

## Veeam Disaster Recovery Orchestrator Release Notes

This document provides last-minute information about Veeam Disaster Recovery Orchestrator (Orchestrator), including system requirements and relevant information on technical support, documentation, online resources and so on.

The current version of Veeam Disaster Recovery Orchestrator is available for download at [veeam.com/disaster-recovery-orchestrator-download.html](https://www.veeam.com/disaster-recovery-orchestrator-download.html) starting from November 23, 2021.

### **NOTE:**

The release build of Veeam Disaster Recovery Orchestrator is 5.0.0.2768, while the embedded Veeam Backup & Replication build is 11.0.1.1261 and the embedded Veeam ONE build is 11.0.1.1880.

### **See next:**

- [System Requirements](#)
- [Required Permissions](#)
- [What's New](#)
- [Known Issues](#)
- [Technical Documentation References](#)
- [Technical Support](#)
- [Company Contacts](#)

# System Requirements

Unless otherwise stated, all 3rd party software must be at the latest update or patch level.

## Hardware

Hardware requirements depend on the size of the managed infrastructure.

Number of Virtualization Hosts*	100	100-500	500-1000	>1000
<b>CPU</b>	6 vCPUs (minimum) – 8 vCPUs (recommended) for the Orchestrator server  4 vCPUs (minimum) – 8 vCPUs (recommended) for the Microsoft SQL Server database	8 vCPUs (minimum) – 12 vCPUs (recommended) for the Orchestrator server  8 vCPUs (minimum) – 12 vCPUs (recommended) for the Microsoft SQL Server database	12 vCPUs (minimum) – 16 vCPUs (recommended) for the Orchestrator server  12 vCPUs (minimum) – 16 vCPUs (recommended) for the Microsoft SQL Server database	>16 vCPUs for the Orchestrator server  >16 vCPUs for the Microsoft SQL Server database
<b>Memory</b>	6 GB (minimum) – 8 GB (recommended) for the Orchestrator server  4 GB (minimum) – 8 GB (recommended) for the Microsoft SQL Server database	8 GB (minimum) – 40 GB (recommended) for the Orchestrator server  8 GB (minimum) – 40 GB (recommended) for the Microsoft SQL Server database	40 GB (minimum) – 70 GB (recommended) for the Orchestrator server  40 GB (minimum) – 70 GB (recommended) for the Microsoft SQL Server database	>70 GB for the Orchestrator server  >70 GB for the Microsoft SQL Server database
<b>Hard Disk Space</b>	50 GB for product installation and sufficient disk space for Microsoft SQL Server (if installed locally)			

\* Under the condition that 1 host has 20 VMs deployed.

## OS

Only 64-bit versions of the following operating systems are supported:

- Microsoft Windows Server 2022
- Microsoft Windows Server 2019
- Microsoft Windows Server 2016
- Microsoft Windows Server 2012 R2
- Microsoft Windows Server 2012
- Microsoft Windows Server 2008 R2 SP1

It is not recommended to install Orchestrator on a machine running Microsoft Windows Server Core, or on a Domain Controller.

## User Management

Windows domain-joined machine.

## Microsoft SQL Server

Local and remote installations of the following versions of Microsoft SQL Server are supported:

- Microsoft SQL Server 2019
- Microsoft SQL Server 2017
- Microsoft SQL Server 2016 (2016 SP2 Express Edition is included in the setup)
- Microsoft SQL Server 2014
- Microsoft SQL Server 2012 (2012 SP4 Express Edition is included in the setup)

It is not recommended to use the Express Edition in any production Orchestrator deployments – it should only be used for product evaluation.

## Veeam Software

If a standalone Veeam Backup & Replication server is used for backup and replication task management, it must be version 11 or later. Version 11a is fully supported.

If you connect a standalone Veeam Backup & Replication server to the Orchestrator server, it must be joined to either the same domain, a trusted domain or a workgroup. Only connections to Veeam Backup & Replication servers deployed on Microsoft Windows Server OS are supported.

## Additional Software

All components will be installed during setup.

For inline Report Template editing, Microsoft Word component of SP2 for Microsoft Office 2010 or later is required.

## Virtualization Platform

- VMware vSphere 6.0, 6.5, 6.7, 7.0 (up to 7.0 Update 3)
- NSX-T Virtual Distributed Switch (N-VDS) networks v2.4 and later

The Orchestrator server must be connected to VMware vCenter Servers. Direct connections to vSphere hosts are not supported.

## Storage System

- HPE 3PAR 3.3.1 MU5
- HPE Primera 4.2 and 4.3
- HPE Alletra 9000
- NetApp ONTAP 9.3, 9.4, 9.5, 9.6, 9.7, 9.8, 9.9
- Lenovo DM 9.3, 9.4, 9.5, 9.6, 9.7, 9.8, 9.9

# Required Permissions

The accounts used for installing and using Veeam Disaster Recovery Orchestrator must have the following permissions.

Account	Required Permission
<b>Setup Account</b>	The account used for product installation must be a domain user who has the Local Administrator permissions on the target machine.
<b>Orchestrator Service Accounts</b>	The accounts used to run Orchestrator services, Veeam Backup & Replication services and Veeam ONE services must have Local Administrator permissions on the Orchestrator server. The accounts must also be granted the <i>Log on as a service</i> right. For more information on Windows security policy settings, see <a href="#">Microsoft Docs</a> .
<b>Orchestrator Agent Account</b>	The account used to install and run the Orchestrator agent on a Veeam Backup & Replication server must have both the local Administrator and the Veeam Backup Administrator permissions on the server.
<b>Orchestrator User Accounts</b>	The accounts used to log in to the Orchestrator UI must be granted the <i>Allow log on locally</i> right. For more information on Windows security policy settings, see <a href="#">Microsoft Docs</a> .
<b>vCenter Server Permissions</b>	The account used to connect the vCenter Server to the Orchestrator infrastructure must have administrative permissions. You can either grant the Administrator role to the account or configure more granular permissions. For more information, see <a href="#">Veeam Backup &amp; Replication Required Permissions</a> and <a href="#">Veeam ONE Required Permissions</a> . To be able to open sessions on the vCenter Server system, the account must also have the <i>Sessions.Validate session</i> privilege on the root vCenter Server. For more information on session privileges, see <a href="#">VMware Docs</a> .
<b>NetApp Storage System Permissions</b>	The account used to connect the storage system to the Orchestrator infrastructure must be granted permissions described in section <a href="#">NetApp Data ONTAP Permissions</a> .
<b>HPE Storage System Permissions</b>	The account used to connect the storage system to the Orchestrator infrastructure must be assigned the <i>Super</i> or <i>Edit</i> role. If the account is assigned the <i>Edit</i> role, both the account and the storage resources that you plan to access must belong to the same domain. <b>Note:</b> Multiple connections to a storage system using different credentials are not supported.
<b>Microsoft SQL Server</b>	Different sets of Microsoft SQL permissions are required in the following cases: <ul style="list-style-type: none"> <li>• <b>Installation</b> (remote or local): the current account needs the <i>CREATE ANY DATABASE</i> permission on the SQL Server level. After the database is created, this account automatically gets a <i>db_owner</i> role and can perform all operations with the database.</li> <li>• <b>Operation:</b> the account used to run Orchestrator services, Veeam Backup &amp; Replication services and Veeam ONE services requires the <i>db_owner</i> role, as well as permissions to execute stored procedures for the configuration databases on the Microsoft SQL Server.</li> </ul> For more information, see <a href="#">Veeam Backup &amp; Replication Required Permissions</a> and <a href="#">Veeam ONE Required Permissions</a> .
<b>Orchestrator Step Accounts</b>	The account used to run the <b>Verify SharePoint URL</b> step, must be assigned the <i>SharePoint_Shell_Access</i> role and must be a member of the <i>WSS_ADMIN_WPG</i> group on the processed VM. The account used to run the <b>Verify Exchange Mailbox</b> step, must be assigned the <i>ApplicationImpersonation</i> role on the processed VM.

# What's New

## Storage orchestration plans for HPE Primera, 3PAR and Alletra 9000 storage systems

Directly orchestrate failover from Hewlett Packard Enterprise (HPE) storage systems, leveraging HPE native snapshot replication, to reduce the complexity of the recovery process.

## Streamlined operations with NEW enhanced dashboard

At-a-glance view of your target environment state alerts you to detect potential RTO and RPO violations before they impact your recovery.

## Simplified plan workflow with NEW inventory view

Accelerate orchestration plan creation with visibility into your whole environment.

# Known Issues

## Technical limitations

- Orchestrator does not support failover orchestration for volumes protected by SnapVault replication. If you create and try to run a NetApp storage plan containing at least one datastore protected by SnapVault replication, the whole plan will fail to execute.
- Orchestrator does not support failover orchestration for SVMs protected by SnapMirror SVM replication. NetApp storage failover is orchestrated at the volume level only.
- Orchestrator does not support storage failover of VMs with RDM disks. When Orchestrator tries to register such a VM in a vCenter Server while executing a storage plan, the server returns the following error: *"Unable to enumerate all disks."*
- Restore plans created in Orchestrator can orchestrate recovery from backup files stored only in repositories of the following storage types: direct-attached storage (both Microsoft Windows and Linux), network-attached storage (SMB shares) and deduplicating storage appliances (Dell EMC Data Domain, ExaGrid, HPE StoreOnce and Quantum DXi).

If you plan to run DataLab tests using deduplicating storage appliances, note the limitations described in [this Veeam KB article](#)).

- Orchestrator does not support tape storage.
- Orchestrator does not support testing of orchestration plans in DataLabs with network mapping from VSS\VDS to N-VDS.

## Installation

- During product installation, it is possible to set any custom port numbers except 9402, 8543 and 20443. This is because the embedded Veeam Backup & Replication is hardcoded to use those ports.
- When you follow the Orchestrator installation wizard as described in the [Veeam Disaster Recovery Orchestrator User Guide](#), you have an option to use an existing Microsoft SQL Server instance to host the Orchestrator database. When choosing the instance, keep in mind that Orchestrator does not support case-sensitive SQL Server databases.
- By default, all Orchestrator Server components are installed all together on a single machine. If you want to install Veeam ONE Monitor Client separately on a Windows Desktop OS, run the .MSI file available in the Orchestrator installation folder by the following path:  
`<Orchestrator>\monitor\VeeamONE.Monitor.Client.x64.msi`

If you want to install Veeam ONE Monitor Client separately on a Windows Server OS, simply launch the Orchestrator setup wizard, and follow the instructions provided in the Veeam Disaster Recovery Orchestrator User Guide, section [Installing Additional Veeam ONE Clients](#).

## Upgrade

- Performance issues on the Orchestrator Server that occur during upgrade may result in issues with Orchestrator agents:
  - The embedded Veeam Backup & Replication server may show an incorrect version. To work around the issue, restart the Veeam Orchestrator Server Service.
  - Remote Veeam Backup & Replication servers may show a timeout issue. To work around the issue, repair the affected Orchestrator agents.
- After upgrade, remote users logged in to the Orchestrator UI using the Firefox browser with the default file cache settings may experience errors.

To work around the issue, reload the page by pressing F5.

- If you try to upgrade Orchestrator to version 5.0 while the existing Orchestrator 4.0 databases are protected with SQL replication, the upgrade process will fail with an error. This is a known issue related to specific aspects of SQL replication.
- If you try to upgrade Orchestrator to version 5.0 while Orchestrator 4.0 is installed to a directory with a non-ASCII name, Orchestrator will not function properly anymore even though the upgrade process will complete successfully.

To work around the issue, reinstall Orchestrator 4.0 to another directory using the existing databases, and then upgrade to Orchestrator 5.0.

## Infrastructure

- If you change the certificate on a vCenter Server added to the Orchestrator infrastructure, Orchestrator may fail to connect to the server.

To work around the issue, edit the vCenter Server connection using the Orchestrator UI — step through the **Edit VMware vCenter Server** wizard, accepting the certificate when prompted.

- If you have a VMware vCloud Director server added to your backup infrastructure, Orchestrator will not be able to process VMs managed by vCenter Servers connected to this server.
- To let Orchestrator connect to the Microsoft SQL Server instance that hosts the Orchestrator database, you must specify authentication credentials during Orchestrator Server installation. However, if you change the password later, you will not be able to modify the provided credentials, and Orchestrator will fail to connect to the SQL Server instance.

To work around the issue, it is recommended that you use Windows Authentication when choosing credentials to connect to the Microsoft SQL Server. Alternatively, [Veeam Customer Support](#) can help you resolve the issue.

## User roles

- When you perform initial configuration of the Orchestrator server, the Initial Configuration Wizard allows you to add users that will be assigned the *Administrator* role for the server. Please keep in mind that you cannot assign the role to a local Administrator. You cannot add local Administrators to any of the user groups provided by Orchestrator. Only domain accounts are supported.

## Recovery locations

- If the name of a Virtual Distributed Switch used in recovery location network mapping contains the backslash character, the recovery of VMs will fail.

The issue will be addressed in a future update.

- After you restart the Veeam Orchestrator Server Service or reboot the Orchestrator Server, storage plans may fail to complete successfully and readiness checks may show that the plans are not ready for failover.

To work around the issue, wait 5-10 minutes for the inventory to be refreshed.

- [Applies only to Veeam Backup & Replication 11; the issue has been resolved in v11a] If network mapping configured for a CDP policy has 2 or more identical target networks, Orchestrator will be unable to fail the VM back to a new location.
- When you create a recovery location, you must specify target hosts where recovered VMs will be registered. If you add an ESXi host to a recovery location as Compute Resources and then move the host to another datacenter, the host will be assigned a new moRef, Orchestrator will consider the host to be another 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.
- Orchestrator can discover only those vSphere Distributed Switches that are connected to at least one ESXi host added to any vCenter Server managed by Orchestrator. Only in this case you will be able to add such a DVS to network resources when configuring recovery locations in the Orchestrator UI.

- When Orchestrator restores VMs to a recovery location which has Instant VM Recovery enabled, VMs are added to the root VM folder.

To work around the issue, move VMs to the required folder manually.

## VM groups

- Under certain circumstances, a plan may be halted after one or more **Register VM** steps have started, but not completed, causing orphaned VMs to be created in the recovery location. These VMs should be removed manually.

The issue will be addressed in a future update.

- If a VM has its data disk (.vmdk) and configuration file (.vmx) stored on separate datastores, and the VM was suspended when storage failover was run, then the datastore where the .vmx file is located must be processed by Orchestrator first. If Orchestrator attempts to process the datastore with the data disk first, the VM will fail to power on as the .vmx file will not contain the correct datastore GUIDs.

To work around the issue, edit the .vmx file manually after all storage processing is complete. Contact Veeam Support for assistance if required.

- If a VM has a name that starts with a dot, Orchestrator will not be able to process that VM. This is a known VMware issue – some functionality is limited for infrastructure objects whose names start with a dot.
- If a VM stores its files on a dedicated swap datastore, Orchestrator will exclude the VM from the VM group related to that datastore – but only in case the VM is protected by storage replication. If the VM is not protected by storage replication, Orchestrator will still include the VM in the VM group.

## Orchestration plans

- Removing a CDP policy will result in the failure to map IP addresses from the failover range to the new location range.

Do not remove CDP policies while the related plans are in progress.

- [Applies only to Veeam Backup & Replication 11; the issue has been resolved in v11a] Renaming a CDP policy will result in the failure to retrieve the SLA for the policy and also the failure to build correct network mapping for the VM.

Do not rename CDP policies while the related plans are in progress.

- Orchestrator does not support a scenario where you add a vCenter Server to the Orchestrator infrastructure using its IP address and to the standalone Veeam Backup & Replication server using its FQDN. In this case, Orchestrator will not be able to execute orchestration plans successfully, and the plans will halt with the following error: *"Cannot find the host on the VBR Server."*
- If you connect a vCenter Server to a standalone Veeam Backup & Replication server added to Orchestrator, you must also connect the same vCenter Server to the Orchestrator server. Otherwise, Orchestrator will fail to map VM inventory correctly and will be unable to locate VM backups and replicas when running orchestration plans. As a result, the plans will fail to complete.
- Another reason why Orchestrator may fail to map VM inventory correctly is that, by design, Veeam ONE collects historical data from the newly connected Veeam Backup & Replication servers for the previous week only. In this case, Orchestrator will also be unable to locate VM backups and replicas when running orchestration plans.

To work around the issue, run backup and replication jobs that has not been run for more than one week.



- If you add new VMs to a VM group, these VMs will not be added to an orchestration plan that contains the VM group until the plan enters the *VERIFIED* or *NOT VERIFIED* state. This applies to all other stable and active states that the plan may acquire.

For more information on plan states, see the Veeam Disaster Recovery Orchestrator User Guide, section [Working with Orchestration Plans](#).

- If you have created multiple backup or replication jobs for one VM, and each job has its own target location, you will not be able to select which backup or replica to use when running orchestration plans. By design, Orchestrator will recover the VM that has the most recent restore point.
- To protect a VM after failover, you must configure a template backup job on the Veeam Backup & Replication server that processes the VM. Otherwise, Orchestrator will not be able to locate the recovered VM, and the *Protect VM Group* step will fail to complete.
- If a plan includes steps that require in-guest scripts to run on virtual machines being recovered, you may get the following error: *"RPC error: There are no more endpoints available from the endpoint mapper."*

To work around the issue, add the *Shutdown Source VM* step to the plan. This will power off source VMs during the recovery process, and the in-guest scripts will be able to run on VMs successfully.

- For datastores connected through the NFSv4.1 protocol, Orchestrator supports failover to a recovery location only in case 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.
- When performing storage failover, Orchestrator does not check whether iSCSI CHAP authentication and NetApp Kerberos authentication is configured for the datastores mounted in the DR site.

To work around the issue, run a DataLab test before performing storage failover to make sure the authentication is configured properly.

- Orchestrator does not support storage failover in a situation where both the target storage system and source storage system have the same name, the source volumes and target volumes have the same name, and the destination LUN and destination volume have the same name too. In this case, Orchestrator will be unable to identify the target volumes, and the NetApp storage plans will halt with errors.

To work around the issue, rename the target storage system, source storage system, destination LUN, source volumes and target volumes. However, when renaming these objects, keep in mind that Orchestrator also does not support NetApp storage failover in a situation where storage systems have alias names.

- If a datastore is backed by several volumes (in a situation where LUNs of the datastore reside on several volumes), you will be able to protect such configuration with different SVMs and to create different storage recovery locations for these SVMs. However, in this case, Orchestrator will be unable to match the locations for the datastore, and the related plans will halt with errors.
- If you suspend a VM that stores its disk files on multiple datastores before creating a SnapMirror snapshot, and then include this VM in a storage plan, Orchestrator will not be able to execute the plan successfully. The plan will halt with the following error: *"Unable to access the virtual disk from the host. Either the host is disconnected from the datastore or has insufficient privilege."*
- At the moment, Orchestrator does not support orchestration of storage failover for objects of the type *Tag* and *VM Folder* on standalone Veeam Backup & Replication servers. No matter if you have a backup job with these objects configured on a standalone Veeam Backup & Replication server, Orchestrator will still execute and test the related storage plans on the embedded Veeam Backup & Replication server.

## Reports

- When you edit a report template, save it first before closing. If you try closing the file before you save it, MS Word will offer you to save changes, but it will lock the file for up to 10 minutes.

## VM console

- After you restart the Veeam Orchestrator Server Service, Orchestrator may fail to open the VM console and may display the "*Connection Failure*" error.

To work around the issue, close the console and try opening it several minutes later.

## Orchestrator agents

On the **Orchestrator Agents** tab of the **Administration** section of the Orchestrator UI you have an option to **Repair** a failed Orchestrator agent installed on a Veeam Backup & Replication server. However, when you try to do that, you may get the following error: "*Repair failed -> The specified service has been marked for deletion.*"

You will encounter this error in case the Orchestrator agent is locked by another service running on the server. To work around the issue, log in to the server and make sure that:

1. The Microsoft Management Console, Task Manager and Process Explorer are closed.

If any of these applications is running, close it.

2. The Veeam Orchestrator Agent Management Service (*VaoAgentSvc*), Veeam Orchestrator Agent for Backup (*VAOBackupAgent*), Veeam Orchestrator Agent for Enterprise Manager (*VAOEmAgent*) have been deleted or have the **StartupType** value set to *Enabled*.

If any of the services has the **StartupType** value set to *Disabled*, restart the server.

Then try repairing the Orchestrator agent again.

## REST API

Orchestrator REST API is backward compatible. However, there is one request that would return an exception — `GET https://<hostname>:<port>/api/v3/Plans`. Use `GET https://<hostname>:<port>/api/v5/Plans` instead.

# Technical Documentation References

If you have any questions about Veeam Disaster Recovery Orchestrator, use the following resources:

- Product webpage: <https://www.veeam.com/disaster-recovery-orchestrator.html>
- User guides: <https://www.veeam.com/documentation-guides-datasheets.html?prd=vao>
- Community forums: <https://forums.veeam.com/>

## Technical Support

Veeam offers email and phone technical support for customers on maintenance and during the official evaluation period. For better experience, please provide the following when contacting Veeam Customer Support:

- Version information for the product and all infrastructure components
- Error message and/or accurate description of the problem you are having
- Log files

For your convenience, the Orchestrator UI allows you to collect logs for each Orchestrator component separately, or all together. To do that:

1. Switch to the **Administration** page.
2. Navigate to **Logs**.
3. Select a check box next to the server where the Orchestrator component runs.
4. Click **Download Logs**. Logs will be saved locally in the default download folder.

### TIP:

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.

To submit your support ticket or obtain additional information, please visit the [Veeam Customer Support Portal](#). Before contacting Veeam Customer Support, consider searching for a resolution on [Veeam Community Forums](#).

## Company Contacts

For the most up-to-date information about company contacts and office locations, please visit the [Veeam Contacts Webpage](#).