VERITAS

Veritas InfoScale™ Deployment on Microsoft Azure Using Solution Templates

Veritas InfoScale™ Enterprise 8.0 - Linux

Last updated: 2025-02-25

About Veritas InfoScale Enterprise

Veritas InfoScale Enterprise enables organizations to provision and manage storage and provides high availability (HA) for business-critical applications. Storage provisioning is independent of hardware types or locations with predictable quality-of-service by identifying and optimizing critical workloads. It increases storage agility and enables you to work with and manage multiple types of storage to achieve better ROI without compromising on performance and flexibility. For application HA, InfoScale Enterprise monitors an application to detect any application or node failure and brings the application services up on a target system in case of a failure.

InfoScale Enterprise also lets you extend workloads to the cloud with Microsoft Azure delivering resiliency, flexibility, and performance at scale.

Some of the key features of InfoScale Enterprise include:

- Accelerated I/O performance using the Veritas SmartIO technology. SmartIO uses instance store-SSD storage, the closest to compute as a data caching device to improve performance. You can use the SmartAssist tool to determine cache size and forecasting of performance gains for custom workloads.
- Efficient and cost-effective method to migrate data to cloud using the Veritas SmartMove technology. SmartMove analyzes storage usage by looking up the Veritas File System (VxFS) metadata and moving only the relevant data to the cloud. This helps you to save on bandwidth and storage costs.
- Redefined storage performance and scalability potential using the Veritas Flexible Storage Sharing (FSS) technology. FSS lets you to provision shared volumes or file systems on shared nothing architectures that redefine the storage, performance and scalability potential of your cloud infrastructure. FSS leverages Azure managed disks to create shared storage in the cloud. By using FSS, you can form a cluster that can be configured within or across availability zones in a single region. In the event of failures-node, storage, or availability zone-data will always be available. Organizations can even run multiple applications in a single FSS cluster and set up SLAs to ensure isolation among applications. This type of deployment unlocks valuable use cases such as running real time analytics on incoming transactional data, or fraud detection on credit card transactions.
- Availability of consistent copy of application data at a remote site to address data center break down using Volume Replicator (VVR). VVR provides volume- and file-level replication capabilities, which ensure that application data is replicated to a remote site (that includes Azure availability zones or regions) to protect against large-scale infrastructure outages. Using VVR, data can be replicated from on-premises to cloud and from cloud to on-premises. With the help of VVR, you can leverage the cloud as a DR site.
- Increased application availability using Cluster Server (VCS). VCS connects multiple, independent systems into a management framework for increased availability. Each system, or node, runs its own operating system and cooperates at the software level to form a cluster. VCS links commodity hardware with intelligent software to provide application failover and control. When a node or a monitored application fails, other nodes can take predefined actions to take over and bring up services elsewhere in the cluster.
- VCS detects the failure of an application by issuing specific commands, tests, or scripts to monitor the overall health of an application. VCS also determines the health of underlying resources by supporting the applications such as file systems and network interfaces. VCS uses a redundant network heartbeat to differentiate between the loss of a system and the loss of communication between systems.

- The VCS fencing module performs membership arbitration to ensure that a split-brain situation does not occur and only one functional cohesive cluster continues to run. It implements a quorum-type functionality to ensure that only one cluster survives a split of the private network. InfoScale provides server-based fencing, where the coordination point servers provide a lock mechanism to determine which nodes survive the network split.

About InfoScale Enterprise deployment using a solution template

Veritas provides the InfoScale Enterprise Solution Template on the Azure Marketplace, which is available for the Red Hat Enterprise Linux (RHEL) 8.4 and the SUSE Linux Enterprise Server (SLES) 15 SP2 platforms. The solution template provides an automated deployment of InfoScale Enterprise on Azure. It sets up and configures the Azure environment for Veritas InfoScale with essential Azure resources, like the virtual machines, the virtual network, and so on.

The template lets you specify the following details for the InfoScale deployment:

- InfoScale component to be used: VCS, SFHA, or SFCFSHA
- InfoScale cluster size: two nodes or four nodes
- Availability options for the Azure virtual machines: within a single availability zone or across two zones, or in an availability set
- Resources to be configured: storage and network
- Fencing configuration: coordination point (CP) servers to provide I/O fencing capability

Using this solution template, InfoScale Enterprise is deployed on the virtual machines with the keyless license. To procure a permanent InfoScale Enterprise 8.0 license, visit the Veritas Licensing Support website at: www.veritas.com/licensing/process

This document (Veritas InfoScale™ Enterprise Deployment on Azure Using Solution Templates) provides instructions for deploying InfoScale Enterprise on Azure by using a solution template. The intended audience for this document includes storage administrators, architects, and system administrators who are planning to deploy the InfoScale Enterprise solution on Azure.

About the deployment scenarios and architecture

You can deploy InfoScale Enterprise using the solution template to set up an InfoScale cluster with up to four nodes. An InfoScale cluster can be configured for either Availability Sets or availability zones. In case of availability zones, the cluster may exist within a single availability zone or span across two zones.

You can implement any of the following configurations by using the appropriate InfoScale components:

- Storage Foundation Cluster File System High Availability (SFCFSHA)
- Storage Foundation and High Availability (SFHA)
- Cluster Server (VCS)

Supported InfoScale version:

- InfoScale Enterprise 8.0

Supported operating systems:

- RHEL 8.4
- SLES 15 SP2 BYOS

Figure 1. InfoScale Enterprise deployment in Azure for an InfoScale cluster

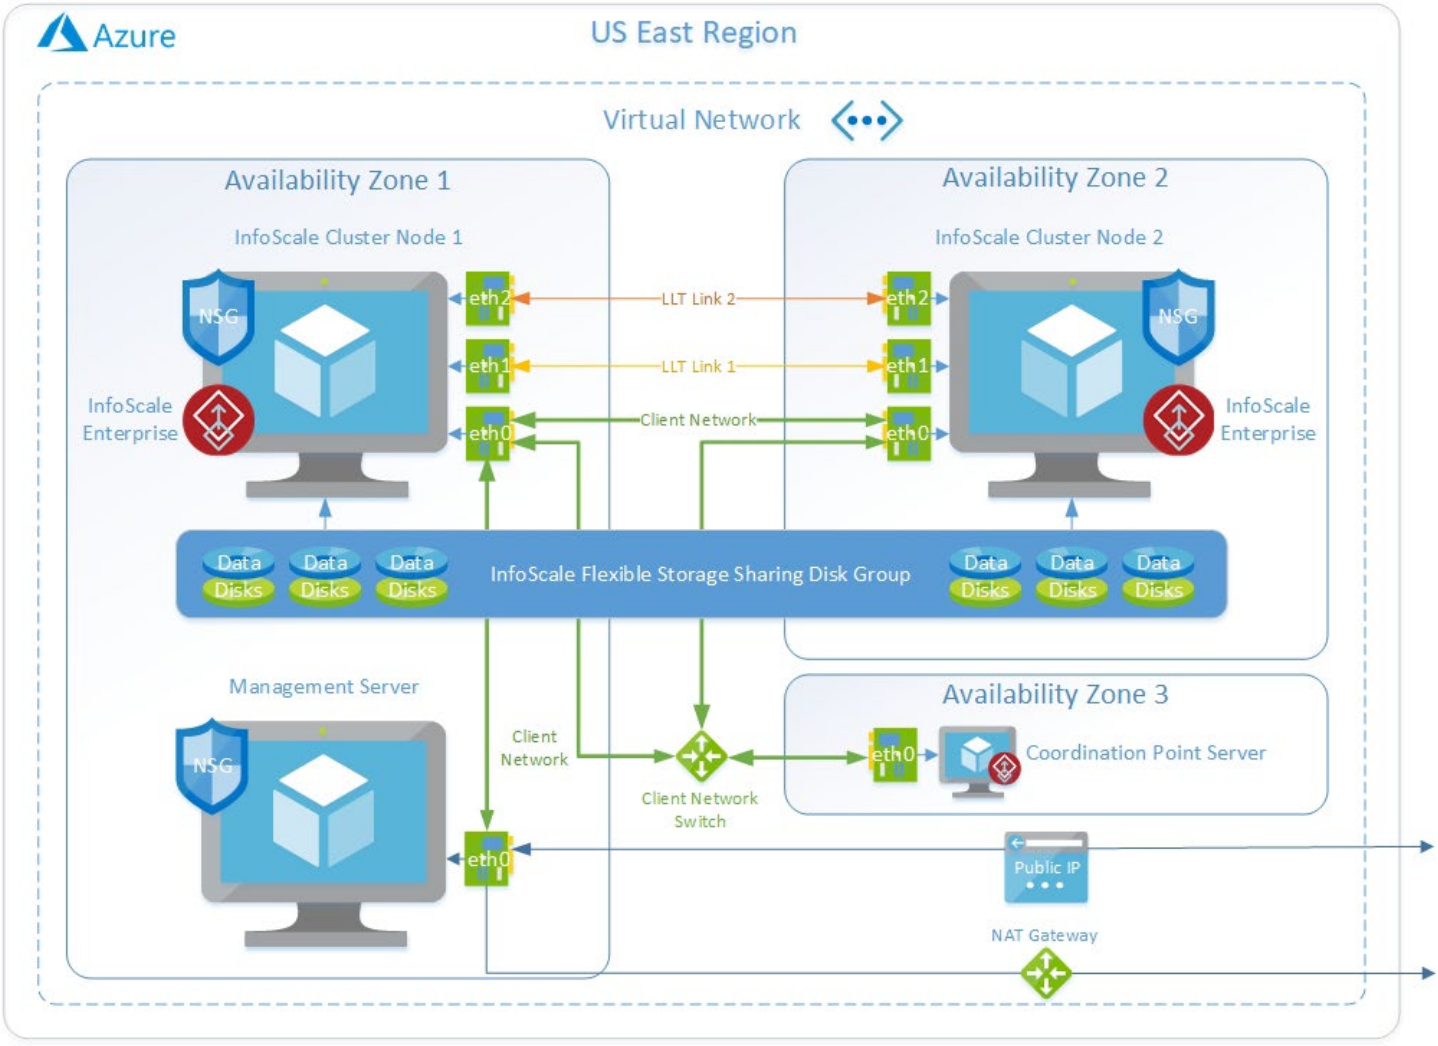

Deploying an InfoScale cluster configuration using the solution template

Use the solution template to deploy an InfoScale cluster in a new virtual network. Accordingly, the various network, storage, and security parameters that you need to set may change.

Figure 2. Workflow for deploying InfoScale Enterprise in a single availability zone or across zones

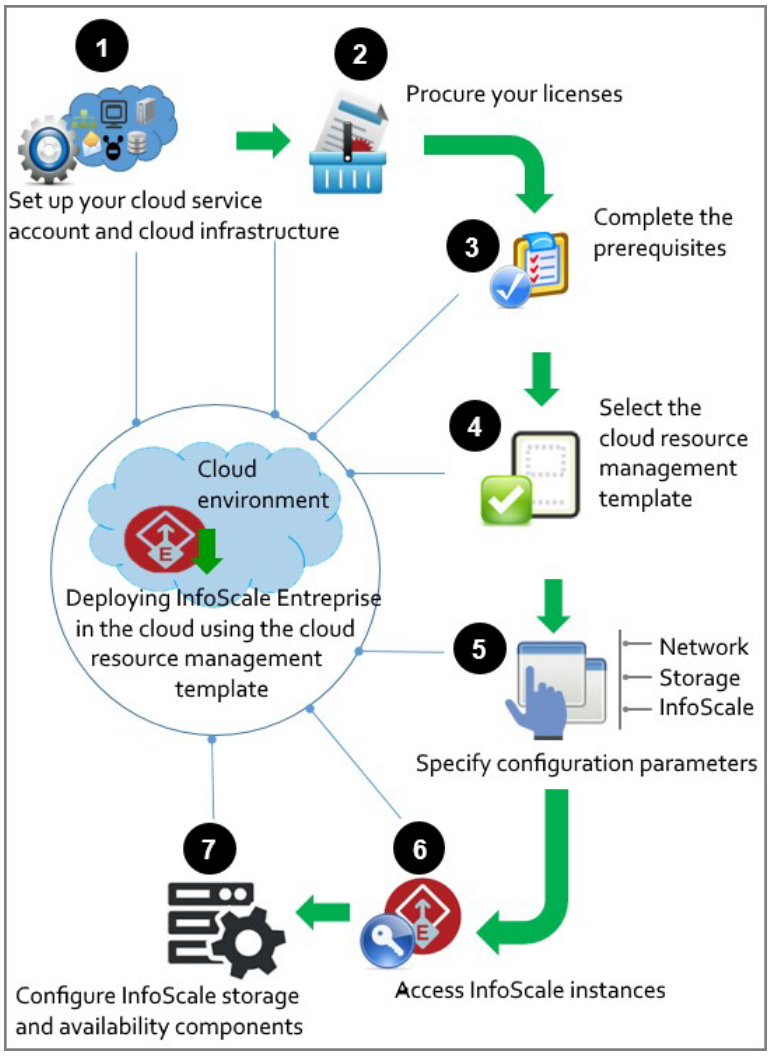

The end-to-end workflow for deploying the cluster consists of the following steps:

- Step1: Address the prerequisites
- Step 2: Deploy InfoScale Enterprise
- Step 3: Access the InfoScale cluster nodes
- Step 4: Perform the post-deployment tasks

Step1: Address the prerequisites

Before you begin to deploy Veritas InfoScale Enterprise on Azure, address the following prerequisites:

- You have an Azure account with an active subscription.
- You have a client ID and a client secret (secret key).
- You have configured the following minimum roles that are required for the deployment of the various resources:
  - Storage Account Contributor
  - Network Contributor
  - Virtual Machine Contributor

Considerations for the storage and the network infrastructure:

- Ensure that the virtual machine size that you choose supports accelerated networking and least three network interfaces (NICs).
  https://docs.microsoft.com/en-us/azure/virtual-network/create-vm-accelerated-networking-cli
- Verify that the virtual machine size and the data disk type that you choose are supported in the region where you plan to deploy Veritas InfoScale Enterprise.
- If you choose to use Azure ultra disks, ensure that they are supported with the virtual machine types that you select. Note that ultra disks are supported with availability zones only and not with availability sets.
  https://docs.microsoft.com/en-us/azure/virtual-machines/linux/disks-enable-ultra-ssd#ga-scope-and-limitations

Step 2: Deploy InfoScale Enterprise

Access the solution template from Azure Marketplace and provide the necessary input to deploy an InfoScale cluster.

To deploy an InfoScale cluster

1. Visit the Azure Marketplace at:
   https://azuremarketplace.microsoft.com/en-us/marketplace/apps/
2. Locate and access the **Veritas InfoScale™ Enterprise** solution template.
3. On the offer page, click **GET IT NOW**.
4. On the **Create this app in Azure** dialog, accept the terms, and click **Continue**.
5. On the **Veritas InfoScale™ Enterprise** page, click **Create**.
6. On the **Basics** blade, provide the following details:

Project details section	
Subscription	Select the subscription ID using which you want to deploy InfoScale Enterprise.
Resource group	Select the resource group under the subscription. Alternatively, you can create a new resource group to be used for the deployment.
Instance details section	
Region	Select the region for the deployment.
Deployment name	Provide the instance name to uniquely identify the deployment and the resources that are associated with it.

AAD application (client) ID	Provide the client ID to be used for authentication with the Azure Active Directory (AAD) application.
AAD application (client) secret	Provide the associated client secret (secret key).
Confirm AAD application (client) secretkey	Provide the associated client secret (secret key) again for confirmation.

7. On the **InfoScale** blade, provide the following details:

Primary details section	
InfoScale version	Select the InfoScale versions to be deployed. Currently, only InfoScale Enterprise 8.0 is supported.
InfoScale component	Select the InfoScale component to be used according to the type of cluster configuration you need to deploy: ▪ Storage Foundation Cluster File System High Availability (SFCFSHA) ▪ Storage Foundation and High Availability (SFHA) ▪ Cluster Server (VCS)
Cluster details section	
Number of cluster nodes	Specify whether you want to configure a two-node or a four-node cluster.
Cluster availability options	Specify whether to use an availability zone or an availability set as the cloud availability configuration. https://docs.microsoft.com/en-us/azure/virtual-machines/linux/availability Not all Azure regions support availability zones; refer to the Azure documentation for details: https://docs.microsoft.com/en-us/azure/availability-zones/az-region **Note**: Ultra disks are not supported with availability sets. https://docs.microsoft.com/en-us/azure/virtual-machines/linux/disks-types
Cluster configuration	If you chose **Availability zone** earlier, specify whether to configure the cluster within a single availability zone or across two zones. You can choose among availability zones **1**, **2**, and **3**. Not all Azure regions support availability zone **3**; the only ones that do are: ▪ Central US ▪ East US ▪ East US 2 ▪ West US 2 ▪ France Central ▪ North Europe ▪ UK South ▪ West Europe ▪ Japan East ▪ Southeast Asia

First availability zone	If you chose to configure the cluster in a single availability zone, specify the zone to be used. If you chose to configure the cluster across two zones, select the first availability zone.
Second availability zone	If you chose to configure the cluster across two zones, select the second availability zone.
Availability set name	If you chose to configure the cluster in an availability set, provide the name to be used for the availability set. Note that the `-set` suffix is already provided.
Availability fault domain	If you chose to configure the cluster in an availability set, provide a value for the availability fault domain.
Availability update domain	Provide a value for the availability update domain.

8. On the **Storage** blade, provide the following details:

Configure storage	Specify whether you want to provision the data disks at the time of deployment. Data disks are used to create the InfoScale volumes or the file system on each virtual machine. If you choose to provision data disks, you can further specify the SKU type, the number of disks, and the size of the disks to be provisioned for each virtual machine. **Note**: The Premium SSD and the Ultra SSD types are not supported for all virtual machine sizes and in all regions. While selecting an SKU type, make sure that it is supported for the region, availability option, and virtual machine size that you specified. https://docs.microsoft.com/en-us/azure/virtual-machines/linux/disks-types **Note**: Remember that Premium SSD v2 disks are shared across all the VMs in a zone, by default. Thus, they can be unintentionally attached to non-InfoScale-cluster VMs, without any error or warning. To restrict the disk access to InfoScale cluster nodes and thus prevent unauthorized use, set the **Disk.MaxShares** property of the Azure disk to **1**.
Premium SSD section	
Configure Premium SSD	Specify whether you want to use the Premium SSD type of resources.
Data disks per virtual machine	If you choose to use Premium SSD, specify the number of such data disks to be provisioned for each virtual machine.
Disk size (in GB)	Specify the disk size in gigabytes. If you choose to associate multiple disks of this type, the same size is used for all the disks.
Ultra SSD section	
Configure Ultra SSD	Specify whether you want to use the Ultra SSD type of resources.
Data disks per virtual machine	If you choose to use Ultra SSD, specify the number of such data disks to be provisioned for each virtual machine.
Disk size (in GB)	Specify the disk size in gigabytes. If you choose to associate multiple disks of this type, the same size is used for all the disks.

Standard SSD section	
Configure Standard SSD	Specify whether you want to use the Standard SSD type of resources.
Data disks per virtual machine	If you choose to use Standard SSD, specify the number of such data disks to be provisioned for each virtual machine.
Disk size (in GB)	Specify the disk size in gigabytes. If you choose to associate multiple disks of this type, the same size is used for all the disks.
Standard HDD section	
Configure Standard HDD	Specify whether you want to use the Standard HDD type of resources.
Data disks per virtual machine	If you choose to use Standard HDD, specify the number of such data disks to be provisioned for each virtual machine.
Disk size (in GB)	Specify the disk size in gigabytes. If you choose to associate multiple disks of this type, the same size is used for all the disks.
Configure storage group section	
Configure FSS disk group	In case of an SFCFSHA configuration, specify whether you want to use an FSS disk group.
InfoScale disk group name	In case of an SFCFSHA configuration, provide the name to be used for the InfoScale disk group. The group is created only if the appropriate data disks are also provisioned with the virtual machines.
Configure volume group	In case of a VCS configuration, specify whether you want to create an LVM volume group.
InfoScale volume group name	In case of a VCS configuration, provide the name to be used for the LVM volume group. The group is created only if the appropriate data disks are also provisioned with the virtual machines.

9. On the **Network** blade, provide the following details:

Virtual network	Select an existing virtual network in which to deploy the InfoScale product. If you select an existing network, ensure that it contains at least three subnets and a Network Security Group, and a NAT gateway associated with it. Alternatively, create a new virtual network for this purpose.
eth0 subnet	Select the subnet to be used for eth0.
LLT link 1 subnet	Select the subnet to be used for Low Latency Transport (LLT) link 1.
LLT link 2 subnet	Select the subnet to be used for LLT link 2.
AzureIP (virtual IP)	Provide the virtual IP, which is an unused IP address from the eth0 subnet of the virtual network. When an application is configured for HA, this IP address is used to fail over the application to another cluster node within the same subnet in the same virtual network.

Network security group name	If you select an existing virtual network, provide the name of the pre-created network security group that is associated with the eth0 subnet.
NAT gateway name	If you select an existing virtual network, provide the name of the pre-created NAT gateway that is associated with the eth0 subnet.
CIDR for remote access	If you choose to create a new virtual network, provide the CIDR block from which the management server can be accessed.

10. On the **Fencing** blade, provide the following details:

Configure fencing	Specify whether you want to configure the I/O fencing option for the Infoscale cluster. When an application is configured for fencing, one or more CP servers may be used for membership arbitration among the cluster nodes. If you choose to configure fencing, you can further specify the details for the required CP servers.
Fencing details section	
Number of CP servers	Specify whether you want to configure one CP server or three. Based on your choice, you may need to provide one or more unused IP addresses from the eth0 subnet of the virtual network to be used for the servers.
Virtual machine size for CP servers	Specify the virtual machine size for the cluster nodes. Choose an appropriate size such that it supports all of the following resources and features: ▪ Type of data disk that you want to provision ▪ Accelerated networking ▪ At least three network interfaces The following virtual machine sizes are recommended: ▪ Standard_A2_v2 ▪ Standard_A4_v2 ▪ Standard_B2ms ▪ Standard_B4ms The following virtual machine sizes are allowed: ▪ Standard_D2s_v3 ▪ Standard_D4s_v3 ▪ Standard_DS2_v2 ▪ Standard_DS3_v2
VIP for CP server	If you chose to configure a single CP server, specify the virtual IP address to be used. If you chose to configure three of them, specify the virtual IP address to be used for the first CP server.
VIP for second CP server	If you chose to configure three CP servers, specify the virtual IP address to be used for the second CP server.
VIP for third CP server	If you chose to configure three CP servers, specify the virtual IP address to be used for the third CP server.

11. On the **Virtual machine** blade, provide the following details:

Image	Select the image to be used as the base operating system for the InfoScale virtual machines. Only RHEL 8.4 and SLES 15 SP2 BYOS are currently supported.
Virtual machine size	Specify the virtual machine size for the cluster nodes. Choose an appropriate size such that it supports all of the following resources and features: ▪ Type of data disk that you want to provision ▪ Accelerated networking ▪ At least three network interfaces The following virtual machine sizes are recommended: ▪ Standard_D8s_v3 ▪ Standard_D16s_v3 ▪ Standard_D32s_v3 ▪ Standard_DS3_v2 ▪ Standard_DS4_v2 ▪ Standard_DS5_v2
User name	Provide the administrator user name for the virtual machine.
Authentication type	Specify whether to use a password-based or an SSH key-based authentication.
Password	In case of password-based authentication, provide the administrator password for the virtual machine.
Confirm password	In case of password-based authentication, reenter the password for confirmation.
SSH public key	In case of SSH key-based authentication, provide the SSH public key for authentication.
Public IP address (Standard SKU) for management server	Provide the public IP address for the management server; only the Standard SKU is supported.
DNS prefix for public IP address	Provide the DNS prefix for the public IP address, which must be a globally unique value.
Diagnostic storage account	Specify the storage account to which any diagnostic data should be sent. Metrics for the virtual machine are sent to a storage account so that you can analyze them with your own tools.

12. On the **Review + create** blade, review the configuration details that are consolidated from the previous blades, and click **Create** to proceed with the deployment.

To monitor the deployment progress, navigate to the resource group in the Microsoft Azure Portal, and select **Deployment** from the navigation pane.

Step 3: Access the InfoScale cluster nodes

After successful deployment, perform the following steps to access the virtual machines that are configured as Veritas InfoScale cluster nodes.

Note: For a new virtual network, the management server is accessible only from remote access CIDR block that you provided on the Network configuration blade in the previous of deployment.

To access the InfoScale cluster nodes using a password

1. Use the appropriate password to log on to the management server first.
2. Use the same password to log on to the InfoScale cluster nodes.
3. Change the ownership of the login session to root.
   ```
   # sudo su
   ```

To access the InfoScale cluster nodes using an SSH key

1. Copy the SSH key to the management server.
   ```
   scp -i <ssh_key> <ssh_key> <user_name>@<public_ip>:/home/<user_name>
   ```
2. Log on to the management server.
   ```
   ssh -i <ssh_key> <user_name>@<public_ip>
   ```
3. From the management server log on to any of the cluster nodes:
   ```
   ssh -i <ssh_key> <user_name>@<public_ip>
   ```
4. Change the ownership of the login session to root.
   ```
   # sudo su
   ```

To view the status of the InfoScale cluster

- Run the following command:
  ```
  # /opt/VRTS/bin/hastatus -sum
  ```

To check the status of the diskgroup

- In case of an SFCFSHA configuration, an FSS diskgroup is created. To check the status of the diskgroup, run:
  ```
  # vxdg list
  # vxdisk list
  ```

Step 4: Perform the post-deployment tasks

When InfoScale Enterprise is successfully deployed using the solution template, the CVM and the Azure service group are already configured. In case of an SFCFSHA configuration, the shared FSS diskgroup is created. In case of a VCS configuration, the LVM volume group is created. You can perform further storage and HA configurations based on the needs of your application.

For general information about configuring the storage and application availability resources, refer to the InfoScale Enterprise 8.0 product documentation.

For information about the configurations that are supported on Azure, refer to the *Veritas InfoScale 8.0 Solutions in Cloud Environments* document.

Troubleshooting

This section lists some of the issues that you may encounter with a InfoScale Enterprise deployment using solution templates and the corresponding troubleshooting steps or further actions. The errors and warnings that are listed here appear in the **`/var/lib/waagent/custom-script/download/0/stdout`** file.

While deploying InfoScale Enterprise, if you run into any issues other than the ones listed here, contact Veritas Technical Support at: https://www.veritas.com/support.

Error: Few systems are not reachable even after multiple retries. Exiting...
Verify whether the management server can successfully connect to the cluster nodes by using host names over an SSH connection. If the cluster nodes cannot be accessed by using host names, contact Microsoft Azure support in case of the RHEL 8.4 platform, or contact SUSE in case of the SLES15 SP2 platform.

Error: Unable to setup communication between InfoScale cluster nodes. Exiting...
Verify whether the management server can successfully connect to the cluster nodes by using hostnames over an SSH connection. If the cluster nodes cannot be accessed by using hostnames, contact Microsoft Azure support in case of the RHEL 8.4 platform or contact SUSE in case of the SLES15 SP2 platform.

Warning: Failed to clear temporary communication between InfoScale cluster nodes
Contact Veritas Technical Support, and share the contents of the following location from each cluster node:
`/var/lib/waagent/custom-script/download/0/*`

Error: InfoScale is not configured within expected time on <virtual_machine_name>
Contact Veritas Technical Support, and share the contents of the following location from each cluster node:
`/var/lib/waagent/custom-script/download/0/*`

Error: InfoScale product configuration failed. For more details check installer logs located at /opt/VRTS/install/logs/installer-*
Contact Veritas Technical Support, and share the contents of the following locations from each cluster node:
`/var/lib/waagent/custom-script/download/0/*`
`/opt/VRTS/install/logs/installer-*`

In an SFCFSHA configuration, a diskgroup is not created or it does not include all the data disks from all the nodes

Initialize the disks that should be part of the diskgroup, if they are not already initialized, by using the command:

```
/etc/vx/bin/vxdisksetup -if <disk_name>
```

Then, add disks to the diskgroup by using the command:

```
vxdg -g <disk_group_name> adddisk <disk_name1> ... <disk_name2>
```

References

Microsoft Azure documentation

- Availability options
  https://docs.microsoft.com/en-us/azure/virtual-machines/linux/availability
- Virtual networks
  https://docs.microsoft.com/en-us/azure/virtual-machines/linux/network-overview
- Virtual machines
  https://docs.microsoft.com/en-us/azure/virtual-machines/linux/sizes
- Managed disks
  https://docs.microsoft.com/en-us/azure/virtual-machines/linux/disks-types

Veritas InfoScale documentation

To access documentation for InfoScale:

1. Access the Veritas Support site at:
   https://www.veritas.com/content/support/en_US
2. Click the **Documentation** icon.
3. Click the **Documentation** tab.
4. Select **InfoScale & Storage Foundation** from the **Product** filter in the left pane.
5. Select the appropriate platform from the **Platform** filter.
6. Select the appropriate type from the **Document Type** filter.
7. Select the appropriate version from the **Choose Version** dropdown in the right pane.
8. Alternatively, type the document title in the search bar at the top to narrow down the search results.

ABOUT VERITAS

Veritas Technologies is a global leader in data protection and availability. Over 50,000 enterprises—including 99 of the Fortune 100—rely on us to abstract IT complexity and simplify data management. Veritas Enterprise Data Services Platform automates the protection and orchestrates the recovery of data everywhere it lives, ensures 24/7 availability of business-critical applications, and provides enterprises with the insights they need to comply with evolving data regulations. With a reputation for reliability at scale and a deployment model to fit any need, Veritas supports more than 500 data sources and over 150 storage targets, including 60 clouds. Learn more at www.veritas.com. Follow us on Twitter at @veritastechllc.

2625 Augustine Drive, Santa Clara, CA 95054
+1 (866) 837 4827
veritas.com

For specific country offices
and contact numbers,
please visit our website.