



**Hewlett Packard**  
Enterprise

# HPE Nimble Storage Deployment Guide for Microsoft Azure

# Contents

<b>Overview.....</b>	<b>5</b>
Benefits of a Hybrid Cloud Architecture.....	5
<b>Cloud Characteristics.....</b>	<b>6</b>
Cloud Types and Service Models.....	6
Deployment Models.....	7
Service Models.....	7
Storage as a Service.....	8
HPE Cloud Volumes.....	8
<b>Microsoft Azure Services.....</b>	<b>9</b>
Azure Compute Services.....	9
Azure Networking Services.....	9
Azure ExpressRoute.....	9
Azure Virtual Network.....	10
<b>HPE Nimble Storage Hybrid Architecture for Azure ExpressRoute.....</b>	<b>11</b>
Deploying an Azure Virtual Network.....	12
Create and Configure a Virtual Network.....	12
Deploying Azure ExpressRoute.....	13
Create an Azure ExpressRoute Instance.....	13
Set Up a Virtual Network Gateway.....	14
Deploying an Azure Instance.....	14
Create an Azure Instance.....	14
Connect To and Configure the Azure Instance.....	16
Deploying the HPE Nimble Storage Array.....	19
Configure an Initiator Group and Deploy a Volume: Classic NimbleOS GUI.....	19
Configure an Initiator Group and Deploy a Volume: New NimbleOS GUI.....	20
Deploying a Host Data Volume.....	20
Connect the Storage Volume.....	20
Assign a Drive Letter to the Windows Volume.....	22
<b>Summary.....</b>	<b>23</b>
<b>About the Author.....</b>	<b>24</b>
Freddy Grahn.....	24

**Version History.....25**

© Copyright 2018 Hewlett Packard Enterprise Development LP. All rights reserved worldwide.

## Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

## Acknowledgments

Intel<sup>®</sup>, Itanium<sup>®</sup>, Pentium<sup>®</sup>, Intel Inside<sup>®</sup>, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft<sup>®</sup> and Windows<sup>®</sup> are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe<sup>®</sup> and Acrobat<sup>®</sup> are trademarks of Adobe Systems Incorporated. Java<sup>®</sup> and Oracle<sup>®</sup> are registered trademarks of Oracle and/or its affiliates.

UNIX<sup>®</sup> is a registered trademark of The Open Group.

## Publication Date

Monday March 19, 2018 13:51:38

## Document ID

kni1503412200233

## Support

All documentation and knowledge base articles are available on HPE InfoSight at <https://infosight.hpe.com>. To register for HPE InfoSight, click the *Create Account* link on the main page.

Email: [support@nimblestorage.com](mailto:support@nimblestorage.com)

For all other general support contact information, go to <https://www.nimblestorage.com/customer-support/>.

# Overview

The topic of cloud can be a complex discussion. There are widespread misconceptions and assumptions about what the cloud is, what it can and cannot do, and how best to adopt it.

For many years, shared infrastructure has been a popular, well-known architecture for IT. The arrival of mainstream virtualization enabled IT departments to realize the ability to share integrated resources that were traditionally overprovisioned and underutilized. However, the problem with a shared infrastructure is that resources are still manually provisioned. The cloud, at its most fundamental level, solves this problem by applying a highly automated provisioning and management paradigm to a shared infrastructure.

This guide contains an overview of key cloud concepts, a brief introduction to the Microsoft Azure cloud computing platform, and guidance for deploying HPE Nimble Storage arrays in a hybrid cloud architecture that is connected to Azure.

## Benefits of a Hybrid Cloud Architecture

The deployment of on-premises HPE Nimble Storage arrays with the Azure cloud platform provides many benefits:

- On-premises HPE Nimble Storage arrays offer superior performance and latency for critical workloads.
- Cloud compute resources can access the on-premises data without the need to migrate that data to the cloud.
- HPE InfoSight offers visibility into the data and predictive analytics.
- Remote offices or remote users can directly access cloud compute resources that access the on-premises data.
- Data volumes can be quickly and nondisruptively moved between different tiers of storage. For example, they can be moved from HPE Nimble Storage adaptive flash arrays to all-flash arrays and vice versa.
- HPE Nimble Storage enterprise-class features such as deduplication and compression for datasets can be leveraged together with the functionalities of the Azure cloud compute infrastructure.

# Cloud Characteristics

What makes a cloud, a cloud? In several ways, the popular joke that “the cloud is just someone else’s computer” is very much accurate. Realistically, however, it goes a bit beyond that.

The National Institute of Standards and Technology (NIST) has carefully defined several aspects of the cloud—characteristics, deployment models, and service models. [NIST special publication 800-145](#) describes the following features as essential cloud characteristics:

## On-demand self-service

On-demand self-service is the ability to automatically provision resources in the cloud without the need for human interaction. This characteristic is perhaps what most differentiates the cloud from a traditional shared infrastructure architecture.

## Broad network access

Broad network access refers to the ability to access cloud resources from anywhere over standard network connections, whether they are private or public in nature. Most companies access public cloud resources over a secure connection. For example, the Azure Virtual Network service is used to provide private cloud resources in a public cloud architecture.

## Resource pooling

Resource pooling corresponds to the traditional concept of shared infrastructure and secure multitenancy. It enables cloud service providers to deploy large amounts of physical resources that can be partitioned logically and securely across multiple customers. The ability to realize every bit of usage from physical resources is one of the key benefits provided by low-cost cloud services.

## Rapid elasticity

Rapid elasticity describes the ability to provision and release resources immediately, without limitations that are based on the number of available physical resources. Alternately, rapid elasticity can be described as rapid scaling of the logical and physical infrastructure. Physically, this scaling means deploying enough hardware to meet or to exceed expected customer demand. This aspect of the cloud is a key differentiator between large cloud service providers such as Amazon Web Services (AWS) and Microsoft Azure, who can afford to deploy vast amounts of physical resources, and smaller cloud service providers, who must plan physical expansion more conservatively.

## Measured service

Measured service is the ability to easily manage and monitor cloud resources. Resource manageability and simplicity have had a large resurgence in popularity in recent years. These concepts have always been important, but they are even more so today for two main reasons:

- The growing number of cloud-based services consumed by millennials, who expect services to be delivered quickly
- The pace of growth in demand for data, compute power, bandwidth, and other resources

Moreover, customers must be able to plan and predict the costs that are associated with operating workloads in the cloud. An often overlooked aspect of a measured service is that “you get what you pay for;” if you pay for x IOPS, you get x IOPS, no more and no less.

## Cloud Types and Service Models

The IT industry uses several terms to talk about cloud computing concepts. For example, clouds are usually classified into four categories:

- Public cloud
- Private cloud
- Hybrid cloud

- Community cloud

These different types of cloud fit into three common service models:

- Software as a service (SaaS)
- Platform as a service (PaaS)
- Infrastructure as a service (IaaS)

NIST provides definitions for the cloud types, which are known as deployment models, and for the cloud-service models.

## Deployment Models

*NIST special publication 800-145* defines cloud deployment models in the following way:

### Public cloud

“The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.”

### Private cloud

“The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.”

### Hybrid cloud

“The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).”

### Community cloud

“The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.”

## Service Models

*NIST special publication 800-145* defines cloud service models in the following way:

### Software as a service

“The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user specific application configuration settings.”

### Platform as a service

“The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.”

### Infrastructure as a service

“The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).”

## Storage as a Service

The commonly used terminology for cloud computing is not limited to the terms defined by NIST. Storage as a service (STaaS) is another service model that is of particular interest to Hewlett Packard Enterprise (HPE) and other storage vendors. HPE currently offers HPE Cloud Volumes, an on-demand cloud service for customers' storage needs.

The STaaS cloud service model can easily extend the capacity of existing storage arrays into a cloud service. If companies deploy storage arrays to a colocation data center to gain broad network access, they can extend those arrays to the STaaS model by addressing three key cloud characteristics:

- On-demand self-service: Building a portal to provision storage
- Rapid elasticity: Deploying enough storage to meet demand
- Measurable service: Creating mechanisms for easily managing and monitoring the storage

## HPE Cloud Volumes

HPE Cloud Volumes provide an enterprise-grade multicloud storage service for running applications in AWS and Azure. The feature is as easy to use as native cloud storage but with the enterprise-grade reliability and features that applications need.

HPE Cloud Volumes are designed for easy data mobility, so you have the freedom to move data to, from, or between public clouds and the data center without being locked in. With HPE InfoSight predictive analytics, you gain global visibility and insights across the stack, no matter where the data lives.



# Microsoft Azure Services

Microsoft Azure is one of the largest public cloud providers. It boasts over 600 services to help you deploy, develop, and manage cloud offerings. Compute, mobile, storage, and data management services are some of the main services that are offered by Azure. Each service contains several subcomponents. For example, Virtual Machines, Container Service, Disks, Snapshots, and Images are some of the services that are displayed under the main Compute section of Azure.

To connect an HPE Nimble Storage array to Azure, you need to understand a small subset of Azure services. To learn about the wide range of cloud services that are available from Azure, visit the [Microsoft Azure](#) website.

One important aspect of deploying cloud services is determining cost. Azure, like most cloud service providers, offers a [Pricing calculator](#) that can be used to estimate monthly Azure costs.

## Azure Compute Services

The core product of the Azure compute service is the virtual machine (VM) instance. You can run many kinds of compute instances:

- General purpose instances (A and D series)
- Compute optimized instances (F series)
- Memory optimized instances (D, E and G series)
- Storage optimized instances (L series)

Most instance categories contain a number of instance types, which provide a range of offerings for number of CPUs, amount of memory, type and amount of storage, and networking performance. For more information about instance types, see [Windows Virtual Machines Pricing](#).

## Azure Networking Services

Two key Azure networking services help you connect on-premises HPE Nimble Storage arrays with cloud-based compute resources:

- Azure ExpressRoute
- Azure Virtual Network

### Azure ExpressRoute

The Azure ExpressRoute service enables you to establish a dedicated network connection between the on-premises data center and Azure. A dedicated network connection provides consistent network connectivity to Azure, increased network security, and increased bandwidth.

Before requesting an ExpressRoute connection from Azure, you must find out which ExpressRoute location is nearest to your data center. ExpressRoute connections can be public, private, or both. IPsec is used to create secure tunnels between the on-premises network and the ExpressRoute location. Multiple virtual interfaces can be created to partition the connectivity. VM instances that are running within an Azure virtual network can be accessed over an ExpressRoute connection.

The cost for Azure ExpressRoute breaks down into four core items:

- **Internet service provider (ISP) connectivity costs.** Costs related to bandwidth and setup charges.
- **Billable port hours.** Measured by Azure; the number of billed hours depends on how many Azure network ports you use.
- **Data transfer.** Billed the following way: Data into the cloud is free; data out is charged (that is, data transferred into the cloud incurs no cost, but data transferred out of the cloud is billed).

- **Azure-specific resource costs.** Costs related to the particular services that you access, such as virtual instances or storage services.

For more information about ExpressRoute, see [ExpressRoute overview](#). For information about ExpressRoute costs, see [ExpressRoute pricing](#).

## Azure Virtual Network

The Azure Virtual Network (VNet) service enables you to create a private virtual network within Azure. With VNet, you have control over IP address ranges, gateways, subnets, and routing tables for connections. If you combine a VPN connection with ExpressRoute, you can deploy the VPN to connect a virtual network to an on-premises data center network. You can further secure VPN connections by encrypting them.

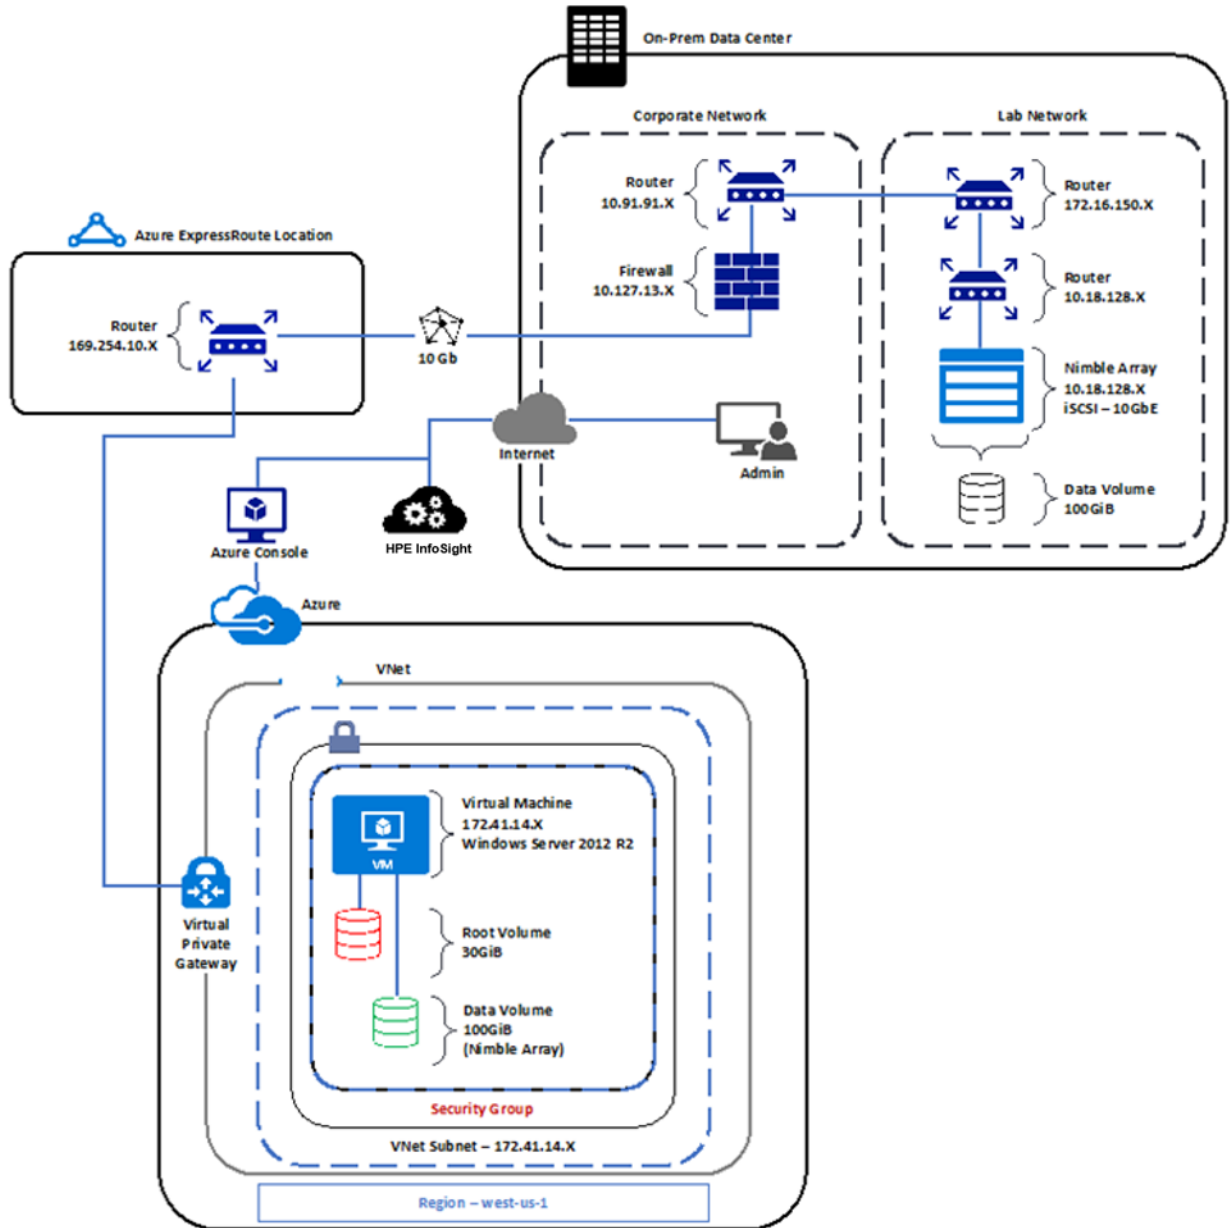
A virtual network can be peered to another virtual network in Azure to bridge the networks together. You can attach elastic IP addresses (public IP addresses) to a network interface to make an Azure instance publicly accessible over the Internet. However, although it is possible to attach public IP addresses to a specific instance, doing so creates a security risk for VNet, which is meant to provide a private network in Azure. Exercise caution when assigning public IP addresses to Azure resources.

For more information about Azure VNet, see [Virtual Network](#).

# HPE Nimble Storage Hybrid Architecture for Azure ExpressRoute

The following diagram illustrates the storage architecture for a solution that combines Azure ExpressRoute and HPE Nimble Storage arrays. The on-premises data center in this architecture must traverse multiple internal networks. An essential characteristic of the architecture is that the storage array discovery IP address and data interfaces are routable to the Azure ExpressRoute path.

**Figure 1: Azure ExpressRoute storage architecture**



To test this reference architecture, HPE provisioned a 10 Gbps connection for the solution. Although you can provision connections at speeds lower than 10 Gbps through Azure partners, the two available speeds for

initiating an ExpressRoute connection from the Azure portal are up to 10 Gbps. You can provision one or more connections at whatever bandwidth rates are necessary to support the data access requirements.

When planning the design for connectivity between on-premises HPE Nimble Storage arrays and Azure compute resources through ExpressRoute, use this reference storage architecture as guidance. Although your requirements and network connectivity might vary slightly from the example architecture, the basic architectural foundation remains common across installations. Nevertheless, ExpressRoute does support other connectivity designs.

**Note** The latest release of HPE Cloud Volumes automates the process of creating a connection between an Azure VM and an HPE Nimble Storage array to simplify your cloud experience.

## Deploying an Azure Virtual Network

Deploying an Azure instance into a virtual network offers several benefits. For example:

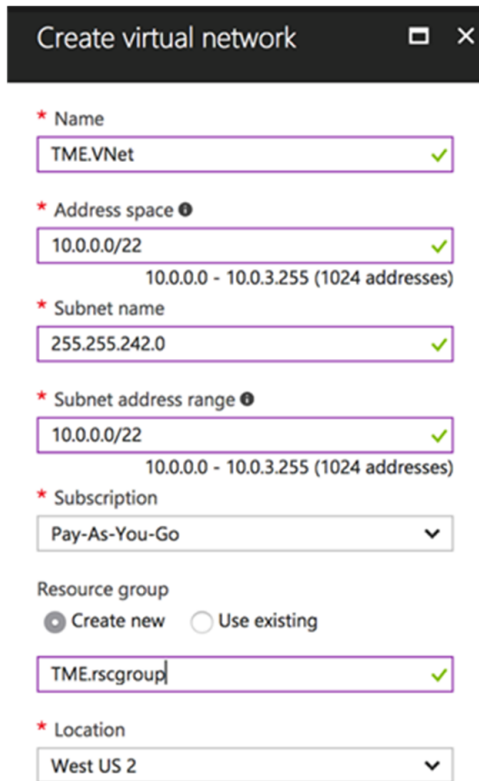
- The instance can use a persistent, private IP address.
- The instance can use security groups to filter inbound and outbound traffic.
- The instance can use network ACLs to filter traffic for the entire virtual network.

You can set up a virtual network in many different ways. For more information about the available options, see [Azure Virtual Network](#).

## Create and Configure a Virtual Network

### Procedure

- 1 Log in to the [Azure](#) portal.
- 2 In the left pane, click **Virtual networks**.
- 3 Click **Add** to create a virtual network.
- 4 Enter a name for the virtual network.
- 5 Select a prefix mask of /22 or lower for the address space.
- 6 Add the subnet name.
- 7 Set the subnet address range.
- 8 Select the correct subscription.
- 9 Create a new resource group or select an existing one.
- 10 Select West US 2 or East US as the location for the virtual network.



**Create virtual network**

\* Name  
TME.VNet ✓

\* Address space ⓘ  
10.0.0.0/22 ✓  
10.0.0.0 - 10.0.3.255 (1024 addresses)

\* Subnet name  
255.255.242.0 ✓

\* Subnet address range ⓘ  
10.0.0.0/22 ✓  
10.0.0.0 - 10.0.3.255 (1024 addresses)

\* Subscription  
Pay-As-You-Go ▼

Resource group  
 Create new  Use existing

TME.rscgroup ✓

\* Location  
West US 2 ▼

#### 11 Click **Create**.

At this point, the virtual network is created and has a single subnet.

## Deploying Azure ExpressRoute

The process for deploying an ExpressRoute circuit is covered in detail in the [Azure ExpressRoute User Guide](#). ExpressRoute enables you to extend on-premises networks into the Azure cloud over a private connection. ExpressRoute connections provide multiple benefits:

- Faster speeds
- Lower latencies
- Higher security
- More reliability

### Create an Azure ExpressRoute Instance

#### Procedure

- 1 Log in to the Azure portal.
- 2 At the top of the left pane, click **New**.
- 3 Select **Networking** > **ExpressRoute** to create an ExpressRoute circuit.
- 4 Enter a name for the ExpressRoute circuit.
- 5 Select an ISP (for example, Equinix).
- 6 Select a peering location where the ISP will provision the connection.

Choosing the best provider and associated location is critical to obtaining a good performance experience.

- 7 Select the required bandwidth for the provisioned network.
- 8 Select the correct SKU tier.
- 9 Select the billing model.
- 10 Click **Create**.

When you click **Create**, charges start processing immediately.

## Set Up a Virtual Network Gateway

### Procedure

- 1 At the bottom of the left pane of the Azure portal, click **More Services**.
- 2 Search for `Virtual network gateways` and select it in the search results list.
- 3 Click **Add** to create a virtual network gateway.
- 4 Enter a name for the gateway.
- 5 Select **ExpressRoute** as the gateway type.
- 6 From the SKU list, select a performance setting of either **Standard** (1000 Mbps), **High performance** (2000 Mbps), or **Ultra performance** (9000 Mbps).
- 7 Click **Virtual network** and select the virtual network that you previously created.
- 8 Click **Public IP address** and select an existing address or create a new address.
- 9 Select the correct subscription.
- 10 If required (as determined by your previous selections), select the location for the gateway.
- 11 Click **Create**.

Azure might take up to 20 minutes to create the virtual network gateway.

## Deploying an Azure Instance

Deploying and configuring an Azure instance involves two tasks:

- Creating an Azure compute instance
- Connecting to the new instance and configuring it

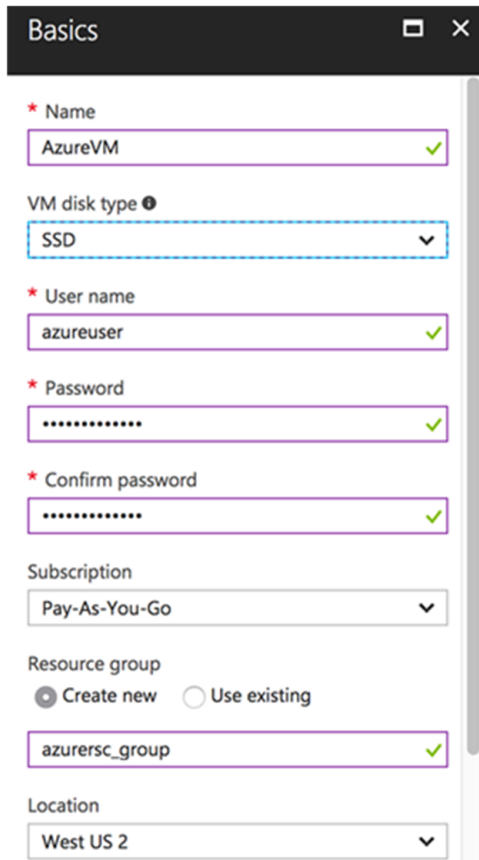
### Create an Azure Instance

#### Procedure

- 1 Log in to the Azure portal.
- 2 At the top of the left pane, click **New**.
- 3 Select **Compute** and then select the type of compute instance that you want to deploy.
- 4 In the Basics blade, set the properties of the new instance:
  - a Enter a name for the instance.
  - b Select a VM disk type (SSD or HDD).
  - c Provide a user name.

Do not use `administrator` as the user name.
  - d Create a password.
  - e Select a subscription setting.

In the example, **Pay-As-You-Go** was selected.
  - f Create a new resource group or select an existing one.
  - g Select the location where the instance will run.



The screenshot shows the 'Basics' configuration window for an Azure VM. The fields and their values are:

- Name: AzureVM
- VM disk type: SSD
- User name: azureuser
- Password: [Redacted]
- Confirm password: [Redacted]
- Subscription: Pay-As-You-Go
- Resource group: azurersc\_group
- Location: West US 2

5 Click **OK**.

6 In the Size blade, choose the instance size that matches your requirements and click **Select**.

For example, the DS1\_V2 Standard size has the following settings:

- 1 core
- 3.5 GB:
  - 2 data disks
  - 3200 maximum IOPS
  - 7 GB local SSD
  - Load balancing

7 In the Settings blade, configure storage, networking, and monitoring options for the instance:

a Select **No** for storage.

This instance will use an HPE Nimble Storage volume that you will create later.

b Select the virtual network that you previously created.

c Select the appropriate subnet.

d Select a public IP address.

e Select a network security group.

f Select any extensions, if applicable.

g Choose whether to create an availability set.

h Select the appropriate monitoring setting.

8 Click **OK**.

After you click **OK**, Azure validates your configuration.

- 9 The last page completes a validation against the configuration you have chosen. When the validation is complete, click **Purchase** at the bottom of the page.

## Connect To and Configure the Azure Instance

This procedure describes the process for connecting to the Azure instance from a Windows based system through Microsoft Remote Desktop Connection. Mac and Linux users should connect to the instance by using the appropriate local OS tools.

### Procedure

- 1 In the Azure portal, open the instances dashboard.
- 2 Locate the instance that you created and select it.
- 3 When the details for the instance are displayed, locate the public IP address.
- 4 Using Microsoft Remote Desktop Protocol (RDP), connect to the instance.

This step requires the user name and password that you provided when you created the instance.

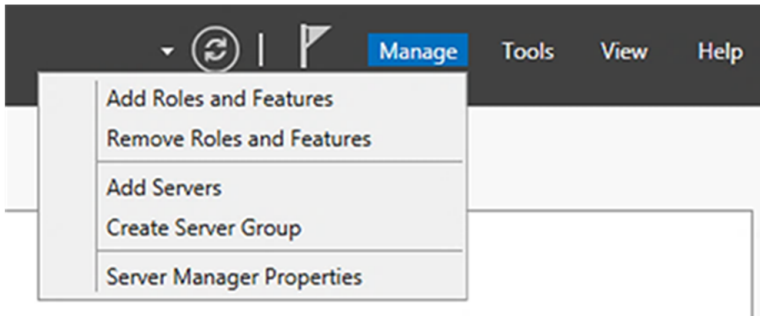
- a Open the Remote Desktop Connection application on the local host.
- b Enter the appropriate IP address in the Computer field and click **Connect**.
- c In the Windows Security dialog box, enter `administrator` for the user name and use the password that was generated.

The system will likely prompt you to accept a certificate before connecting to the remote system.

- d If you are prompted to accept a security certificate, select the **Don't ask me again for connections to this computer** checkbox and click **Yes**.

The remote desktop is displayed.

- 5 To install multipath I/O (MPIO) on the system, open the Windows Start menu and select **Server Manager**.
- 6 In Server Manager, select **Manage > Add Roles and Features**.



- 7 On the Before You Begin page of the Add Roles and Features wizard, click **Next**.
- 8 On the Installation Type page, keep the selection for **Role-based or feature-based installation**. Click **Next**.
- 9 On the Server Selection page, verify that **Select a server from the server pool** is selected and that your Windows server is highlighted in the Server Pool list. Click **Next**.
- 10 On the Server Roles page, click **Next**.
- 11 On the Features page, scroll down the Features list and select **Multipath I/O**. Click **Next**.
- 12 On the Confirmation page, click **Install**.

Wait for the installation to be completed.

- 13 On the Results page, click **Close**.



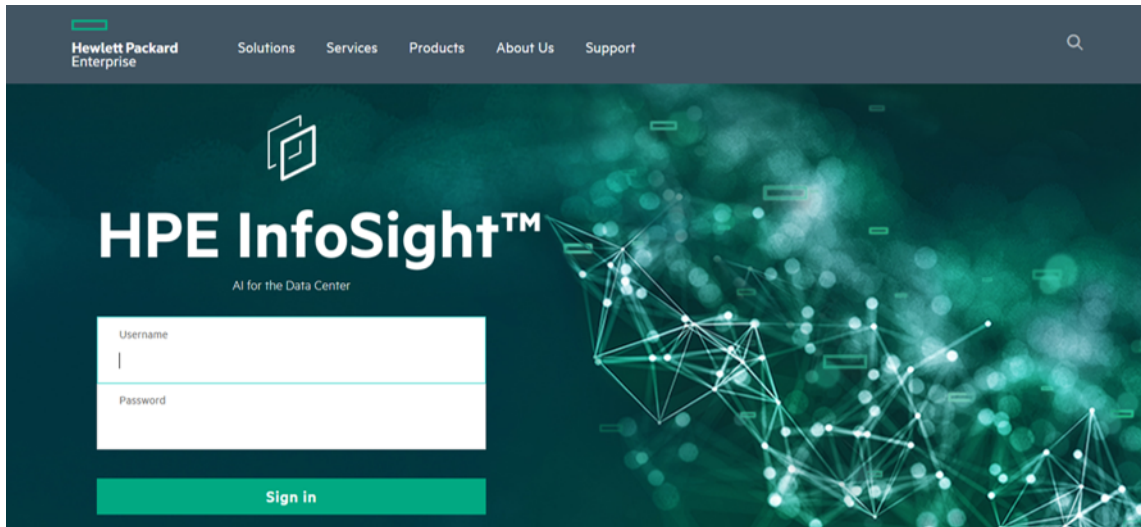
14 Close the Server Manager window.

15 Open Microsoft Internet Explorer and navigate to HPE InfoSight.

**Note** Internet security is set to high by default. To navigate to HPE InfoSight, add the site to the list of trusted sites or lower the Internet security settings.

16 Enter your credentials to log in to HPE InfoSight.

**Note** If you do not have HPE InfoSight credentials, click the **New user? Enroll now** link and register for an account before returning to this step.



17 In the main HPE InfoSight dashboard, select **Resources > Software Downloads**.

18 In the left pane, under Integration Kits, click **HPE Nimble Storage Windows Toolkit (NWT)**.

19 Click **Show other versions** on the right side of the Current Version title.

20 Locate the 3.2.0.410 (GA) version entry and click **Software (64-bit)** to download the toolkit.

**Note** Download a version of the toolkit that is the same as or higher than the version of NimbleOS that is installed on the HPE Nimble Storage array. Although this example uses 3.2.0.410 (GA), you can use any supported version that is appropriate for your configuration.

21 After you download the toolkit, launch the installation file.

22 On the Welcome page of the installation wizard, click **Next**.

23 Read and accept the license agreement. Click **Next**.

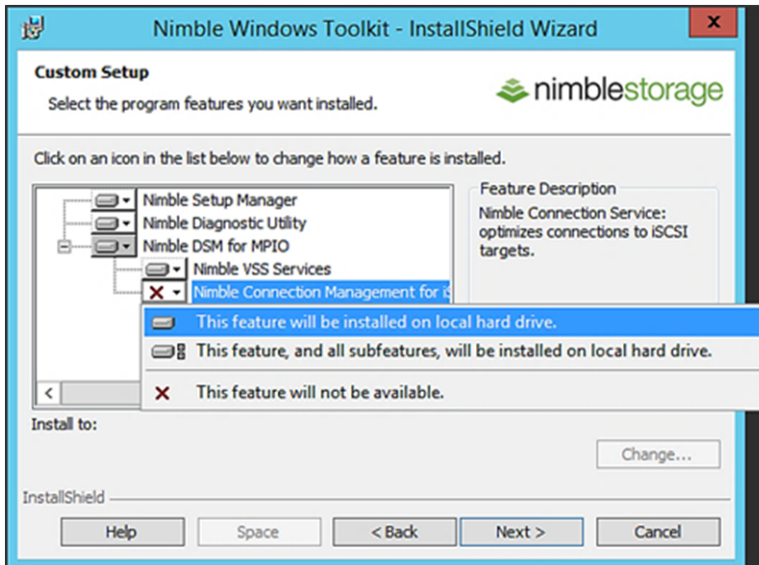
24 Verify that all critical Microsoft hotfixes have been applied to the system and click **Next**.

**Important** You must resolve all hotfix gaps before you continue with the installation. To generate a report that identifies these gaps, click **Hotfix Report**.



25 On the Nimble Logs Directory page, click **Next**.

26 On the Custom Setup page, click **Nimble Connection Management for iSCSI** in the features list and select **This feature will be installed on local hard drive**. Click **Next**.



27 On the Ready to Install the Program page, click **Install**.

The wizard displays a message that lists the registry key values that the installation will adjust.

28 Click **OK** in the message.

29 When the installation is complete, click **Finish**.

A system restart is necessary to finalize the installation process.

30 Click **Yes** to restart the system.

31 When the restart is complete, open the Windows Start menu and locate the HPE Nimble Storage applications.

32 Open the **Nimble Connection Manager** application.

**33** Record the IQN string that is displayed in the Initiator Name field.

You will need the IQN information to configure the HPE Nimble Storage array.

## Deploying the HPE Nimble Storage Array

Instructions for the installation and basic configuration of an HPE Nimble Storage array are outside the scope of this guide. For those instructions, see the *Administration Guide* for the NimbleOS version that is running on your array. The documentation is available on [HPE InfoSight](#).

When the basic setup and configuration of the array is complete, proceed to the array configuration tasks that are specific to deploying a storage volume for the Azure instance.

NimbleOS 3.4 and later releases run a newer version of the NimbleOS GUI. NimbleOS releases earlier than 3.4 run the classic GUI. The steps for configuring the array differ slightly between the two interfaces.

### Configure an Initiator Group and Deploy a Volume: Classic NimbleOS GUI

This procedure applies to releases of NimbleOS earlier than release 3.4.

#### Procedure

- 1 Open a web browser and navigate to the management IP address of the HPE Nimble Storage array.
- 2 Enter your user name and password and click **Log In**.
- 3 On the array dashboard, select **Manage > Initiator Groups**.
- 4 On the Initiator Groups page, click **Create**.
- 5 Enter a name for the initiator group (for example, `Azure.TME.VM`).
- 6 Verify that **Use all configured subnets** is selected.

**Note** A single initiator group can be used for all IQNs that are associated with Azure instances.

- 7 In the Initiators pane, set the general properties of the initiator group:
  - a Give the IQN a descriptive name that refers to the Azure instance (for example, `VM.ME.M4`).
  - b Enter the IQN string that you recorded from the Azure instance (for example, `iqn.1991-05.com.microsoft:win-9jtu9s7f4v2`).
  - c Enter the private IP address of the Azure VM (for example, `172.41.14.218`).
  - d Click **Create**.
- 8 Select **Manage > Volumes**.
- 9 On the Volumes dashboard, click **New Volume**.
- 10 On the General page of the Create Volume wizard, set the general properties of the volume:
  - a Enter a volume name (for example, `Azure.TME.VM.M4`).
  - b Select the **Windows File Server** performance policy.
  - c Select the iSCSI initiator group that you previously created (for example, **Azure.TME.VM**).
  - d Click **Next**.
- 11 On the Space page, configure the size of the volume:
  - a Enter a numeric value for the Size field (for example, 1).
  - b Set the unit of measurement for capacity (for example, TiB).
  - c Click **Next**.
- 12 On the Protection page, select a data protection policy.

If the volume is not in a replication relationship, select **No volume collection**.
- 13 Click **Finish**.

## Configure an Initiator Group and Deploy a Volume: New NimbleOS GUI

This procedure applies to NimbleOS 3.4 and later releases.

### Procedure

- 1 Open a web browser and navigate to the management IP address of the HPE Nimble Storage array.
- 2 Enter your user name and password and click **Log In**.
- 3 On the array dashboard, select **Manage > Data Access**.
- 4 On the Data Access page, click the plus sign to add an initiator group.
- 5 Enter a name for the initiator group (for example, `Azure.TME.VM`).
- 6 Verify that **Use all configured subnets** is selected.

**Note** A single initiator group can be used for all IQNs that are associated with Azure instances.

- 7 In the Initiators pane, click **Add**.
- 8 Set the general properties of the initiator group:
  - a Give the IQN a descriptive name that refers to the Azure instance (for example, `VM.ME.M4`).
  - b Enter the IQN string that you recorded from the Azure instance (for example, `iqn.1991-05.com.microsoft:win-9jtu9s7f4v2`).
  - c Enter the private IP address of the Azure VM (for example, `172.41.14.218`).
  - d Click **Create**.
- 9 Select **Manage > Data Storage**.
- 10 On the Data Storage page, click the plus sign to add a new volume.
- 11 In the Create Volume dialog box, set the general properties of the volume:
  - a Enter a volume name (for example, `Azure.TME.VM.M4`).
  - b Select the **Windows File Server** performance policy.
  - c Enter a numeric value for the Size field (for example, `1`).
  - d Set the unit of measurement for capacity (for example, `TiB`).
  - e Select a data protection policy (for example, **Not protected**).
  - f In the Access field, select the initiator group that you previously created (for example, **Azure.TME.VM**).
- 12 Click **Create**.

## Deploying a Host Data Volume

After you prepare the Windows environment and configure the HPE Nimble Storage array, you can configure the host volume, which involves two tasks:

- Connecting the storage volume to the host
- Assigning a drive letter to the volume

## Connect the Storage Volume

### Procedure

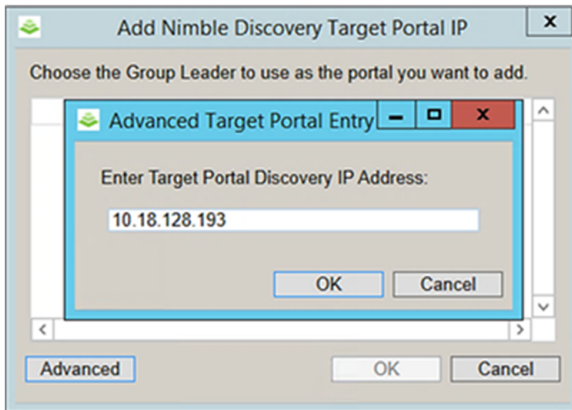
- 1 Open the Remote Desktop Connection application on the local host.
- 2 Enter the appropriate IP address in the Computer field and click **Connect**.
- 3 In the Windows Security dialog box, enter `administrator` as the user name and use the password that was generated.

The system will likely prompt you to accept a certificate before connecting to the remote system.

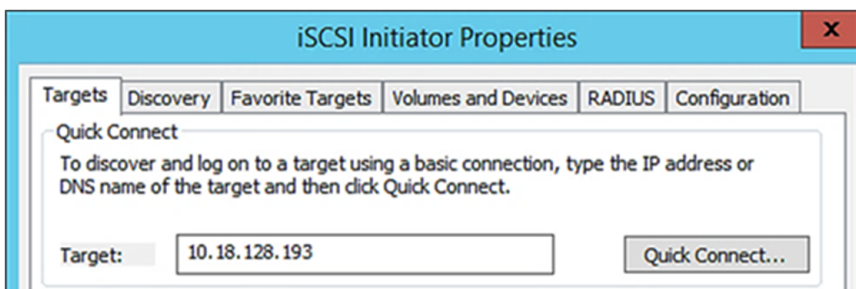
- 4 If you are prompted to accept a security certificate, select the **Don't ask me again for connections to this computer** checkbox and click **Yes**.

The remote desktop is displayed.

- 5 If HPE Nimble Storage Connection Manager (NCM) is not already running, open the Windows Start menu and locate the HPE Nimble Storage applications.
- 6 Open the **Nimble Connection Manager** application.
- 7 In the main window of NCM, in the Nimble Discovery IP (Port 3260) pane, click **Add**.
- 8 In the Add Nimble Discovery Target Portal IP dialog box, click **Advanced**.
- 9 In the Advanced Target Portal Entry dialog box, enter the iSCSI discovery IP address of the HPE Nimble Storage array (for example, 10.18.128.193). Click **OK**.



- 10 Back in the Add Nimble Discovery Target Portal IP dialog box, click **OK**.
- 11 In the main window of NCM, verify that the discovery IP address has been added to the Nimble Discovery IP (Port 3260) pane.
- 12 Keep NCM open.
- 13 Open the Windows Start menu and enter `iSCSI initiator` in the search field.
- 14 Click the iSCSI initiator entry from the search results list.
- 15 On the Targets tab of the iSCSI Initiator Properties dialog box, enter the iSCSI discovery IP address of the HPE Nimble Storage array in the Target field (for example, 10.18.128.193).
- 16 Click **Quick Connect**.



- 17 In the Quick Connect dialog box, verify that the correct IQN is listed with a status of Connected. Click **Done**.
- 18 Back at the iSCSI Initiator Properties dialog box, in the Discovered Targets pane, select the IQN that is associated with the HPE Nimble Storage array. While the IQN is selected, click **Connect**.
- 19 In the Connect to Target dialog box, select the **Enable multi-path** checkbox and click **OK**.
- 20 Close the iSCSI Initiator Properties dialog box.
- 21 Go back to NCM (which you should have left open from step 12).

**22** Click the **Nimble Volumes** tab.

**23** Verify that the volume is listed as connected and healthy in the Discovered Nimble Volumes pane.

## Assign a Drive Letter to the Windows Volume

### Procedure

- 1 Open the Windows Start menu and enter `disk management` in the search field.
- 2 Select **Create and format hard disk partitions** from the search results list.
- 3 In the main window of the Disk Management application, locate an unknown disk entry that matches the size of the volume that you created on the HPE Nimble Storage array.

For example, there should be an entry with settings such as Disk 1, Unknown, 1024.00 GB, Offline, Unallocated.

- 4 Right-click the unknown disk and select **Online**.
- 5 Right-click the unknown disk again and select **Initialize Disk**.
- 6 In the Initialize Disk dialog box, click **OK**.

In response, the disk entry should show as online and basic, but still unallocated.

- 7 Right-click the **Unallocated** section of the disk entry and select **New Simple Volume**.
- 8 On the Welcome page of the New Simple Volume wizard, click **Next**.
- 9 On the Specify Volume Size page, verify that the maximum size is selected. Click **Next**.
- 10 On the Assign Drive Letter or Path page, verify that **Assign the following drive letter** is selected and that the drop-down list shows an appropriate drive letter (for example, D). Click **Next**.
- 11 On the Format Partition page, enter a name for the volume label (for example, `TME.Instance.M4`). Click **Next**.
- 12 On the Complete page, click **Finish**.

The disk entry should show the status of the disk as Formatting. When the formatting is complete, the disk status changes to Healthy (Primary Partition).

- 13 Open Windows File Explorer and verify that the new volume is listed and accessible.

**Note** To further verify that the volume is working correctly, try copying or creating files in the new volume.

## Summary


The cloud is not the answer to everything (at least not yet), but it is clear that the new reality of IT revolves around leveraging on-premises and cloud architectures together in a hybrid model. Hybrid cloud architectures that span the public cloud, the private cloud (on-premises or hosted), and on-premises resources deliver the most realistic and robust technology capabilities for businesses today.

There is no doubt that you can gain significant benefits from deploying on-premises or colocated HPE Nimble Storage arrays in support of Azure cloud compute resources. Such an architecture enables you to take advantage of the enterprise features, availability, accessibility, and scalability of best-in-class HPE Nimble Storage solutions to retain control over one of the most—if not *the* most—critical asset that your company has: your data.

This guide will be updated continually to reflect the latest guidance and recommendations for deploying HPE Nimble Storage arrays in conjunction with the Azure platform. For more information about Azure features, consult the multitude of online documents available on the [Microsoft Azure](#) website. For additional information about HPE Nimble Storage technologies, visit the documentation pages on [HPE InfoSight](#).

## About the Author

### Freddy Grahn



*Senior Technical Marketing Engineer, Product Management*  
Hewlett Packard Enterprise

Freddy Grahn has over 20 years of experience working with enterprise databases and over 10 years of experience in the storage industry. He has written many Oracle performance papers and supported hundreds of customers in deploying Oracle solutions. Freddy is known around the office as “The Oracle,” but he has yet to predict the future.



## Version History

Version	Release Date	Description
1.1	March 2018	Terminology updates
1.0	August 2017	Initial release