARA scan using rule file** or both] By design, Orchestrator scans the restore point until a virus or YARA rule match is detected. Then, Orchestrator either completes the scanning session or proceeds to the next restore point in case you have specified several restore points to scan at step 4b. However, you can instruct Orchestrator to continue scanning the restore point until all viruses and YARA rule matches are detected. To do that, select the **Full image scan** check box.

e. To save changes made to the plan malware schedule, click **Save**.

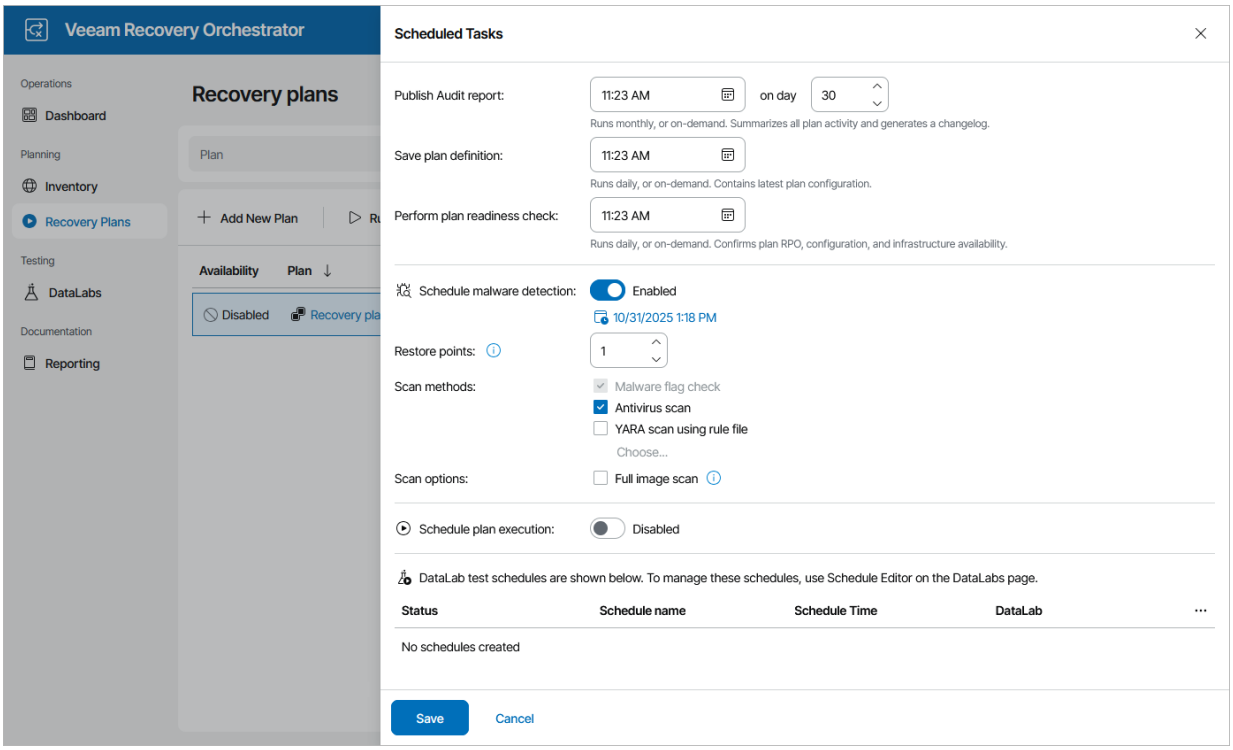

Starting On-Demand Plan Scan

Scanning for malware may be started on-demand for a recovery plan in the *ENABLED* or *DISABLED* state. To start scanning for a plan, perform the following steps:

1. Navigate to **Recovery Plans**.
2. Select the necessary plan and click **Scan**.
3. In the **Malware Scan** window, do the following:
   a. In the **Restore Point** section, choose a restore point that you want to scan.
   b. In the **Malware detection** section, configure the following settings:
      i. In the **Restore points** field, specify the number of restore points on each machine that you want to scan for malware.

         By default, Orchestrator checks the most recent restore point. However, if you specify to scan more than 1 restore point, Orchestrator will perform scanning starting from the earliest available restore point to the most recent one. If all of the restore points are infected, the plan will acquire the *NOT VERIFIED* state after the scan process completes.

         For more information on the way Orchestrator chooses restore points for malware scan, the Veeam Recovery Orchestrator User Guide, section How Orchestrator Selects Restore Points During On-Demand Malware Scan.

      ii. [Applies only to restore and cloud plans] Decide whether you want to scan these restore points with antivirus software, YARA rules or both.

      ii. [Applies only if you have selected **Virus scan**, **YARA scan using rule file** or both] By design, Orchestrator scans the restore point until a virus or YARA rule match is detected. Then, Orchestrator either completes the scanning session or proceeds to the next restore point in case you have specified several restore points to scan at step 4b. However, you can instruct Orchestrator to continue scanning the restore point until all viruses and YARA rule matches are detected. To do that, select the **Full image scan** check box.

c. Review configuration information and click **Run**.

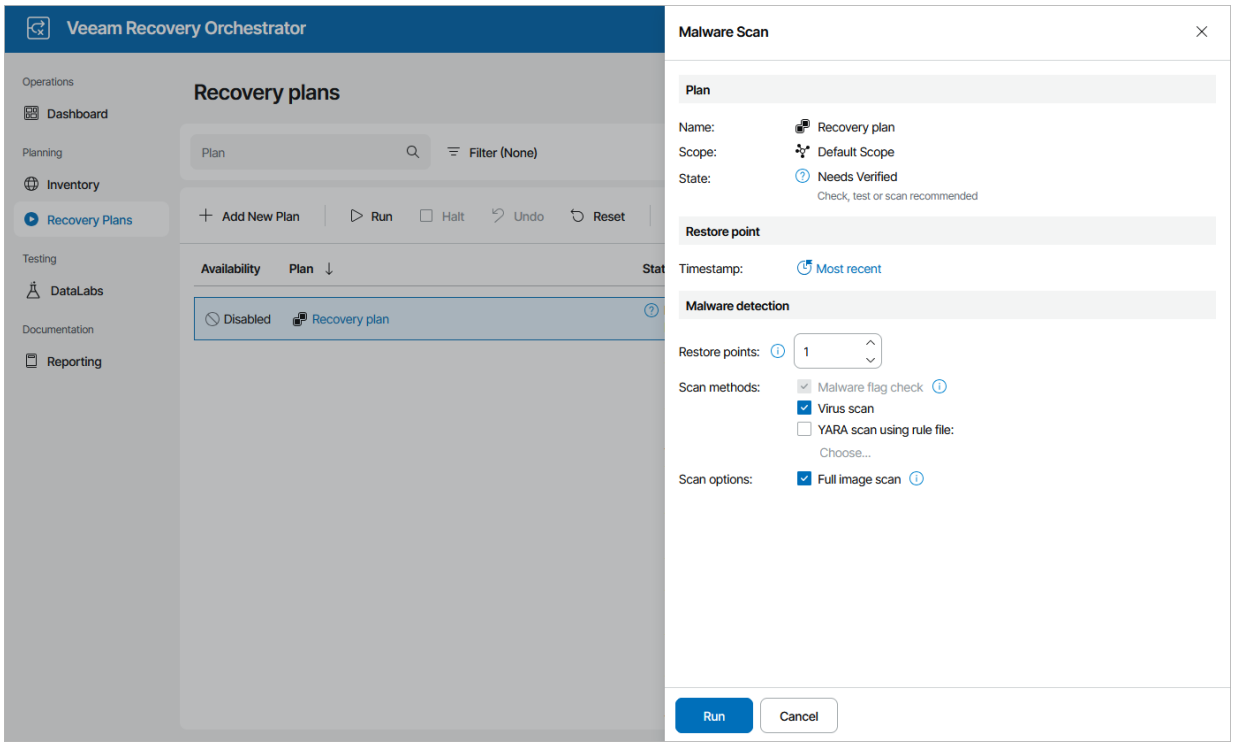

Generating Reports

Orchestrator comes with a number of reports that allow you to:

- Obtain plan configuration and change tracking data. For more information, see Generating Plan Definition Report.
- Check plan configuration before running recovery plans. For more information, see Running Plan Readiness Check.
- Obtain the results of plan testing and execution. For more information, see Viewing DataLab Test Results and Viewing Plan Execution History.
- Obtain data on all activity for a specific period. For more information, see Viewing Audit Report .
- Obtain the results of malware scan. For more information, see Generating Malware Scan Report.

You can then use the reports to send them by email to engineers, auditors and managers, and to troubleshoot issues that prevent the recovery process from completing successfully. Reports are sent as PDF files attached to report notifications. To learn how to add recipients to whom notifications will be sent, see Configuring Notification Settings.

Before You Begin

All reports generated by Orchestrator are prefixed by a cover page — a report template that you select when creating a plan. This template can be edited in-line in Orchestrator using the Microsoft Word integration. When reports are generated, they are appended to the cover page template.

Orchestrator includes 1 instance of the default report template that comes in English. The default template instance contains example text and can be used as is — however, it is recommended that you clone and customize a template for your specific needs.

Note that not all parts of the default template can be modified. Some sections are visible but cannot be edited by users. These sections are automatically filled out with the plan information when a report is generated.

Managing Templates

By design, you cannot customize a Veeam default template instance itself. To generate a recovery plan report based on a modified template, you must create a clone of an out-of-the-box template, edit the template using Microsoft Word integration, and then select it as the **Report Template** for the plan.

1. Navigate to **Reporting** > **Templates**.
2. In the **Template** column, select the default template and click **Clone**. This will create a copy of the template.
3. In the **Clone Template** window, enter a name and description for the new template, select a scope for which the template will be available, and click **Clone**.
5. Select the new template and click **Edit in Word**. This will launch Microsoft Word.

   If you are prompted for a password, specify the credentials that you used to access the Orchestrator UI.

> **IMPORTANT**
>
> To allow the Microsoft Word integration, Microsoft Word component of SP2 for Microsoft Office 2010 or later must be installed on the machine where the Orchestrator UI runs.

6. Customize the default template as required and save the document. All changes will be automatically uploaded to Orchestrator.

   If you want to include plan properties in the report based on the customized template, you can insert the following dynamic variables: *~Created*, *~TimeZone*, *~PlanType*, *~PlanName*, *~PlanDescription*, *~PlanContactName*, *~PlanContactEmail*, *~PlanContactTel*, *~Site*, *~SiteScopeName*, *~SiteDescription*, *~SiteContactName*, *~SiteContactEmail*, *~SiteContactTel*, *~ServerName*, *~MachinesInPlan*, *~GroupsInPlan*, *~ReportType*, *~TargetRTO* and *~TargetRPO*. To populate these variables while generating the report output, Orchestrator will use properties specified during the plan creation process.

   To insert a variable:

   a. Switch to the **Developer** tab. By default, the Microsoft Word ribbon does not show the tab. To display the tab:

      i. Click **File** > **Options**.

      ii. In the **Word Options** window, switch to the **Customize Ribbon** tab, select the **Developer** check box in the **Main Tabs** list, and click **OK**.

   b. Select text areas where you want to insert the variable.

   c. Click the **Rich Text Content Control** button.

   d. In the control field, enter the required variable.

> **NOTE**
>
> Orchestrator reports do not support Microsoft Word interactive elements (such as comments, footnotes and charts). If you include such elements in a template, they will not be included in the resulting report.

7. Navigate to **Recovery Plans**.

8. Select the modified template as a **Report Template** for the plan. To do that, follow the instructions provided in section Creating Replica Plans, Creating CDP Replica Plans, Creating Restore Plans, Creating Storage Plans or Creating Cloud Plans.

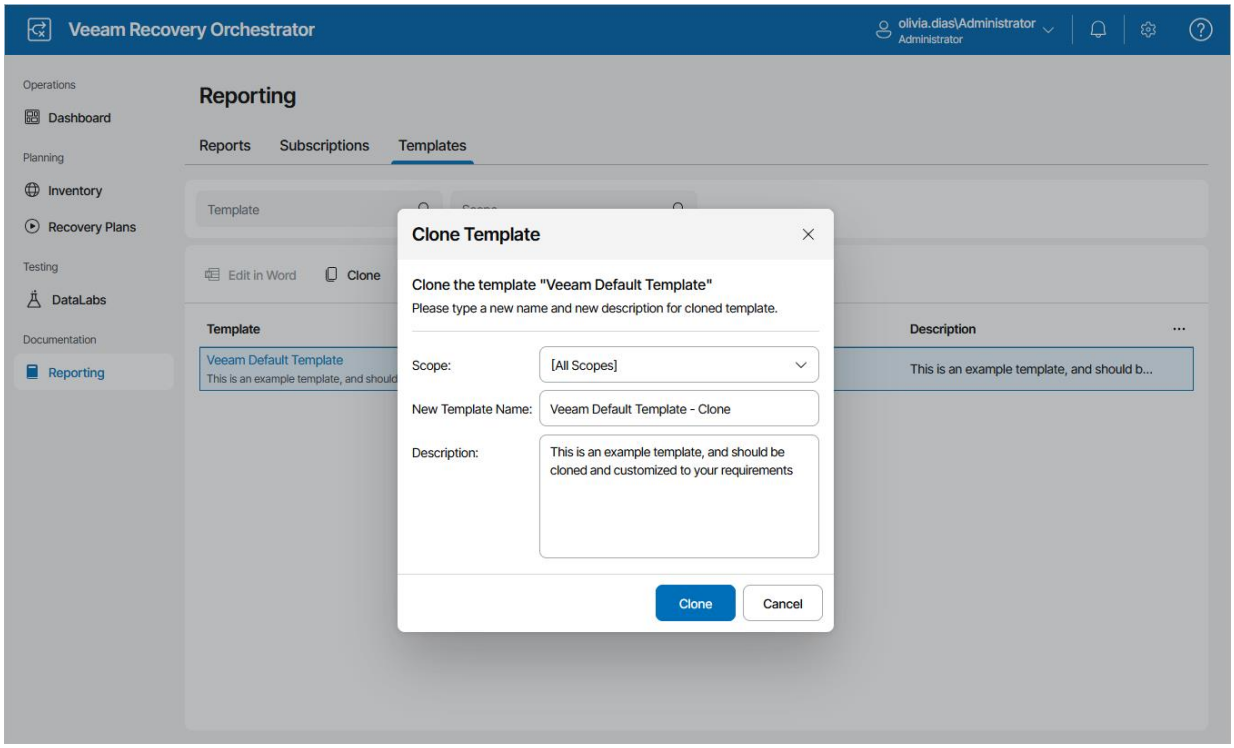

Generating Plan Definition Report

As soon as you create a recovery plan, you will be able to generate the **Plan Definition Report**. The report provides an easily shareable view of all inventory groups, as well as steps and parameters defined by the plan.

This document is ideal for auditors and managers, and can be used to obtain a sign-off from application owners who need to verify plan configuration.

Orchestrator generates two types of reports:

- A summary report that includes a plan overview and a summary of inventory groups included in the plan with drill-down hyperlinks to individual machines.
- A full report that also includes details on the recovery location specified for the plan, information on specific steps that will run during the recovery process and the plan change log, which allows you to track who changed plan settings, when and what was changed.

Updating Definition Reports

By default, Orchestrator runs the Plan Definition Report automatically for every *ENABLED* recovery plan daily. You can also generate the report for a plan on demand:

1. Navigate to **Recovery Plans**.
2. Select the plan.
3. From the **Publish** menu, select **Plan Definition Report**.

   -OR-

   Right-click the plan and select **Plan Definition Report** from the drop-down menu.

> **NOTE**
>
> The **Plan Definition Report** link will be unavailable in case the plan is being edited.

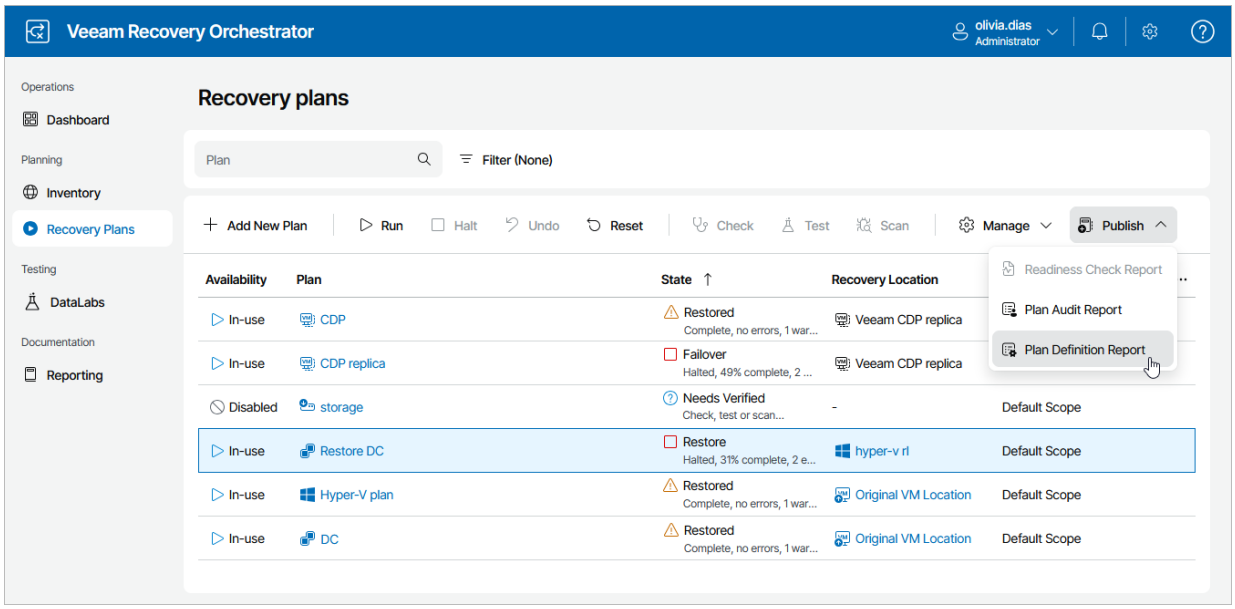

Downloading Plan Definition Reports

To access the report for a recovery plan:

1. Navigate to **Reporting**.
2. Select the report.
3. Click the plan name to download a summary report.

   -OR-

   Click **Download** and choose whether you want to download a summary or full report.

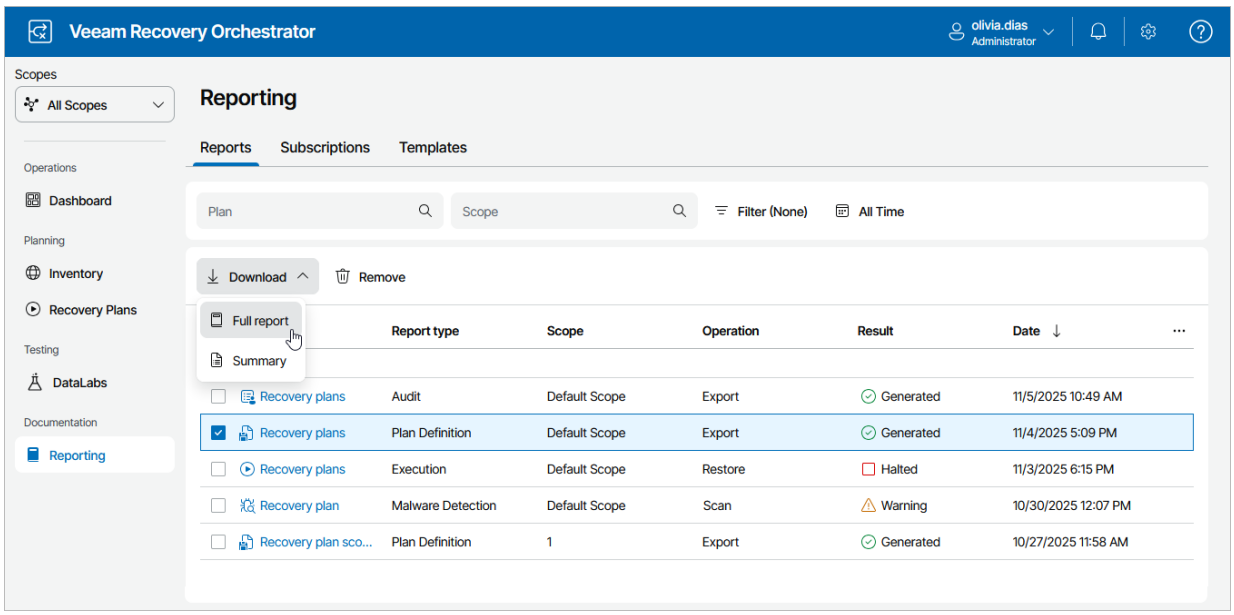

The Plan Definition Report will use the default report template or a custom template. The plan definition will be appended at the end of the template.

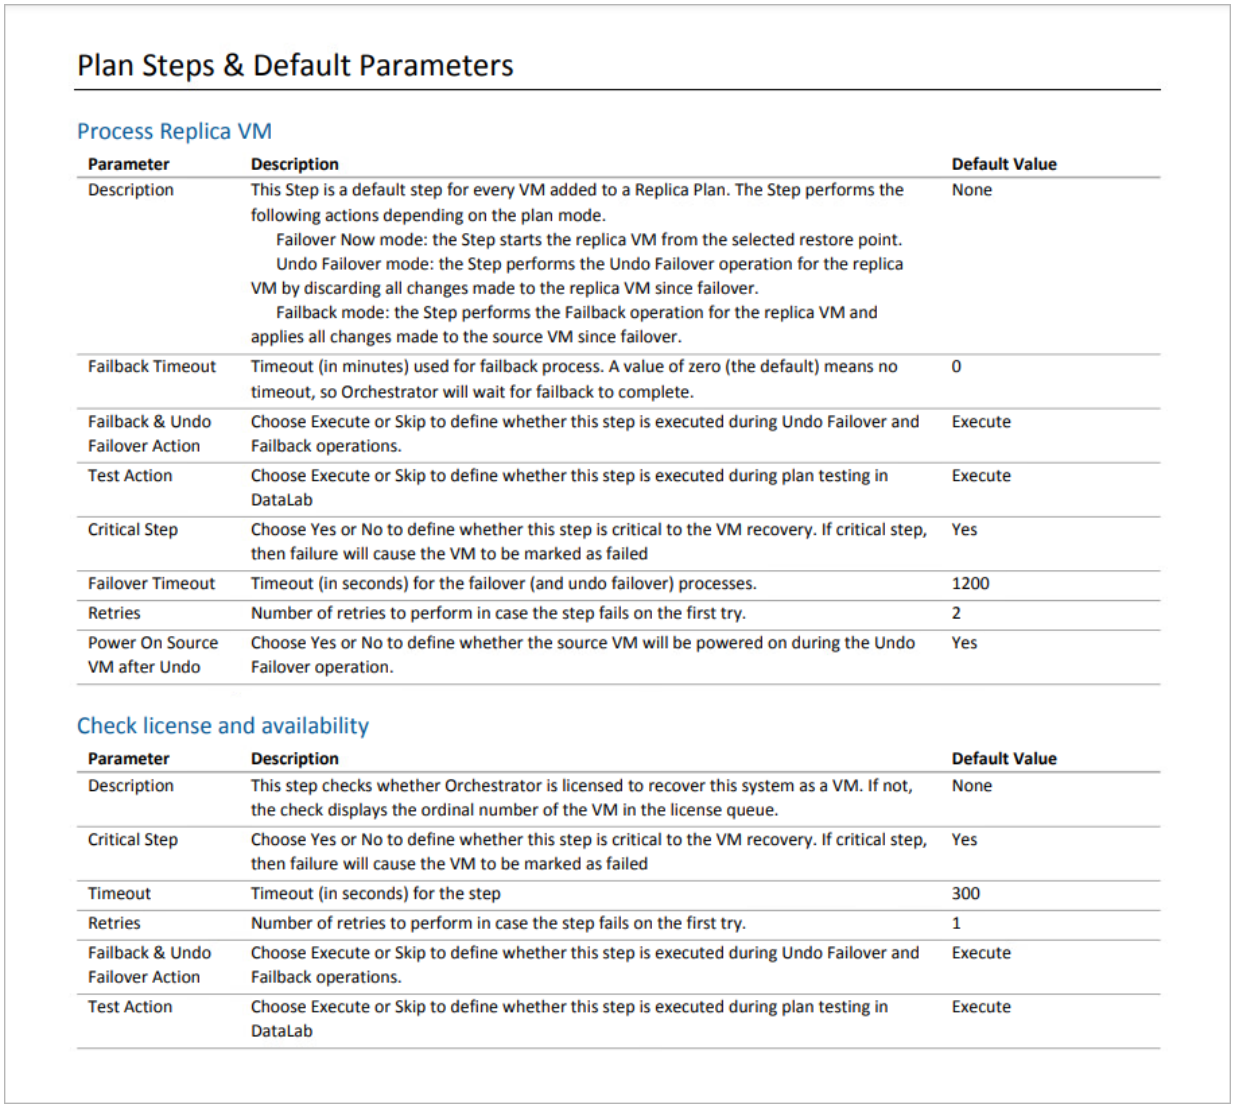

Plan Steps & Default Parameters

Process Replica VM

Parameter	Description	Default Value
Description	This Step is a default step for every VM added to a Replica Plan. The Step performs the following actions depending on the plan mode. Failover Now mode: the Step starts the replica VM from the selected restore point. Undo Failover mode: the Step performs the Undo Failover operation for the replica VM by discarding all changes made to the replica VM since failover. Failback mode: the Step performs the Failback operation for the replica VM and applies all changes made to the source VM since failover.	None
Failback Timeout	Timeout (in minutes) used for failback process. A value of zero (the default) means no timeout, so Orchestrator will wait for failback to complete.	0
Failback & Undo Failover Action	Choose Execute or Skip to define whether this step is executed during Undo Failover and Failback operations.	Execute
Test Action	Choose Execute or Skip to define whether this step is executed during plan testing in DataLab	Execute
Critical Step	Choose Yes or No to define whether this step is critical to the VM recovery. If critical step, then failure will cause the VM to be marked as failed	Yes
Failover Timeout	Timeout (in seconds) for the failover (and undo failover) processes.	1200
Retries	Number of retries to perform in case the step fails on the first try.	2
Power On Source VM after Undo	Choose Yes or No to define whether the source VM will be powered on during the Undo Failover operation.	Yes

Check license and availability

Parameter	Description	Default Value
Description	This step checks whether Orchestrator is licensed to recover this system as a VM. If not, the check displays the ordinal number of the VM in the license queue.	None
Critical Step	Choose Yes or No to define whether this step is critical to the VM recovery. If critical step, then failure will cause the VM to be marked as failed	Yes
Timeout	Timeout (in seconds) for the step	300
Retries	Number of retries to perform in case the step fails on the first try.	1
Failback & Undo Failover Action	Choose Execute or Skip to define whether this step is executed during Undo Failover and Failback operations.	Execute
Test Action	Choose Execute or Skip to define whether this step is executed during plan testing in DataLab	Execute

Running Plan Readiness Check

Readiness Check is a very low-impact and fast method to confirm that configuration of a recovery plan matches the DR environment, and therefore the plan should run successfully.

The readiness check will work through every plan step to perform specific checks against each item included in a plan. It allows you to ensure the following:

- Storage systems are detected and prepared for failover.
- Datastores included in storage plans are protected by storage replication.
- Replica VMs are detected and ready for failover.
- Backups are detected and ready for restore.
- Veeam Backup & Replication servers are online and available.
- Infrastructure such as VMware vCenter, SCVMM server, NetApp and HPE storage is online and available.
- Required credentials are provided.
- Required step parameters are configured.

The readiness check is almost zero-impact and completes very quickly. It can therefore be run very frequently. For example, it is recommended that you run the readiness check in the following cases:

- After you create a plan, run the readiness check to verify whether the plan will be able to run successfully.
- After you edit a plan, run the readiness check to confirm that the changes are valid.
- After you test a plan in a DataLab, run the readiness check to confirm that replicas were shut down successfully and are ready for failover.
- After you make some changes to the virtual infrastructure, run the readiness check to confirm that recovery locations used for plans still have available resources to complete the recovery process.

Orchestrator generates two types of reports:

- A summary report that includes a plan overview and a summary of inventory groups included in the plan with drill-down hyperlinks to specific machines and color-coded results of checking every plan step.
- A full report that also includes details on the recovery location specified for the plan and information on specific steps that will run during the recovery process.

Running Plan Readiness Check Manually

By default, Orchestrator runs the readiness check automatically for every *ENABLED* recovery plan daily. To run the check manually for a plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Check**.

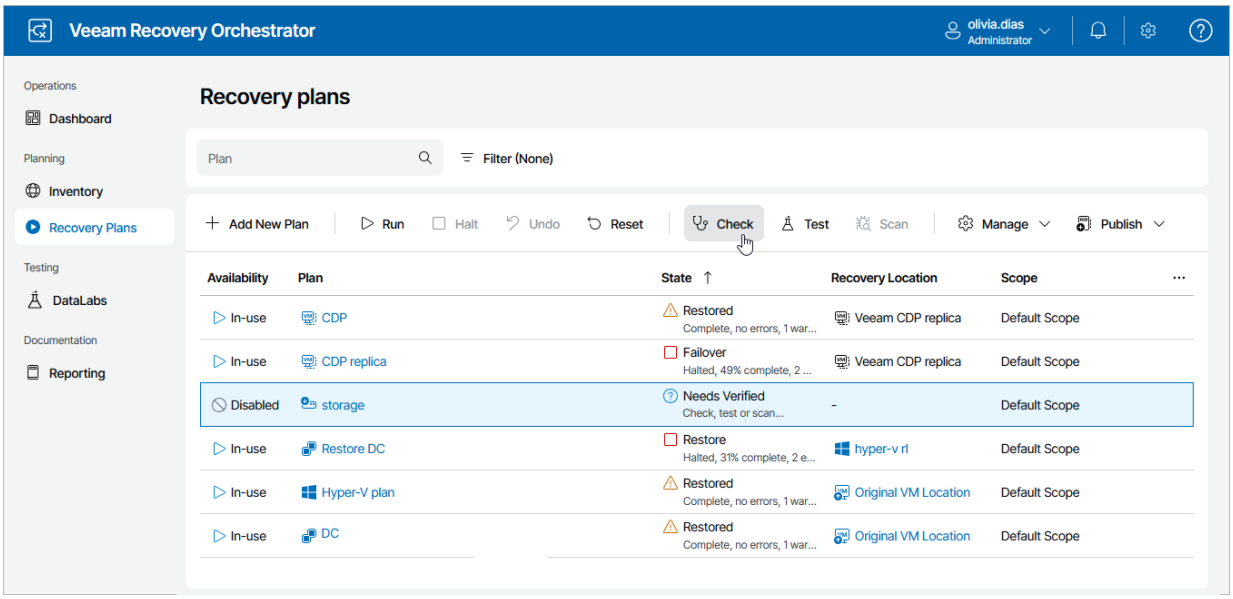

As soon as the readiness check completes, the **State** column will display the check result. The state information (*Verified*, *Needs Verified* or *Not Verified*) is a rollup of the Readiness Check and DataLab test results.

> **TIP**
>
> Summary information on readiness check results over all scopes will be also available on the Home Page Dashboard.

Downloading Plan Readiness Check Reports

To view details of the readiness check for a recovery plan:

1. Navigate to **Reporting**.
2. Select the report.
3. Click the plan name to download a summary report.

   -OR-

Click **Download** and choose whether you want to download a summary or full report.

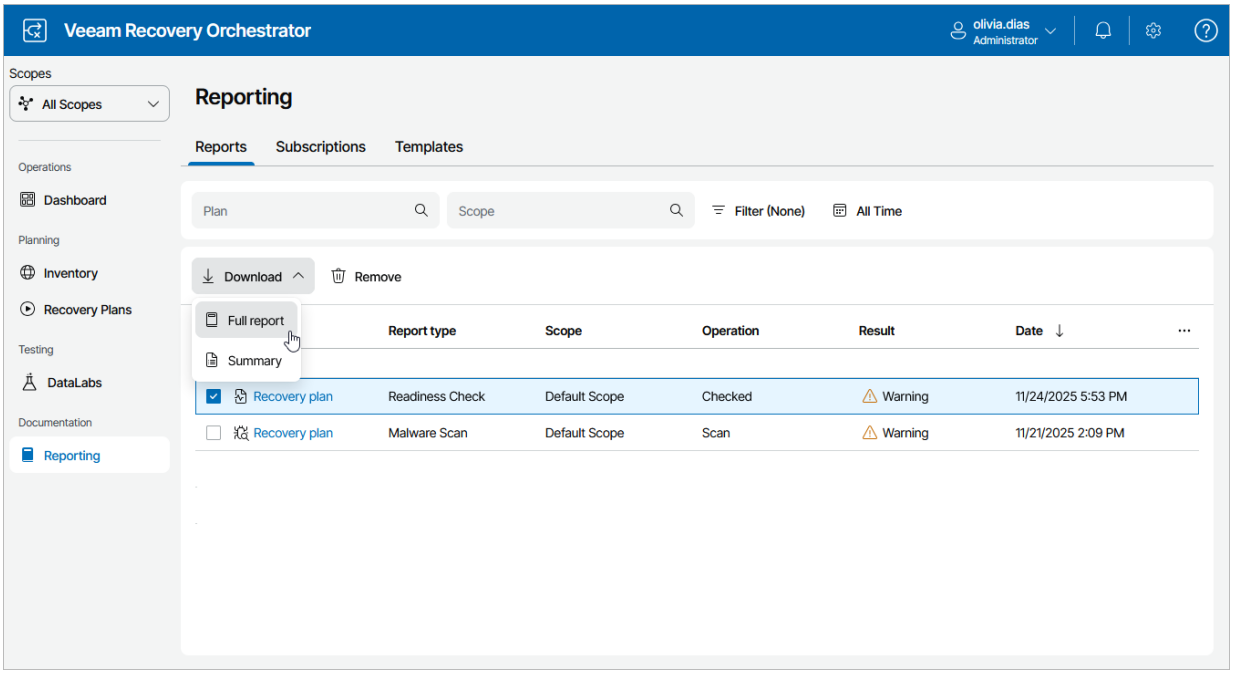

The Readiness Check Report will use the default report template or a custom template. The results of all readiness checks will be appended at the end of the template.

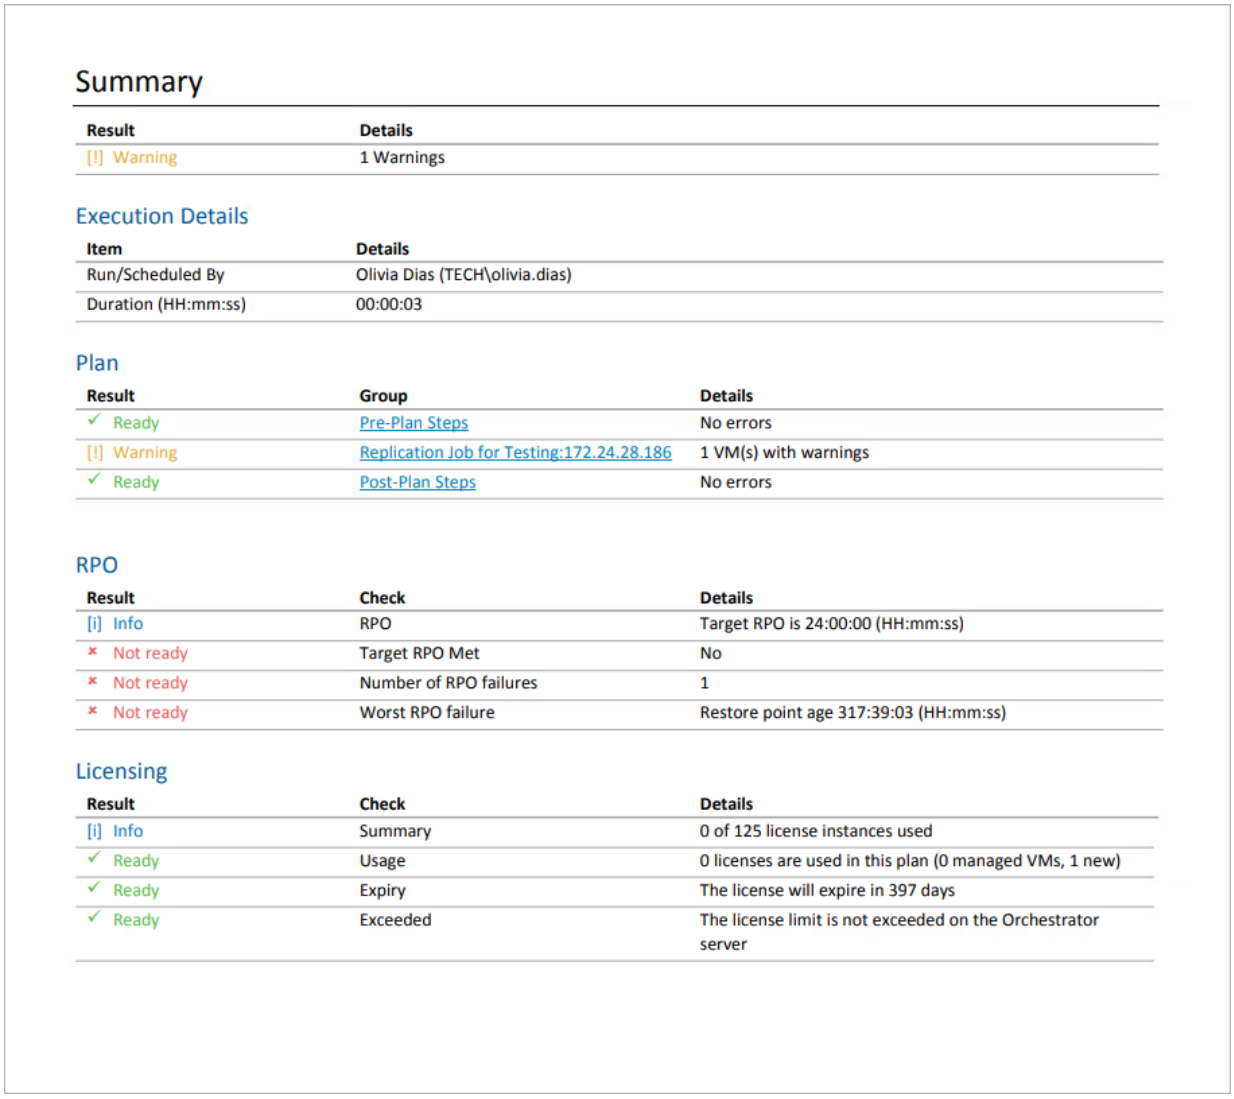

Summary

Result	Details
[!] Warning	1 Warnings

Execution Details

Item	Details
Run/Scheduled By	Olivia Dias (TECH\olivia.dias)
Duration (HH:mm:ss)	00:00:03

Plan

Result	Group	Details
✓ Ready	Pre-Plan Steps	No errors
[!] Warning	Replication Job for Testing:172.24.28.186	1 VM(s) with warnings
✓ Ready	Post-Plan Steps	No errors

RPO

Result	Check	Details
[i] Info	RPO	Target RPO is 24:00:00 (HH:mm:ss)
× Not ready	Target RPO Met	No
× Not ready	Number of RPO failures	1
× Not ready	Worst RPO failure	Restore point age 317:39:03 (HH:mm:ss)

Licensing

Result	Check	Details
[i] Info	Summary	0 of 125 license instances used
✓ Ready	Usage	0 licenses are used in this plan (0 managed VMs, 1 new)
✓ Ready	Expiry	The license will expire in 397 days
✓ Ready	Exceeded	The license limit is not exceeded on the Orchestrator server

Viewing DataLab Test Results

After you test a plan in an isolated Orchestrator DataLab, Orchestrator will generate the **DataLab Test Report**. The report contains test execution details and provides information on configured test environment. Summary information on plan test results for all scopes will be also available on the Home Page Dashboard.

Orchestrator generates two types of reports:

- A summary report that includes a plan overview and a summary of inventory groups included in the plan with drill-down hyperlinks to specific machines and color-coded results of testing every plan step.
- A full report that also includes details on the DataLab appliance and specific steps that will run during the recovery process. For every group, machine and step included in the plan, the processing start time and duration will be recorded.

> **TIP**
>
> Summary information on DataLab test results over all scopes will be also available on the Home Page Dashboard.

Downloading DataLab Test Reports

To access the report for a recovery plan:

1. Navigate to **Reporting**.
2. Select the report.
3. Click the plan name to download a summary report.

   -OR-

   Click **Download** and choose whether you want to download a summary or full report.

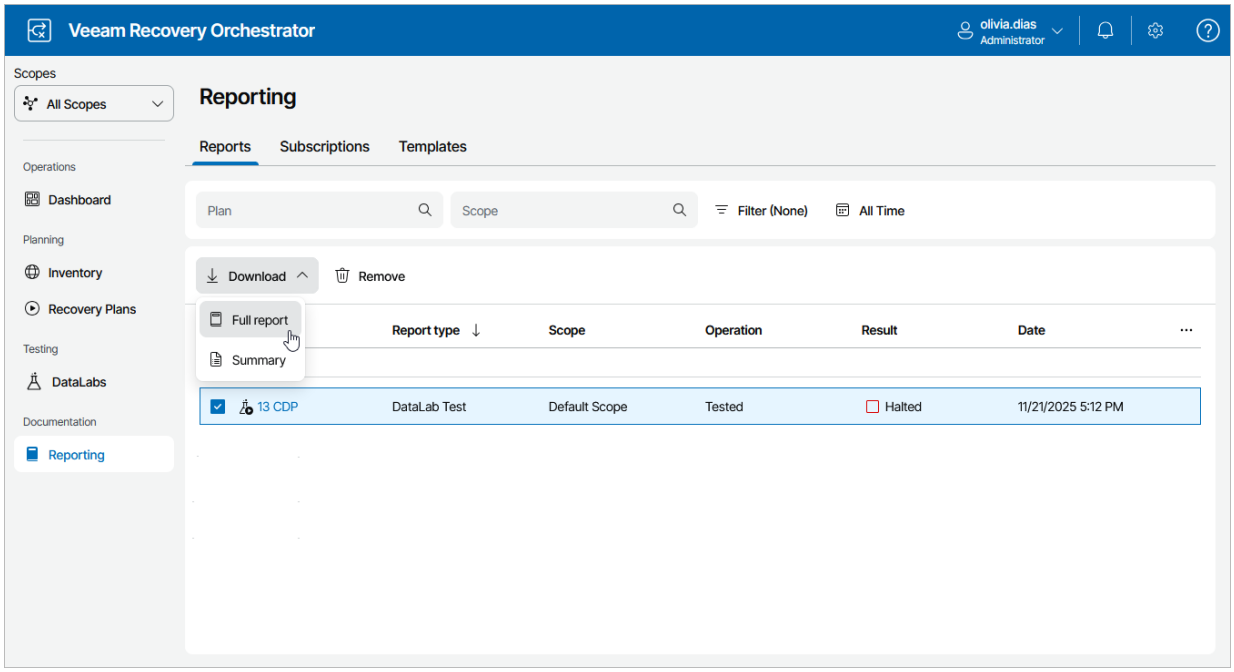

The DataLab Test Report will use the default report template or a custom template. The results of DataLab testing will be appended at the end of the template. The report will contain both the results of starting the DataLab and lab groups, and of testing the plan.

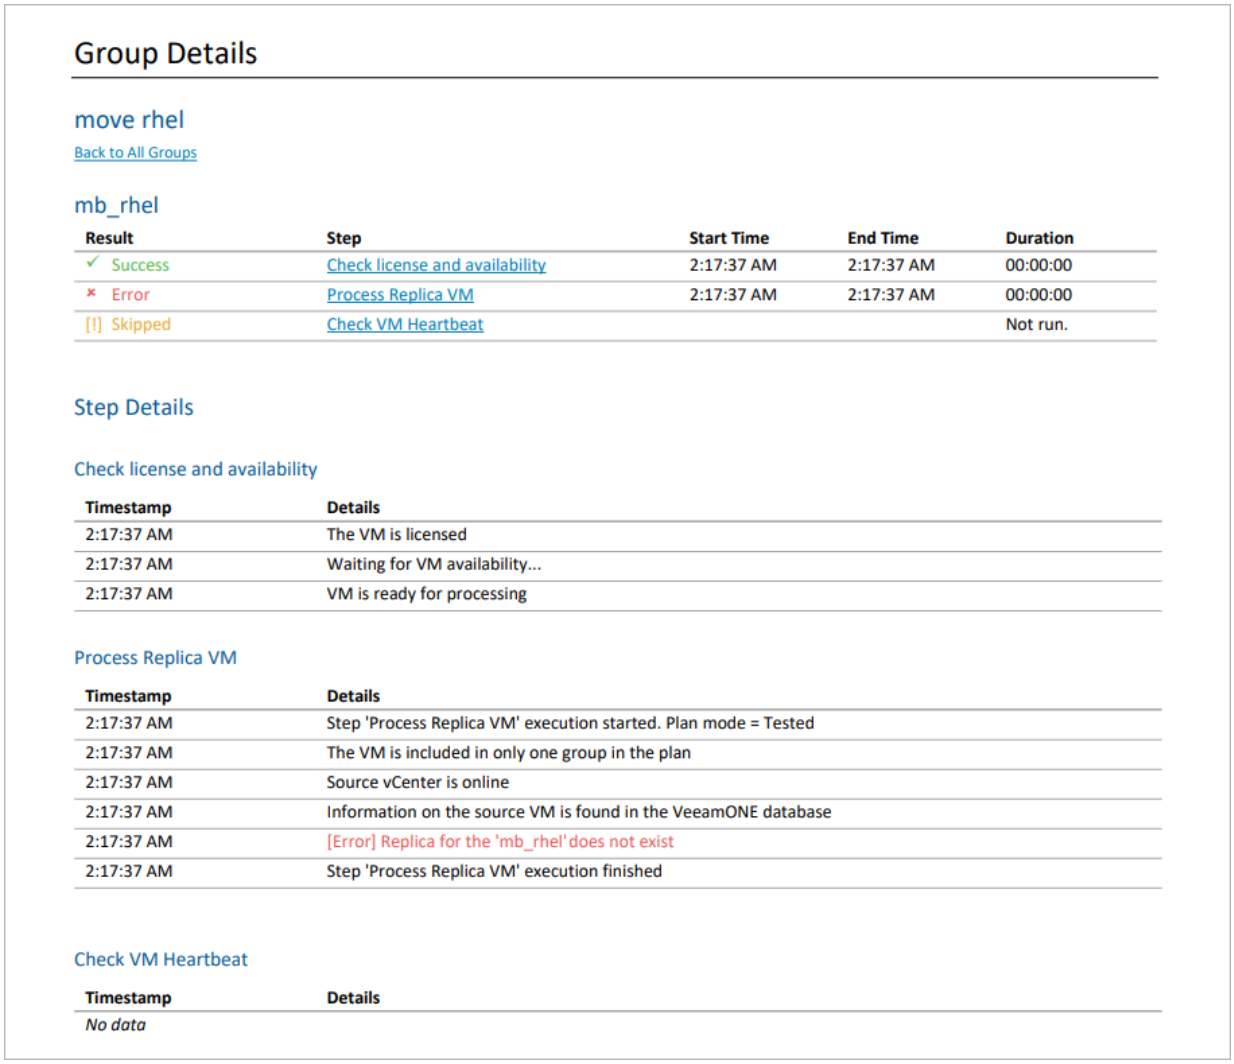

Group Details

move rhel

Back to All Groups

mb_rhel

Result	Step	Start Time	End Time	Duration
✓ Success	Check license and availability	2:17:37 AM	2:17:37 AM	00:00:00
× Error	Process Replica VM	2:17:37 AM	2:17:37 AM	00:00:00
[!] Skipped	Check VM Heartbeat			Not run.

Step Details

Check license and availability

Timestamp	Details
2:17:37 AM	The VM is licensed
2:17:37 AM	Waiting for VM availability...
2:17:37 AM	VM is ready for processing

Process Replica VM

Timestamp	Details
2:17:37 AM	Step 'Process Replica VM' execution started. Plan mode = Tested
2:17:37 AM	The VM is included in only one group in the plan
2:17:37 AM	Source vCenter is online
2:17:37 AM	Information on the source VM is found in the VeeamONE database
2:17:37 AM	[Error] Replica for the 'mb_rhel' does not exist
2:17:37 AM	Step 'Process Replica VM' execution finished

Check VM Heartbeat

Timestamp	Details
No data	

Viewing Plan Execution History

For each executed recovery plan (that is, upon transition from one stable state to another), Orchestrator will generate the **Plan Execution Report**. The report contains plan performance details and provides information on each processed machine and any errors encountered during plan execution.

Orchestrator generates two types of reports:

- A summary report that includes a plan overview and a summary of inventory groups included in the plan with drill-down hyperlinks to specific machines and color-coded results of processing every plan step.
- A full report that also includes information on specific steps that will run during the recovery process. For every group, machine and step included in the plan, the processing start time and duration will be recorded.

> **TIP**
>
> Summary information on plan execution results over all Orchestrator scopes will be also available on the Home Page Dashboard.

Downloading Plan Execution History

To access the report for a recovery plan:

1. Navigate to **Reporting**.
2. Select the report.
3. Click the plan name to download a summary report.

   -OR-

   Click **Download** and choose whether you want to download a summary or full report.

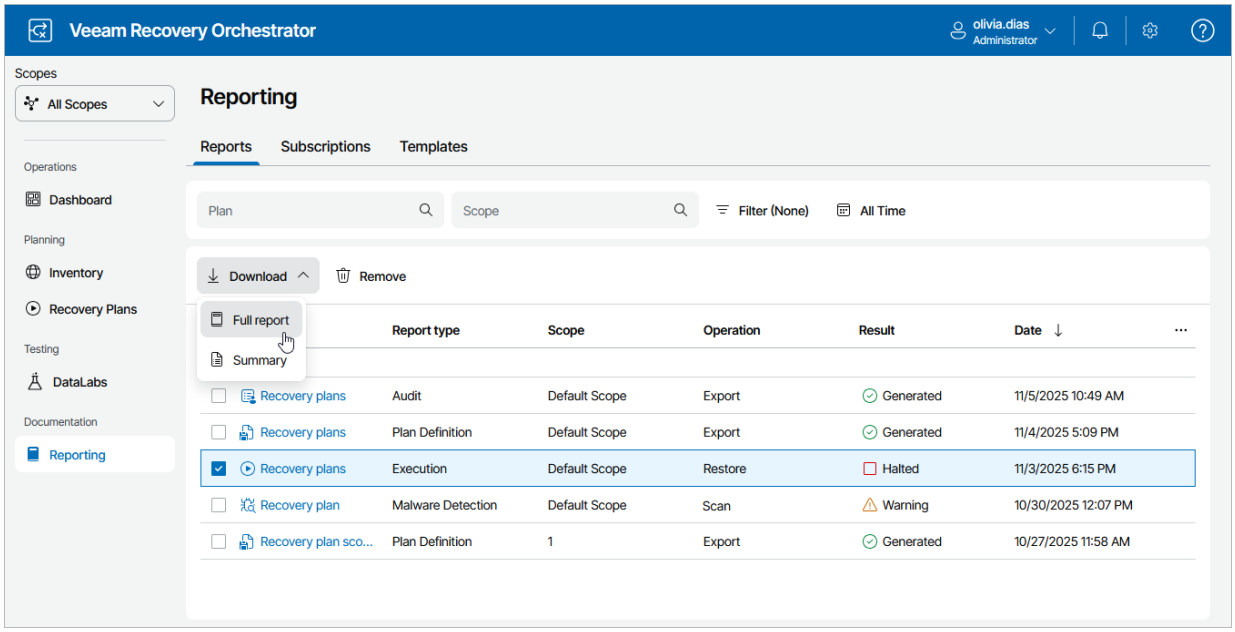

The Plan Execution Report will use the default report template or a custom template. The results of plan execution will be appended at the end of the template.

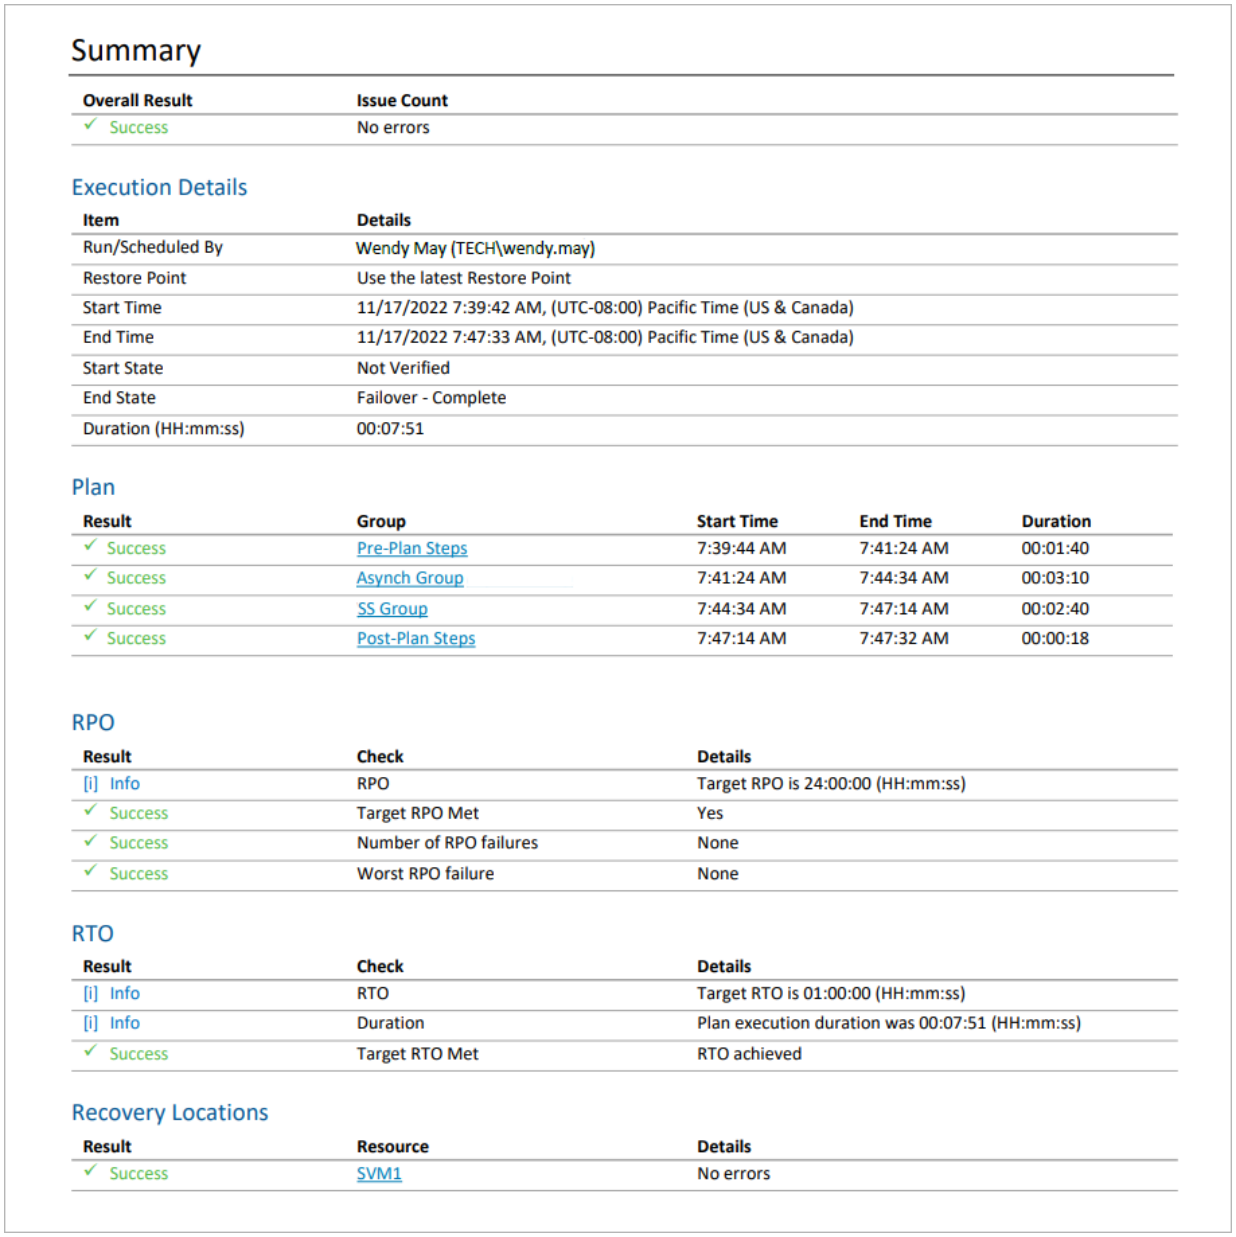

Summary

Overall Result	Issue Count
✓ Success	No errors

Execution Details

Item	Details
Run/Scheduled By	Wendy May (TECH\wendy.may)
Restore Point	Use the latest Restore Point
Start Time	11/17/2022 7:39:42 AM, (UTC-08:00) Pacific Time (US & Canada)
End Time	11/17/2022 7:47:33 AM, (UTC-08:00) Pacific Time (US & Canada)
Start State	Not Verified
End State	Failover - Complete
Duration (HH:mm:ss)	00:07:51

Plan

Result	Group	Start Time	End Time	Duration
✓ Success	Pre-Plan Steps	7:39:44 AM	7:41:24 AM	00:01:40
✓ Success	Asynch Group	7:41:24 AM	7:44:34 AM	00:03:10
✓ Success	SS Group	7:44:34 AM	7:47:14 AM	00:02:40
✓ Success	Post-Plan Steps	7:47:14 AM	7:47:32 AM	00:00:18

RPO

Result	Check	Details
[i] Info	RPO	Target RPO is 24:00:00 (HH:mm:ss)
✓ Success	Target RPO Met	Yes
✓ Success	Number of RPO failures	None
✓ Success	Worst RPO failure	None

RTO

Result	Check	Details
[i] Info	RTO	Target RTO is 01:00:00 (HH:mm:ss)
[i] Info	Duration	Plan execution duration was 00:07:51 (HH:mm:ss)
✓ Success	Target RTO Met	RTO achieved

Recovery Locations

Result	Resource	Details
✓ Success	SVM1	No errors

Viewing Audit Report

The **Plan Audit Report** includes data on all the activity performed within the specified period. The report contains details on plan check and scan operations, provides information on plan RTO and RPO, and lists any errors encountered during plan execution.

Orchestrator generates two types of reports:

- A summary report that includes a plan overview and a summary on all the performed scan and checks.
- A full report that also includes information on each check.

Generating Plan Audit Report

By default, Orchestrator runs the Plan Audit Report automatically for every *ENABLED* recovery plan once a month. You can also generate the report for a plan on demand:

1. Navigate to **Recovery Plans**.
2. Select the plan.
3. Click **Publish** > **Plan Audit Report**.

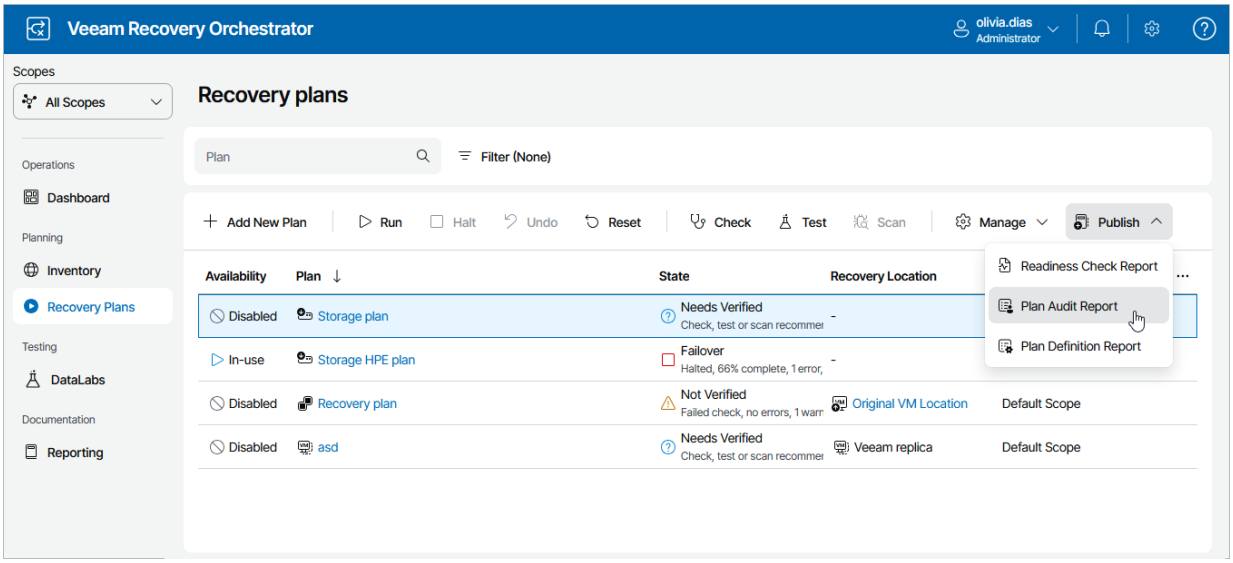

Downloading Plan Execution History

To access the report for a recovery plan:

1. Navigate to **Reporting**.
2. Select the report.
3. Click the plan name to download a summary report.

   -OR-

Click **Download** and choose whether you want to download a summary or full report.

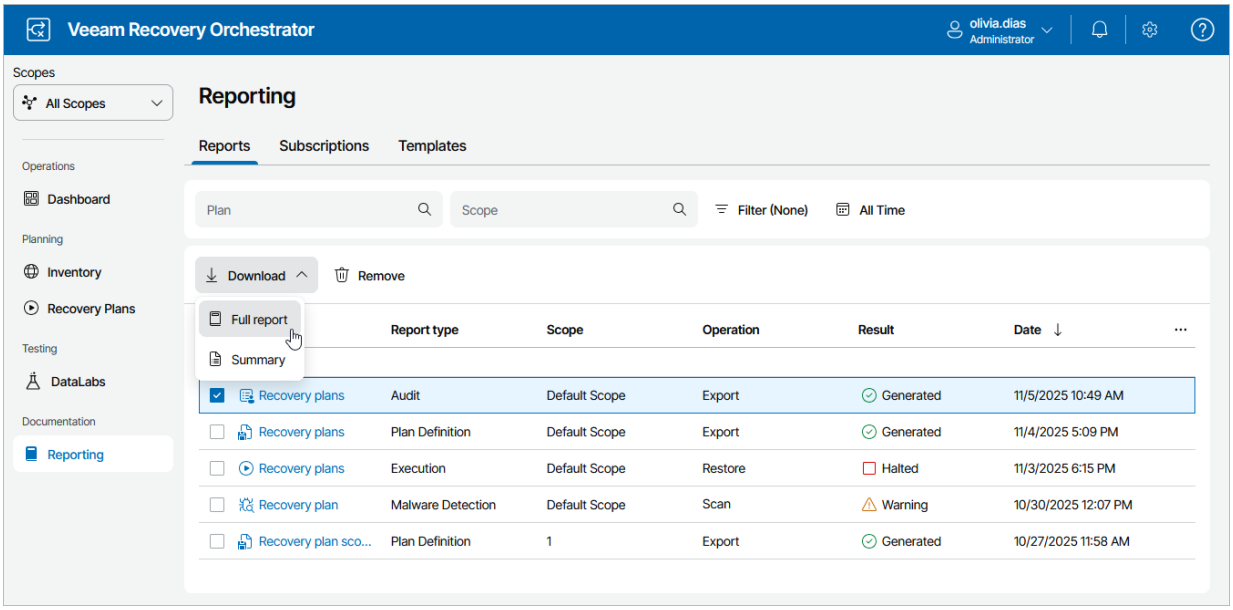

The Plan Audit Report will use the default report template or a custom template. The results of plan audit will be appended at the end of the template.

Summary

Key:
Halted
Malware detected
Error
Warning
Worst RPO/RTO failure
Success
Each symbol represents the worst result in that day
Subscript n shows the number of reports of that type in that day

Date		Malware Scans	Readiness Checks	Executions	Tests	RPO	RTO
12/16/2025	⦸	-	-	-	-	-	-
12/15/2025	⦸	-	-	-	-	-	-
12/14/2025	⦸	-	-	-	-	-	-
12/13/2025	⦸	-	-	-	-	-	-
12/12/2025	⦸	-	-	-	-	✓	-
12/11/2025	⦸	-	-	-	-	-	-
12/10/2025	⦸	-	-	-	-	-	-
12/9/2025	⦸	-	-	-	-	-	-
12/8/2025	⦸	-	-	-	-	-	-
12/7/2025	⦸	-	-	-	-	-	-
12/6/2025	⦸	-	-	-	-	✓	-
12/5/2025	⦸	-	-	-	-	✓	✓
12/4/2025	⦸	-	-	-	-	-	-
12/3/2025	⦸	-	-	-	-	-	-
12/2/2025	⦸	-	-	-	-	-	-
12/1/2025	⦸	-	-	⚠ 9	-	✗	-
11/30/2025	⦸	-	-	-	-	-	✓
11/29/2025	⦸	-	-	-	-	-	-
11/28/2025	⦸	-	-	-	-	-	-
11/27/2025	⦸	-	-	-	-	-	-
11/26/2025	⦸	-	-	⚠ 9	-	✗	✓
11/25/2025	⦸	-	-	-	-	-	-
11/24/2025	⦸	-	-	-	-	-	-
11/23/2025	⦸	-	-	-	-	-	-
11/22/2025	⦸	-	-	-	-	✓	✓
11/21/2025	⦸	-	-	-	-	-	-
11/20/2025	⦸	-	-	-	-	-	-
11/19/2025	⦸	-	-	-	-	-	-
11/18/2025	⦸	-	-	-	-	-	-
11/17/2025	⦸	-	-	-	-	-	-
11/16/2025		-	-	-	-	-	-

⦸= The plan was disabled during that day. Changes for disabled plans are collected, however readiness checks, malware scans and tests are not recorded

Generating Malware Scan Report

After you check a plan for possible malware, Orchestrator will generate the **Malware Scan Report**. The report contains:

- Details on the check for malware flags.
- Results on the scan of restore points with antivirus software.
- Details on the performed YARA scan.

Orchestrator generates two types of reports:

- A summary report that includes a plan overview, information on all VM groups included in the plan and a summary on all the performed malware checks.
- A full report that also includes information on all VMs included in the plan.

Generating Malware Scan Report

To generate the report for a recovery plan:

1. Navigate to **Recovery Plans**.
2. Select the plan and click **Scan**.

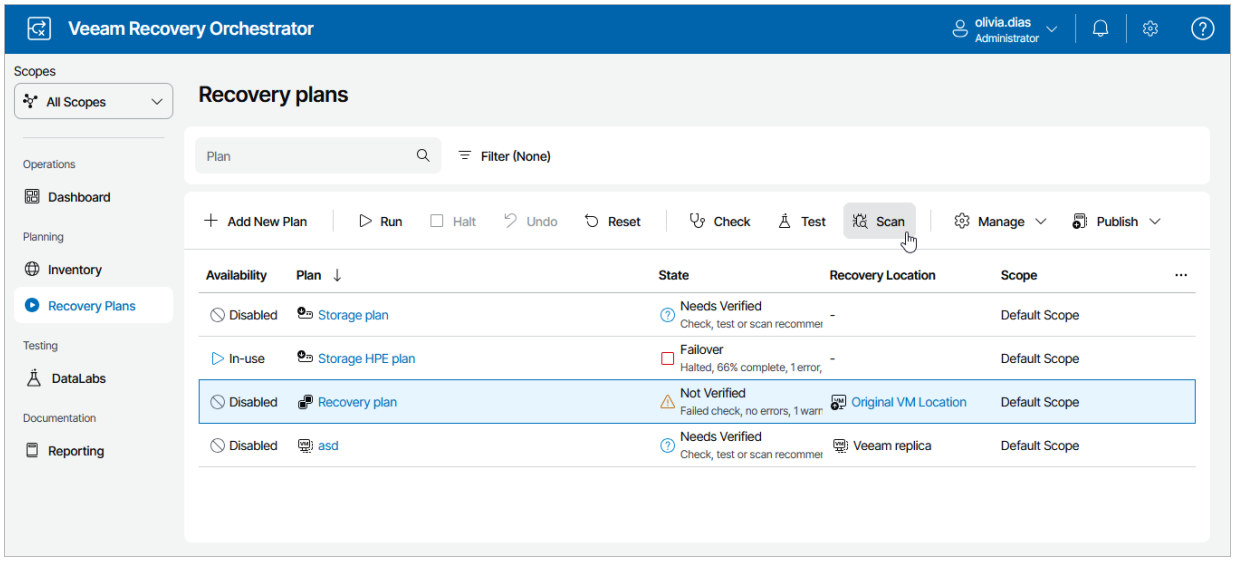

Downloading Malware Scan Report

To access the report for a recovery plan:

1. Navigate to **Reporting**.
2. Select the report.
3. Click the plan name to download a summary report.

   -OR-

Click **Download** and choose whether you want to download a summary or full report.

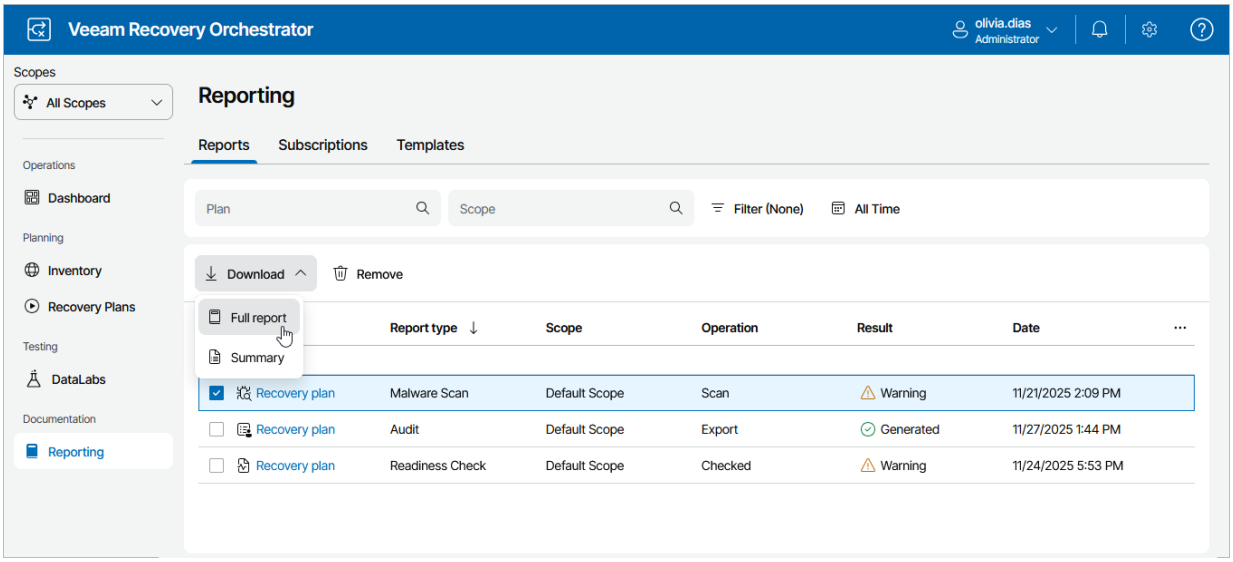

The Malware Scan Report will use the default report template or a custom template. The results of malware scan will be appended at the end of the template.

Reviewing Dashboards

Orchestrator comes with 2 dashboards that allow you to:

- Track the health state of the connected infrastructure. For more information, see Administration Dashboard.
- Analyze the readiness for disaster recovery operations across different scopes. For more information, see Home Page Dashboard.

Administration Dashboard

The dashboard on the **Administration** page of the Orchestrator UI provides at-a-glance real-time overview of your infrastructure:

- Shows the state of the connected Veeam Backup & Replication servers, vCenter Servers, Microsoft Hyper-V servers, storage systems and Microsoft Azure compute accounts.
- Displays information on license usage across the whole infrastructure.
- Shows the replication status of datastores included in storage plans.
- Displays all configured recovery locations.

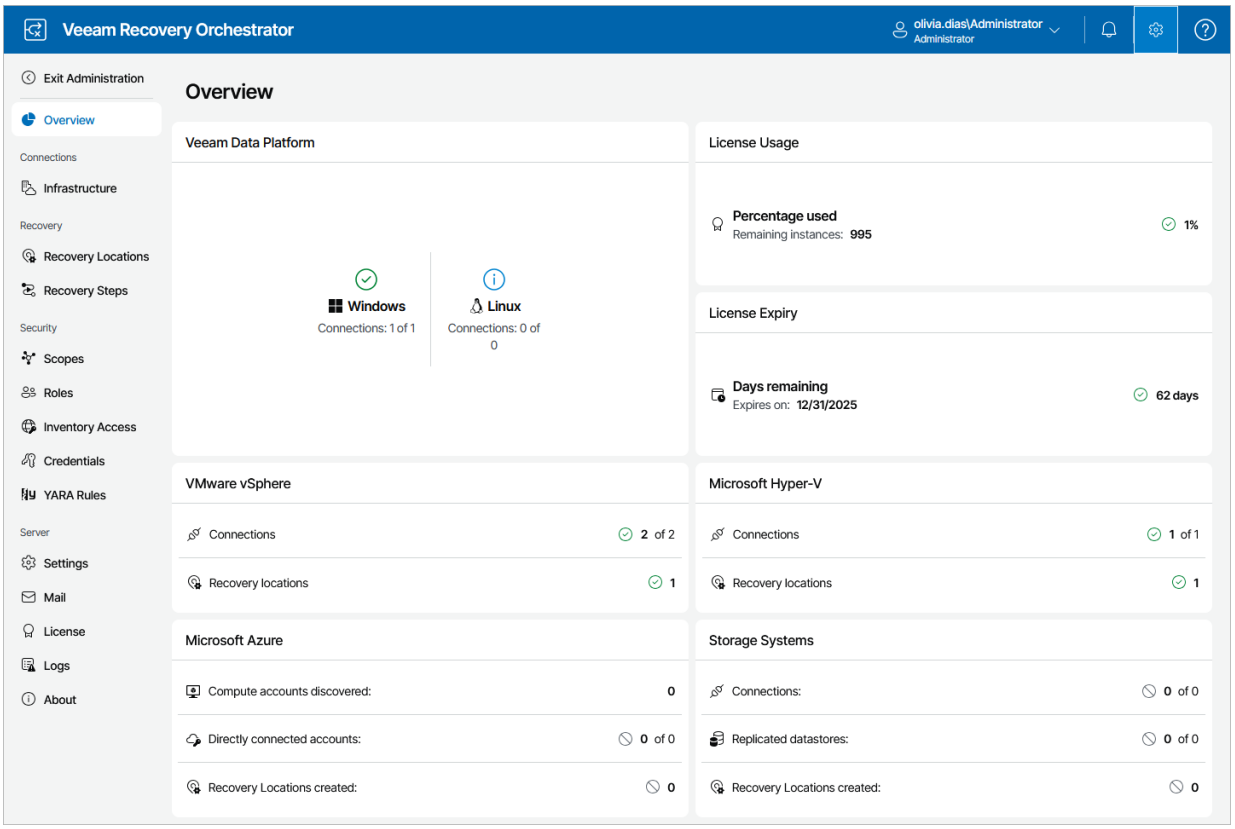

Home Page Dashboard

The dashboard on the home page of the Orchestrator UI provides an overview of all recovery plans for the selected scope:

- The **Plan Verification** chart shows the number of:
  - Failed checks
  - Checks completed successfully
  - Plans not checked yet

  The worst state of a plan readiness check is *Not verified*. It means that the plan is not in the ready-to-run state.
- The **Plan Execution** chart shows the number of:
  - Halted plans
  - Plans completed successfully
  - Plans completed with warnings
  - Plans completed with errors
  - Plans completed with malware issues

  The worst state of a plan execution is *Halted*. It means that the plan has stopped processing because of a critical error for a machine from a critical inventory group in the plan.
- The **Plan Testing** chart shows the number of:
  - Failed and halted plan tests
  - Plan tests completed successfully
  - Plan tests completed with warnings
  - Plan tests completed with errors
  - Plan tests completed with malware issues

  The worst state of a plan testing is *Failed*. It means that the test has stopped because of a critical error for a machine from a critical inventory group in the lab.
- The **Plan Scanning** chart shows the number of:
  - Halted plan scans
  - Plan scans completed successfully
  - Plan scans completed with warnings
  - Plan scans completed with errors
  - Plan scans completed with malware issues

  The worst state of a plan scanning is *Malware issue*. It means that the scan has stopped because an infected restore point was detected on a machine included the plan.

- The **Recovery Point Objectives (RPO)** and **Recovery Time Objectives (RTO)** panes show recovery plans with the RPO and RTO, allowing you to track the achieved objectives versus targets for 10 plans to ensure you are meeting business service level agreements (SLAs).

To switch to the **Plan Details** page and see the list of issues that occurred while performing plan steps for a plan, click the plan name in the **Plan** column.

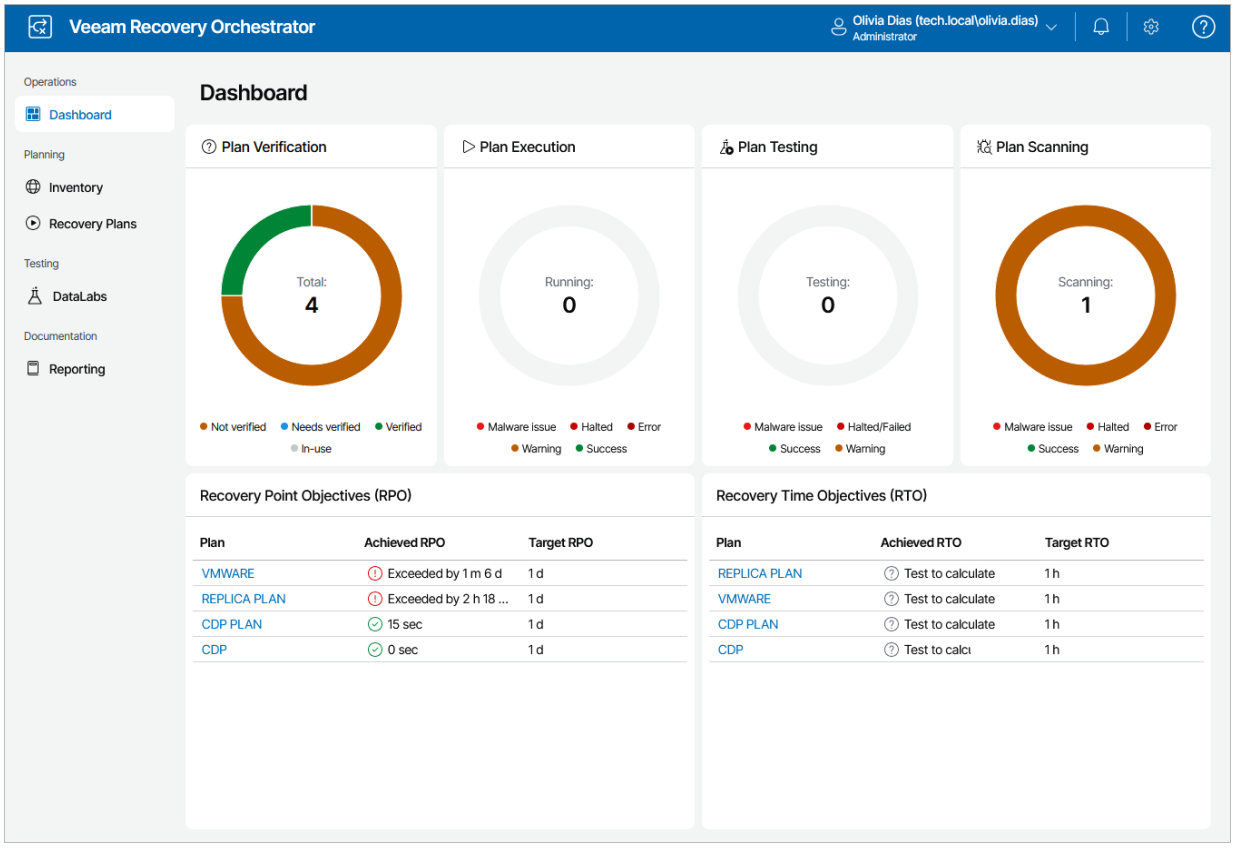

Managing Custom Scripts to Veeam Recovery Orchestrator

If you have a PowerShell script that you want to run as part of the recovery process, you can upload your script into Orchestrator, and it will be executed when you run your plan.

The script can run on a Veeam Backup & Replication server, on the Orchestrator server or inside each machine included in the plan. You can customize settings required for script execution and pass various parameters into the script: credentials, runtime variables (such as *vm_name* or *plan_state*) and any other custom parameters you require. Script output will be captured in plan details in the Orchestrator UI, and in Plan Execution and DataLab Test reports.

> **NOTE**
>
> Due to Microsoft Azure limitations, script output for cloud plans is limited to 4 Kb.

Requirements

If you want to create a custom script and execute it when running a recovery plan, you must take into account the following considerations.

- Running custom scripts inside guest OSes of Linux-based machines is not supported.
- The script you want to use must be a PowerShell script. Orchestrator 13 supports PowerShell scripts only.
- To allow the script to run inside a machine guest OS, it is required that you have Microsoft PowerShell 3.0 and .Net Framework 4.0 installed on each machine for which you enable the custom script step.
- To allow the script to run inside the guest OS of a machine recovered to Microsoft Azure, the client secret or certificate must be specified for the Microsoft Azure account added to the Veeam Backup & Replication protecting this machine. For more information, see Connecting Microsoft Azure Servers.
- To allow the script to run inside the guest OS of a machine recovered to Microsoft Azure, the compute account (Azure AD application) added to the Veeam Backup & Replication server configuration and used to access the Microsoft Azure resources must be assigned the *Contributor* and *Key Vault Crypto User* user roles. Alternatively, you can create a custom role on the Microsoft Azure portal with the granular permissions listed in the Veeam Backup & Replication User Guide, section Creating Custom Role for Azure and Azure Stack Hub Accounts. Note that this role must also have the following permissions assigned: `Microsoft.Compute/virtualMachines/runCommands/read`, `Microsoft.Compute/virtualMachines/runCommands/write`, `Microsoft.Compute/virtualMachines/runCommands/delete`.
- To allow the script to run on a Veeam Backup & Replication server, no additional software is required.

Adding Custom Scripts

To upload an existing script into Orchestrator as a separate plan step, perform the following steps:

1. Switch to the **Administration** page.
2. Navigate to **Recovery Steps** and click **Add**.
3. In the **Add Custom Script Step** window, click the link in the **Script details** field.
4. In the **Step Configuration** window, click the link in the **Script file** field and browse to the script file in the **Add** window. You can also provide a name and description for future reference. The maximum length of the step name is 128 characters; the following characters are not supported: * : / \ ? " < > | .
5. Click **Apply**.

This section will demonstrate how to upload a simple example script into Orchestrator.

```
Param(
   [Parameter(Mandatory=$true)]
   [string]$folderName
)
try {
   $fileName = "HelloWorld.txt"
   "Hello World!" | Out-File -FilePath "$folderName\$fileName"
   Write-Host "File $fileName was created in folder $folderName"
}
catch {
   Write-Error "Failed to create file in folder $folderName"
   Write-Error $_.Exception.Message
}
```

Additionally, you can configure parameters for step execution and add any other custom parameters that your script requires.

Configuring Step Parameters

In the **Step parameters** section of the **Add Custom Script** window, do the following:

1. Choose whether you want the step to be critical for machine recovery.
2. Select the **Run step during a DataLab test** check box if you want the step to be executed during plan testing in a DataLab.
3. Select the **Run step during Undo and Failback** check box if you want the step to be executed during the Failback and Undo Failover operations.
4. In the **Timeout** field, specify the maximum amount of time (in seconds) for the step to execute.
5. In the **Retries** filed, specify the number of retries that will be attempted if the step fails on the first try.
6. From the **Script Execution Location** drop-down list, choose whether you want the step to be executed on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.

> **IMPORTANT**
>
> To allow the script to run inside the guest OS of a processed machine, it is required that you have Microsoft PowerShell 3.0 and .Net Framework 4.0 installed on each machine for which you enable this step.

7. [This step applies only if the **Script Execution Location** parameter value is set to *In-Guest OS*] In the **Windows Credentials** field, click the link and choose credentials that will be used to gain access to the in-guest OS.

For more information on script parameters, see Configuring Common Parameters.

Configuring Script Parameters

By default, Orchestrator automatically detects all mandatory parameters of uploaded scripts — but only in case these parameters are of the *Credentials*, *Boolean*, *Text* and *Integer* types. You can also add any other custom parameters that your script requires. To do that, click **Add** in the **Script parameters** section and do the following the **Add Script Parameter** window:

1. Select a type of the parameter that you want to add. In our example, the parameter *folderName* is required — that is why the parameter type will be *Text*.
2. Use the **Parameter name** and **Parameter description** fields to enter a name for the parameter and to provide a description for future reference. In our example, the parameter name will be *folderName*.
3. Enter a default value that you want to assign to the parameter. You can leave this field empty for the value to be set when the step is added to a plan.
4. Click **Apply**.

You can also pass runtime variables into the script. For more information, see Using Runtime Parameter Variables.

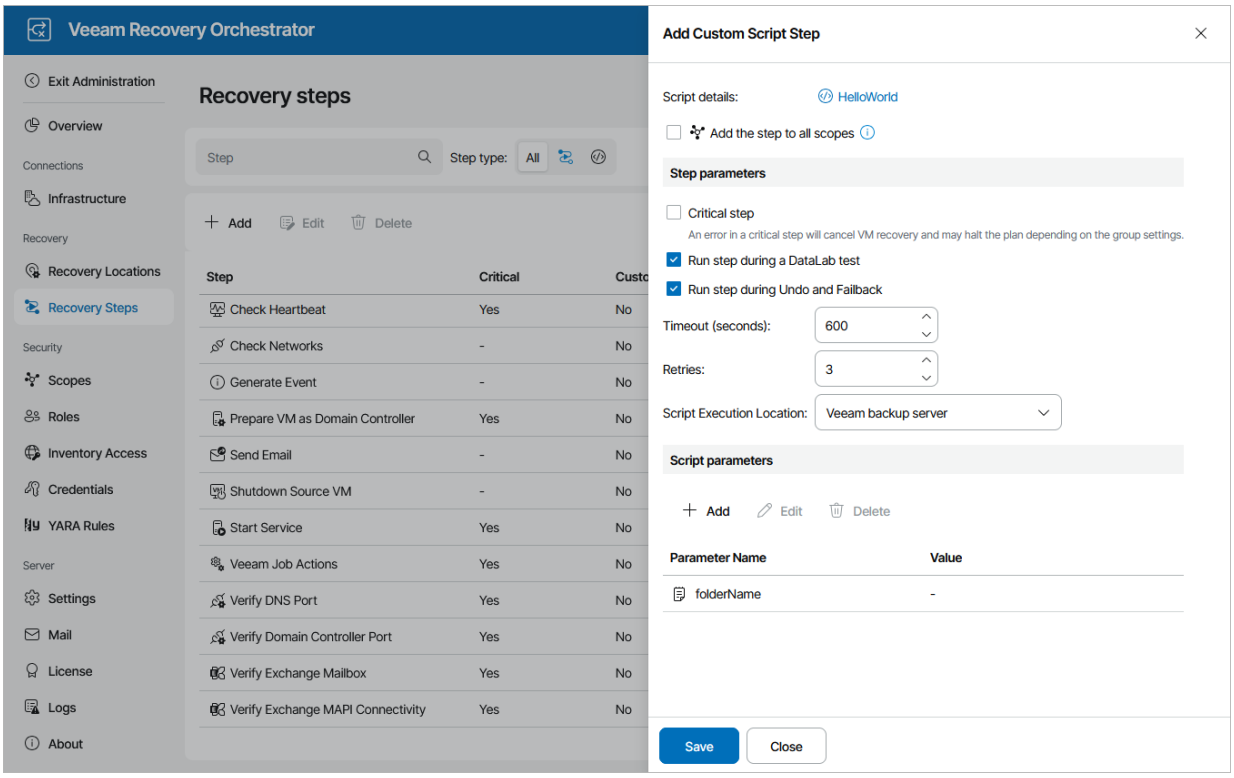

Configuring Common Parameters

When you create a custom script step, you can set a list of default parameters for script execution. Orchestrator already includes a number of out-of-the-box common default parameters that you can modify as described in section Adding Custom Scripts.

> **IMPORTANT**
>
> To allow the script to run inside the guest OS of a processed machine, it is required that you have Microsoft PowerShell 3.0 and .Net Framework 4.0 installed on each machine for which you enable this step.

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	No
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	Yes
Timeout	Maximum amount of time (in seconds) for the step to execute.	600
Retries	Number of retries that will be attempted if the step fails on the first try.	3
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS. **Note**: This parameter cannot be changed when restoring VMs to a Microsoft Hyper-V environment.	Veeam backup server
Windows Credentials*	Credentials required to gain access to the in-guest OS. **Note**: Applies only if the **Script Execution Location** parameter value is set to *In-Guest OS*.	—

*This parameter is not required for custom scripts running inside guest OSes of machines included in cloud plans because these scripts are executed under the *NT AUTHORITY\SYSTEM* account.

To specify credentials that the script will use to run within the guest OS of a processed machine, follow the instructions provided in section Configuring Windows Credentials Parameter.

Configuring Windows Credentials Parameter

If you want to provide credentials that the script will use to run within the guest OS of a machine included in the plan, do the following:

1. Add credentials that the script will use to connect to the machine when performing recovery as described in section Managing Credentials.
2. Navigate to **Recovery Steps**.
3. In the **Steps** column, select the step and click **Edit**.
4. In the **Step Editor** window, do the following
   a. In **Script Execution Location** field, select *In-Guest OS*.
   b. Click the link in the **Windows credentials** field, select the necessary credentials in the **Choose Credentials** window and click **Apply**.
   c. To save changes made to the script, click **Save**.

> **NOTE**
>
> If you do not specify any credentials, the script will fail to run, and the Readiness Check test will report that the **Windows Credentials** parameter settings are not configured.

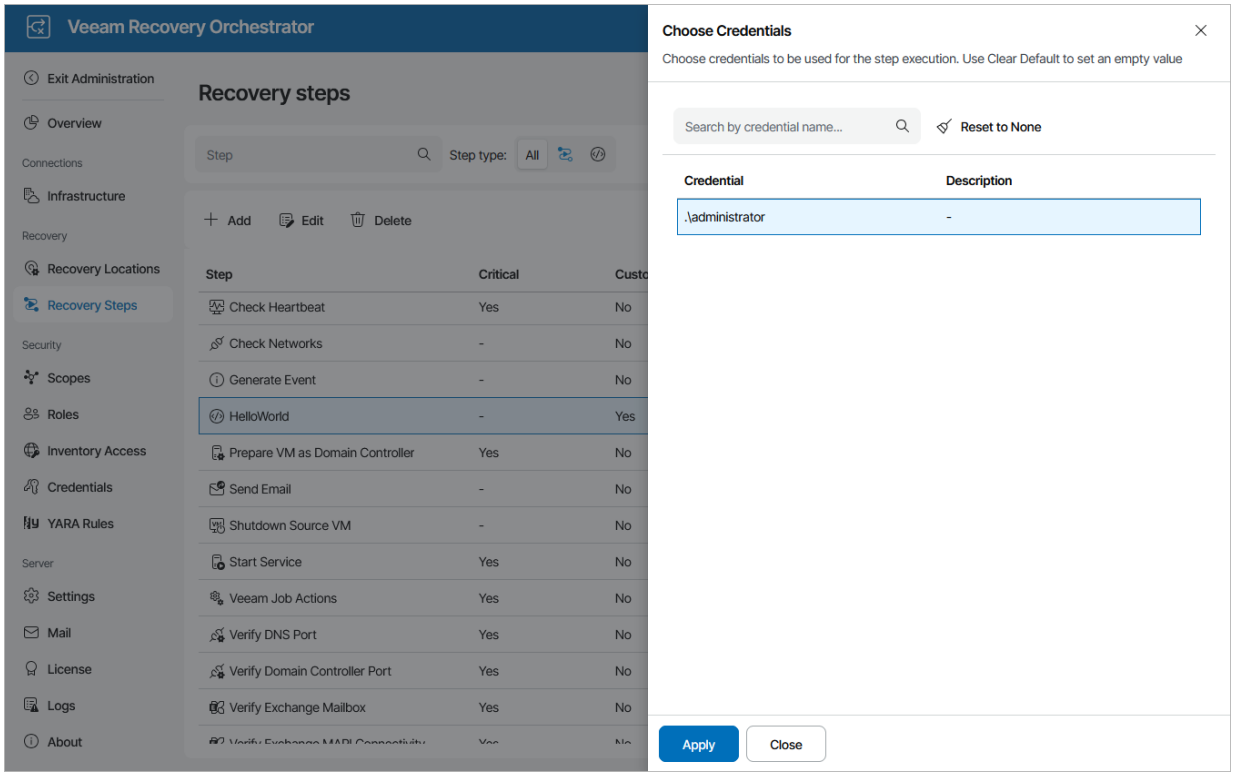

Adding Credentials Parameter to Your Script

You may add multiple custom parameters of the *Credentials* type. To pass the credentials to the script, Orchestrator uses the following parameter name convention.

In the script file, a parameter with the *Credentials* type will split into 2 parameters: the first one will contain a user name and the second one will contain a password. For example, if you add a credential parameter named *SQLCreds*, Orchestrator will pass it to the script as *$SQLCredsUsername* and *$SQLCredsPassword*.

```
Param(
 [string]$SQLCredsUsername,
 [string]$SQLCredsPassword
)
```

IMPORTANT

The statement must contain a comma-separated list of variables prefixed with a data type. Default values are optional.

Using Runtime Parameter Variables

Orchestrator allows you to pass runtime variables into the script.

In our example, the *folderName* custom parameter has been added recently, and it is required to specify a default value for it. To set a custom variable as the default value, do the following:

1. Switch to the **Administration** page.
2. Navigate to **Recovery Steps**.
3. In the **Steps** column, select the script step and click **Edit**.
4. In the **Step Editor** window, click **Add** in the **Script parameters** section.
5. In the **Add Script Parameter** window, do the following:
   a. Use the **Parameter name** field to enter a name for the new parameter. The maximum length of the location name is 128 characters; the following characters are not supported: * : / \ ? " < > | .
   b. From the **Parameter type** drop-down list, select *Text / Macro*.
   c. In the **Default value** field, click **Choose** and do the following in the **Set Default Text / Macro** window:
      i. Set the **Show macros** toggle to *On*
      ii. From the list of available variables, double-click the value that you want to assign to the parameter, and click **Apply**.
   d. Click **Apply**.
6. To save changes made to the script, click **Save**.

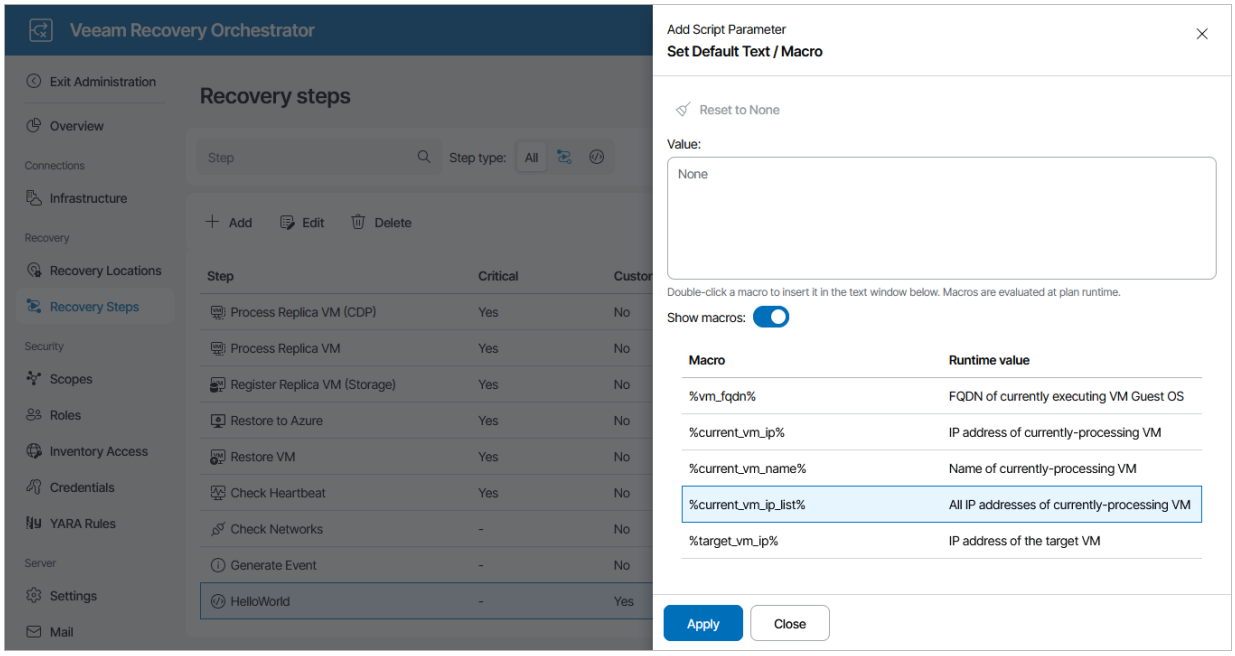

For more information on parameter variables that you can pass into a script, see Appendix A. Recovery Plan Steps.

Adding Custom Script Step to Plan

For each machine included in a recovery plan, you can add a custom script step to be performed when processing the machine:

1. Navigate to **Recovery Plans**.
2. Select the plan to which you want to add the custom step and click **Manage** > **Edit**.
3. On the **Edit Plan** page, in the **Recovery plan** column, expand the plan to see all its inventory groups. Then, select the necessary inventory group and click **Step editor**.
4. In the **Step Editor** window, click **Add**, select the custom step that you want to add to the plan and click **Save**.

> **TIP**
>
> You can also add a custom step to a specific machine as described in section Configuring Steps.

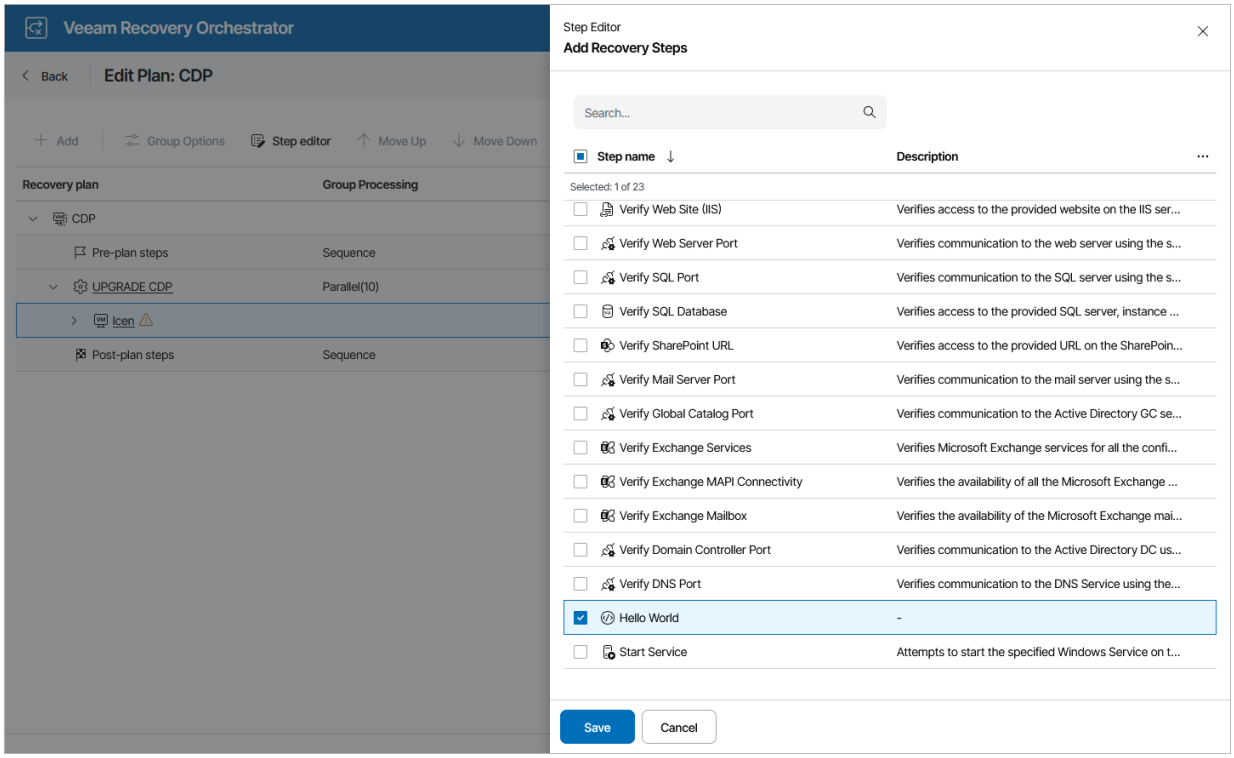

After you add the custom step for the machine in the plan, check step parameter settings and modify them if required. For more information, see Configuring Step Parameters.

Capturing Script Errors and Warnings

To log any custom script information into step details and execution reports, make sure to use the Write-Host cmdlet in the script. To indicate errors and warnings occurred during script execution and to pass this data to Orchestrator, make sure to use the Write-Error and Write-Warning cmdlets in the script. This will post script output in the Orchestrator UI, and also in Plan Execution and DataLab Test reports.

If no errors and warnings occur during script execution, Orchestrator will report the *Success* execution state. Otherwise, in case a number of warnings and errors occurs, Orchestrator will report the worst state.

> **NOTE**
>
> For a script that runs inside the guest OS of a machine included in a cloud plan, the `Write-Warning` cmdlet is not supported.

Appendices

See in this section:

Appendix A. Recovery Plan Steps

For every machine included in a recovery plan, there are steps to be performed in sequence. This section provides information on these steps.

Recovery plan steps require various parameters that are passed between steps, are available when editing and adding a step, and are used in messages. Steps can accept results from previous steps as input, and output to subsequent steps. Step execution results are displayed on the **Plan Details** page.

To learn how to configure step parameters, see Editing Recovery Plans.

Steps Available

During failover to the DR site and restore to a recovery location, the Orchestrator server performs the list of steps described in this section.

Plan Step	Restore Plan		Replica Plan	CDP Replica Plan	Storage Plan	Cloud Plan
	To VMware	To Hyper-V				
Check VM Heartbeat	●	●	●	●	●	○
Restore to Azure	○	○	○	○	○	●
Generate Event	●	●	●	●	●	●
Check Networks	●	○	●	●	●	○
Prepare VM as Domain Controller*	●	○	●	●	●	○
Process Replica VM (CDP)	○	○	○	●	○	○
Process Replica VM	○	○	●	○	○	○
Register Replica VM (Storage)	○	○	○	○	●	○
Restore VM	●	●	○	○	○	○
Send Email	●	◐	●	●	●	●
Shutdown Source VM	●	◐	●	●	●	●
Start Service	●	◐	●	●	●	○
Veeam Job Actions	●	●	●	●	●	●

Plan Step	Restore Plan		Replica Plan	CDP Replica Plan	Storage Plan	Cloud Plan
	To VMware	To Hyper-V				
Verify DNS Port	●	○	●	●	●	○
Verify Domain Controller Port	●	○	●	●	●	○
Verify Exchange Mailbox	●	○	●	●	●	○
Verify Exchange MAPI Connectivity	●	○	●	●	●	○
Verify Exchange Services	●	○	●	●	●	○
Verify Global Catalog Port	●	○	●	●	●	○
Verify Mail Server Port	●	○	●	●	●	○
Verify SharePoint URL	●	○	●	●	●	○
Verify SQL Database	●	○	●	●	●	○
Verify SQL Port	●	○	●	●	●	○
Verify Web Server Port	●	○	●	●	●	○
Verify Web Site (IIS)	●	○	●	●	●	○
VM Power Actions	●	●	●	●	●	●
Custom Script	●	●	●	●	●	●

Plan Step	Restore Plan		Replica Plan	CDP Replica Plan	Storage Plan	Cloud Plan
	To VMware	To Hyper-V				
Protect VM Group						

*This step is available for DataLab testing only.

Check VM Heartbeat

This step checks heartbeat of the recovered VM using VMware Tools. If VMware Tools are not installed, remove this step from the plan.

You can override the following parameters for the **Check VM Heartbeat** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	Yes
Heartbeat Count	Number of heartbeats to execute.	4
Heartbeat Poll Time	Amount of time (in seconds) to wait between heartbeats.	30

Restore to Azure

This step restores the selected machine from a vSphere or Veeam agent backup to the Microsoft Azure recovery location specified for the plan.

You can override the following parameters for the **Restore to Azure** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Restored VM Name	Name under which the machine will be recovered and registered.	*%source_machine_name%*
Public IP	Defines whether a public IP address will be assigned to the recovered VM.	No
VM Configuration	Defines the VM configuration to be used to recover machines.	Configuration 1
Restore Timeout (minutes)	Maximum amount of time (in minutes) for the step to execute. **Note**: The default parameter value is set to *0*, which means that Orchestrator will wait for the step to complete for as long as required. However, if you run the Halt action to halt the restore process, Orchestrator will halt the plan immediately, despite the infinite timeout value.	0
Retries	Number of retries that will be attempted if the step fails on the first try.	2

Generate Event

This step generates events in the Windows event log on the Orchestrator server.

You can override the following parameters for the **Generate Event** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	No

Parameter	Description	Default Value
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab. **Note**: If you set the parameter value to *Execute*, keep in mind that all actions performed while testing the plan will not be reverted when the test is over.	No
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	Yes
Timeout	Maximum amount of time (in seconds) for the step to execute.	60
Retries	Number of retries that will be attempted if the step fails on the first try.	3
Event Text	Event description. **Note**: You can use the default text or define *%values%* that will populate during plan execution.	Plan *%plan_name%* is in state *%plan_state%*

Check Networks

This step pings the selected machine and VMNICs until a stable response is received or timeout exceeds.

You can override the following parameters for the **Check Networks** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	No
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Execute
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	Execute
Timeout	Maximum amount of time (in seconds) for the step to execute.	600

Parameter	Description	Default Value
Retries	Number of retries that will be attempted if the step fails on the first try.	0

Prepare VM as Domain Controller

This step prepares a domain controller running on the selected VM to perform an authoritative restore. This will always be the first step to execute for the VM.

For more information on performing an authoritative restore of a domain controller, see Microsoft Docs and this Veeam KB article.

> **IMPORTANT**
>
> Before you restore a domain controller, you must enable application-aware processing for the selected VM in Veeam Backup & Replication as described in the Veeam Backup & Replication User Guide, section Enable Application-Aware Processing.

The **Prepare VM as Domain Controller** step has the following parameters, but they are not editable:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for VM recovery. If you mark the step as *Critical*, its failure for a VM from a critical inventory group will halt the plan.	Yes

Process Replica VM (CDP)

This step performs a number of actions that depend on the plan state:

- **Failover** – you will fail over from the source VM to the VM replica using the selected restore point. The VM replica will be powered on.
- **Failback** – you will fail back from the VM replica to the source VM. Changes made to the VM replica will be synchronized with the source VM. The source VM will be powered on.

For more information on other actions that can be performed with by step, see the Veeam Backup & Replication User Guide, section Failover and Failback for CDP.

You can override the following parameters for the **Process Replica VM (CDP)** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for VM recovery. If you mark the step as *Critical*, its failure for a VM from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab. **Note**: The default parameter value is set to *Skip* since the current product version does not support testing of CDP replica plans.	Execute (not editable)
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	Execute (not editable)
Power On Source VM After Undo	Defines whether the source VM will be powered on during the Undo Failover operation.	Yes
Timeout	Timeout (in seconds) used for the Failover and Undo Failover operations.	1200
Failback Timeout	Timeout (in minutes) used for the Failback operation. **Note**: The default parameter value is set to *0*, which means that Orchestrator will wait for the step to complete for as long as required. However, if you run the Halt action to halt the failback process, Orchestrator will halt the plan immediately, despite the infinite timeout value.	0
Retries	Number of retries that will be attempted if the step fails on the first try.	2

Process Replica VM

This step performs a number of actions that depend on the plan state:

- **Failover** – you will fail over from the source VM to the VM replica using the selected restore point. The VM replica will be powered on.
- **Failback** – you will fail back from the VM replica to the source VM. Changes made to the VM replica will be synchronized with the source VM. The source VM will be powered on.

For more information on other actions that can be performed by this step, see the Veeam Backup & Replication User Guide, section Failover and Failback for Replication.

You can override the following parameters for the **Process Replica VM** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for VM recovery. If you mark the step as *Critical*, its failure for a VM from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	Yes (not editable)
Power On Source VM After Undo	Defines whether the source VM will be powered on during the Undo Failover operation.	Yes
Timeout	Timeout (in seconds) used for the Failover and Undo Failover operations.	1200
Failback Timeout	Timeout (in minutes) used for the Failback operation. **Note**: The default parameter value is set to *0*, which means that Orchestrator will wait for the step to complete for as long as required. However, if you run the Halt action to halt the failback process, Orchestrator will halt the plan immediately, despite the infinite timeout value.	0
Retries	Number of retries that will be attempted if the step fails on the first try.	2

Register Replica VM (Storage)

This step registers the selected VM on a host from a storage recovery location, changes the IP address configuration of the VM, applies the mapping specified for the location, and then powers the VM on.

> **NOTE**
>
> Before powering a recovered VM on, Orchestrator removes all unused swap (.VSWP) files from the default VM directory. If a custom location is used to store swap files, Orchestrator will not be able to remove them.

You can override the following parameters for the **Register Replica VM (Storage)** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for VM recovery. If you mark the step as *Critical*, its failure for a VM from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes (not editable)
Timeout	Maximum amount of time (in seconds) for the step to execute.	600
Retries	Number of retries that will be attempted if the step fails on the first try.	2

Restore VM

This step restores the selected machine to the recovery location specified for the plan.

You can override the following parameters for the **Restore VM** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Power On VM After Restore	Defines whether the VM will be powered on after the restore operation.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Restore Timeout (minutes)	Maximum amount of time (in minutes) for the step to execute. **Note**: The default parameter value is set to *0*, which means that Orchestrator will wait for the step to complete for as long as required. However, if you run the Halt action to halt the restore process, Orchestrator will halt the plan immediately, despite the infinite timeout value.	0

Parameter	Description	Default Value
Retries	Number of retries that will be attempted if the step fails on the first try.	2
Restored VM Name	Name under which the machine will be recovered and registered. **Note**: By default, the recovered VM has the name of the source machine. If you want to recover the machine to the same datacenter where the source machine is registered and still resides, it is recommended that you change the parameter value to avoid conflicts.	*%source_machine_name%*

Send Email

This step sends an email using the configured SMTP server and subscribed email addresses.

> **NOTE**
>
> For restore plans, this step is not available when restoring VMs from vSphere backups to a Microsoft Hyper-V environment.

You can override the following parameters for the **Send Email** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	No
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab. **Note**: If you set the parameter value to *Execute*, keep in mind that all actions performed while testing the plan will not be reverted when the test is over.	No
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	Yes
Timeout	Maximum amount of time (in seconds) for the step to execute.	600
Retries	Number of retries that will be attempted if the step fails on the first try.	0

Parameter	Description	Default Value
Recipients	Recipients of the email. **Note**: To add multiple addresses, use commas.	—
Subject	Subject of the email. **Note**: You can use the default text, or define *%values%* which will populate during plan execution.	Orchestrator Email notification for *%plan_name%*
Body	Body of the email. **Note**: You can use the default text, or define *%values%* which will populate during plan execution.	Plan *%plan_name%* is in state *%plan_state%*

Shutdown Source VM

This step shuts down the selected source VM. It does not affect the recovered VM.

> **NOTE**
>
> For restore plans, this step is not available when restoring VMs from Hyper-V backups to a Microsoft Hyper-V environment.

You can override the following parameters for the **Shutdown Source VM** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for VM recovery. If you mark the step as *Critical*, its failure for a VM from a critical inventory group will halt the plan.	No
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab. **Note**: If you set the parameter value to *Execute*, keep in mind that all actions performed while testing the plan will not be reverted when the test is over.	No
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No (not editable)
Timeout	Maximum amount of time (in seconds) for the step to execute.	300

Parameter	Description	Default Value
Retries	Number of retries that will be attempted if the step fails on the first try.	1
Shutdown Action	Defines whether the step will shut down the guest OS of the source VM. **Note**: The *Shutdown OS* option requires VMware Tools to be installed in the guest OS.	Shutdown OS

Start Service

This step starts the Windows Service on the processed machine.

> **NOTE**
>
> For restore plans, this step is not available when restoring VMs to a Microsoft Hyper-V environment.

You can override the following parameters for the **Start Service** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	Yes
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	10
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.	In-Guest OS (not editable)
Credentials	Credentials required to gain access to the in-guest OS.	—

Parameter	Description	Default Value
Service Name	Name of the service to start. **Note**: A short ServiceName must be used.	–

Veeam Job Actions

This step allows you to perform the following job actions required to support plans: enable, disable, start and stop. For example, this step can be useful if you need to disable existing jobs that use storage snapshots in Veeam Backup & Replication.

> **NOTE**
>
> Orchestrator can recover machines protected by backup and replication jobs only.

You can override the following parameters for the **Veeam Job Actions** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No
Wait for Job to Complete	Defines whether Orchestrator will wait for the action to complete before proceeding to the next step.	Yes
Timeout	Maximum amount of time (in seconds) for the step to execute.	600
Retries	Number of retries that will be attempted if the step fails on the first try.	3
Job action	Action to perform on the job.	Enable
Job Name	Name of the job to perform actions on.	–

Verify DNS Port

This step verifies the port used to connect to the recovered VM with the Domain Naming Service role.

You can override the following parameters for the **Verify DNS Port** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	10
Port Number	Port number to check for access to the DNS Service.	53
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.	Veeam backup server (not editable)
Server	Name of the server to check. **Note**: The DNS name or IPv4 address should be used.	*%current_vm_ip_list%*

Verify Domain Controller Port

This step verifies the port used to connect to the recovered VM with the Active Directory DC role.

You can override the following parameters for the **Verify Domain Controller Port** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes

Parameter	Description	Default Value
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	10
Port Number	Port number to check for access to the LDAP AD Service.	389
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.	Veeam backup server (not editable)
Server	Name of the server to check. **Note:** The DNS name or IPv4 address should be used.	*%current_vm_ip_list%*

Verify Exchange Mailbox

This step verifies the Exchange Mail Server accessibility.

IMPORTANT

1. To allow the script to run inside the guest OS of a processed machine, it is required that you have Microsoft PowerShell 3.0, .Net Framework 4.0 and Exchange Server 2010 (or later) installed on each machine for which you enable this step.
2. To allow the script to gain access to the Exchange Mail Server to verify the Exchange mailbox, the Exchange Server must run Microsoft Exchange Web Services Managed API 2.1 (or later). The script will use the EWS Managed API to access the server.
3. To allow the script to verify the Exchange mailbox, the Exchange Mail Server must run Microsoft Windows Server 2008 R2 (or later).
4. The account used to run the script must have the *ApplicationImpersonation* permissions on the Exchange Mail Server. However, keep in mind that once you assign these permissions to the account, the Active Directory synchronization process may take up to 15 minutes for one Active Directory Site (and longer if there are multiple AD Sites involved).

   After the synchronization process is over, replicate or back up your Lab Group Domain Controller so the account used to test plans also has the permissions. To learn how to manage impersonation rights, see this CodeTwo KB article.

You can override the following parameters for the **Verify Exchange Mailbox** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	10
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.	In-Guest OS (not editable)
Windows Credentials	Credentials required to gain access to the in-guest OS.	—
Exchange Credentials	Credentials required to gain access to the Exchange mailbox.	—
Exchange Server	Name of the machine where the Microsoft Exchange Web Services Managed API runs.	*%vm_fqdn%*
Email Address	Email address of the mailbox to check.	—

Verify Exchange MAPI Connectivity

This step logs on to all active databases on the local server to verify connectivity to the system mailbox.

IMPORTANT

1. To allow the script to run inside the guest OS of a processed machine, it is required that you have Microsoft PowerShell 3.0, .Net Framework 4.0 and Exchange Server 2013 (or later) installed on each machine for which you enable this step.
2. To allow the script to verify connectivity to the mailbox, the local server must run Microsoft Windows Server 2008 R2 (or later).

You can override the following parameters for the **Verify Exchange MAPI Connectivity** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	10
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.	In-Guest OS (not editable)
Credentials	Credentials required to gain access to the in-guest OS.	—

Verify Exchange Services

This step verifies that Microsoft Exchange services are running on the recovered VM.

IMPORTANT

1. To allow the script to run inside the guest OS of a processed machine, it is required that you have Microsoft PowerShell 3.0, .Net Framework 4.0 and Exchange Server 2013 (or later) installed on each machine for which you enable this step.
2. To allow the script to verify that the services are running on the selected machine, the machine must run Microsoft Windows Server 2008 R2 (or later).

You can override the following parameters for the **Verify Exchange Services** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes

Parameter	Description	Default Value
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	10
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.	In-Guest OS (not editable)
Credentials	Credentials required to gain access to the in-guest OS.	—
Exchange Server	Name of the machine where Microsoft Exchange Server runs.	*%vm_fqdn%*

Verify Global Catalog Port

This step verifies the port used to connect to the recovered VM with the Global Catalog role.

You can override the following parameters for the **Verify Global Catalog Port** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No
Timeout	Maximum amount of time (in seconds) for the step to execute.	300

Parameter	Description	Default Value
Retries	Number of retries that will be attempted if the step fails on the first try.	10
Port Number	Port number to check for access to the LDAP GC Service.	3268
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.	Veeam backup server (not editable)
Server	Name of the server to check. **Note**: The DNS name or IPv4 address should be used.	*%current_vm_ip_list%*

Verify Mail Server Port

This step verifies the port used to connect to the recovered VM with the Mail Server role.

You can override the following parameters for the **Verify Mail Server Port** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	10
Port Number	Port number to check for access to the SMTP Service.	25
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.	Veeam backup server (not editable)

Parameter	Description	Default Value
Server	Name of the server to check. **Note:** The DNS name or IPv4 address should be used.	*%current_vm_ip_list%*

Verify SharePoint URL

This step verifies the SharePoint Server accessibility.

IMPORTANT

1. To allow the script to run inside the guest OS of a processed machine, it is required that you have Microsoft PowerShell 3.0, .Net Framework 4.0 and SharePoint 2013 (or later) installed on each machine for which you enable this step.
2. To allow the script to verify the SharePoint Server accessibility, the server must run Microsoft Windows Server 2008 R2 (or later).
3. The account used to run the script must be assigned the *SharePoint_Shell_Access* role and must be a member of the *WSS_ADMIN_WPG* group on each processed machine.

You can override the following parameters for the **Verify SharePoint URL** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	10
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.	In-Guest OS (not editable)
SharePoint URL	Name of the SharePoint Server to check.	—

Parameter	Description	Default Value
Credentials	Credentials required to gain access to the in-guest OS.	—

Verify SQL Database

This step verifies the SQL database instance accessibility.

IMPORTANT

1. To allow the script to run inside the guest OS of a processed machine, it is required that you have Microsoft PowerShell 3.0, .Net Framework 4.0 and SQL Server 2008 (or later) installed on each machine for which you enable this step.
2. To allow the script to verify the SQL database instance accessibility, the verified SQL Server instance must run Microsoft Windows Server 2008 R2 (or later).
3. To allow the script to connect to the verified SQL Server instance, Orchestrator uses an account whose credentials are specified as values for either the Windows Credentials or SQL Credentials parameter. The account must have the *VIEW ANY DATABASE* permission. For more information, see Microsoft Docs.

You can override the following parameters for the **Verify SQL Database** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	10
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.	In-Guest OS (not editable)

Parameter	Description	Default Value
Windows Credentials	Credentials required to gain access to the SQL Server instance that uses *Windows Authentication*. If the SQL Server instance you want to check uses *Windows Authentication*, the provided credentials will be used to connect to both the in-guest OS and the SQL instance. In this case, specify a value for the **Windows Credentials** parameter, and leave the SQL Credentials parameter value empty.	—
SQL Credentials	SQL account credentials required to gain access to the SQL Server instance that uses *SQL Server Authentication*. If the SQL Server instance you want to check uses *SQL Server Authentication*, Orchestrator will use SQL Server credentials to access the SQL instance, and Windows credentials to access the in-guest OS. That is why you must specify values for both the **SQL Credentials** and Windows Credentials parameters. **Note**: The **SQL Credentials** parameter does not accept Windows credentials.	—
SQL Instance	Name of the SQL instance to check. **Note**: To check all detected instances, select *ALL*.	ALL
SQL Database	Name of the SQL database to check. **Note**: To check all detected databases, select *ALL*.	ALL

Verify SQL Port

This step verifies the port used to connect to the recovered VM with the SQL Server role.

You can override the following parameters for the **Verify SQL Port** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes

Parameter	Description	Default Value
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	10
Port Number	Port number to check for access to the SQL Service.	1433
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.	Veeam backup server (not editable)
Server	Name of the server to check. **Note**: The DNS name or IPv4 address should be used.	*%current_vm_ip_list%*

Verify Web Server Port

This step verifies the port used to connect to the recovered VM with the Web Server role.

You can override the following parameters for the **Verify Web Server Port** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	10

Parameter	Description	Default Value
Port Number	Port number to check for access to the Web Service.	80
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.	Veeam backup server (not editable)
Server	Name of the server to check. **Note**: The DNS/NetBIOS name or IPv4 address should be used.	*%current_vm_ip_list%*

Verify Web Site (IIS)

This step verifies the website accessibility.

IMPORTANT

1. To allow the script to run inside the guest OS of a processed machine, it is required that you have Microsoft PowerShell 3.0, .Net Framework 4.0 and IIS 8.0 (or later) installed on each machine for which you enable this step.
2. To allow the script to verify the website accessibility, the processed machine must run Microsoft Windows Server 2008 R2 (or later).

You can override the following parameters for the **Verify Web Site (IIS)** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	Yes
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	No
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	10

Parameter	Description	Default Value
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS.	In-Guest OS (not editable)
Credentials	Credentials required to gain access to the in-guest OS.	—
Website Name	Name of the website to check.	Default Web Site

VM Power Actions

This step allows you to perform the following VM power actions required to support recovery plans: power on, power off, shutdown, suspend and resume. For example, this step can be useful if you need to power off non-critical VMs in the DR site to free resources required for recovery.

> **IMPORTANT**
>
> Do not use this step to perform power actions on replica VMs. These actions are automatically performed during the Process Replica VM step execution.

You can override the following parameters for the **VM Power Actions** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for VM recovery. If you mark the step as *Critical*, its failure for a VM from a critical inventory group will halt the plan.	No
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab. **Note**: If you set the parameter value to *Execute*, keep in mind that all actions performed while testing the plan will not be reverted when the test is over.	No
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations. **Note**: If you set the parameter value to *Execute*, the step will perform an action opposite to that specified for the **Power Action** parameter. For example, if you set the **Power Action** parameter value to *Power On*, the step will power off the selected VMs during the Failback and Undo Failover operations.	Yes

Parameter	Description	Default Value
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	1
Action	Power action to perform on the VMs.	Power On
vCenter Name	Name of the vCenter Server where the VMs are located. **Note**: Use separate step instances for each VMware connection.	—
VM Names	Names of the VMs to perform power actions on. **Note**: To add multiple VMs, use commas.	—

Custom Script

If you have customized your own script to be used when running a recovery plan, you can override the following parameters for the **Custom Script** step:

> **IMPORTANT**
>
> To allow the script to run inside the guest OS of a processed machine, it is required that you have Microsoft PowerShell 3.0 and .Net Framework 4.0 installed on each machine for which you enable this step.

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	No
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Yes
Run Step During Undo and Failback	Defines whether the step will be executed during the Failback and Undo Failover operations.	Yes
Timeout	Maximum amount of time (in seconds) for the step to execute.	600

Parameter	Description	Default Value
Retries	Number of retries that will be attempted if the step fails on the first try.	3
Script Execution Location	Defines whether the script will run on the Veeam Backup & Replication server, on the Orchestrator server or on the in-guest OS. **Note**: The script cannot be executed on the in-guest OS when restoring VMs to a Microsoft Hyper-V environment.	Veeam backup server
Windows Credentials	Credentials required to gain access to the in-guest OS. Applies only if the **Script Execution Location** parameter value is set to *In-Guest OS*.	—

Protect VM Group

When you add an inventory group to an existing plan, you have an option to run the **Protect VM Group** step to protect machines included in the plan.

This step creates a new template-based backup or replication job to back up or replicate machines in the specified inventory group as soon as the recovery process completes. Note that for replica plans, the template job will run only after the plan enters the *PERMANENT FAILOVER* state.

> **IMPORTANT**
>
> The Veeam Backup & Replication server on which the template backup or replication job has been created must be connected to the vCenter Server that manages target machines.

You can override the following parameters for the **Protect VM Group** step:

Parameter	Description	Default Value
Critical Step	Defines whether the step is critical for machine recovery. If you mark the step as *Critical*, its failure for a machine from a critical inventory group will halt the plan.	No
Run Step During a DataLab Test	Defines whether the step will be executed during plan testing in a DataLab.	Skip (not editable)
Timeout	Maximum amount of time (in seconds) for the step to execute.	300
Retries	Number of retries that will be attempted if the step fails on the first try.	0

Parameter Variables

Parameter variables are passed between steps during plan execution, and are available when editing and adding step parameters.

You can define the following variables for any step parameter that allows plain text to be entered:

Parameter Variable	Content
%vm_fqdn%	FQDN of the currently processed machine
%current_vm_ip%	IP address of the currently processed VM
%current_vm_name%	Name of the currently processed VM
%current_vm_ip_list%	All IP addresses of the currently processed VM
%target_vm_ip%	IP address of the target VM
%target_vm_ip_list%	All IP addresses of the target VM
%target_vm_name%	Name of the target VM
%replica_vm_state%	State of the currently processed replica VM
%source_machine_name%	Name of the source machine
%source_vm_ip%	IP address of the source VM
%source_vm_ip_list%	All IP addresses of the source VM
%source_vm_name%	Name of the source VM
%plan_name%	Name of the recovery plan
%vao_server_name%	Name of the Orchestrator server
%plan_state%	Current state of the recovery plan
%plan_test_mode%	Boolean variable that indicates whether the plan is currently being tested (True/False)

Parameter Variable	Content
%group_name%	Name of the currently processed inventory group
%plan_summary%	Output information on the recovery plan (error/warning/success for all steps)
%group_summary%	Output information on the currently processed inventory group
%vm_summary%	Output information on the currently processed VM
%vao_ui%	URL to access the home page of the Orchestrator UI
%plan_vms%	List of all machines included in the recovery plan

Appendix B. Getting Technical Support

Veeam offers email and phone technical support for customers on maintenance and during the official evaluation period. For a better experience, provide the following details when contacting Veeam Customer Support:

- Version information for the product and its components
- Error message or accurate description of the problem you are facing
- Log files

For your convenience, the Orchestrator UI allows you to collect logs for each Orchestrator component separately. To do that:

1. Switch to the **Administration** page.
2. Navigate to **Logs**.
3. Select check boxes next to the necessary components.
4. Click **Download Logs**. Logs will be saved locally in the default download folder.

> **NOTE**
>
> Every archive with log files that you download contains an anonymized file with the current Orchestrator configuration and statistical information. This file can be used by Orchestrator product management to improve the product. No information will be shared outside of Veeam at any time.