



**Build Scale out storage system
by Windows Server 2019
Azure Stack HCI and
QuantaPlex T21P-4U
Deployment Guide**

Written by Bono Hsueh

Last Update: 3/23/2020

Version: 1.3



Abstract

Azure Stack HCI and Windows Server Software Defined (WSSD) Program are both invitation-only program by Microsoft in which solution providers design hyper-converged infrastructures with Windows Server technologies.

After going through Microsoft's validation and requirements, all the systems are pre-configured with certified components, tested and certified by QCT.

Microsoft Cloud Ready Appliances integrate all of Microsoft Windows Server 2019 SDDC technologies—software-defined compute, storage and networking virtualization and security—into a hyper converged appliances and disaggregated appliances. This optimized modular appliance is easy to use, agile, scalable and cloud-ready platform from four nodes to sixteen nodes in the same cluster to handle a variety of virtualization workloads and to address current needs while preparing for future requirements.

REVISIONS

Version	Date	Description	Authors
1.0	10/08/2019	First Publish	Bono Hsueh
1.1	10/17/2019	Add Materials	Bono Hsueh
1.2	10/24/2019	Add Materials	Bono Hsueh
1.3	10/30/2019	Add Materials	Bono Hsueh



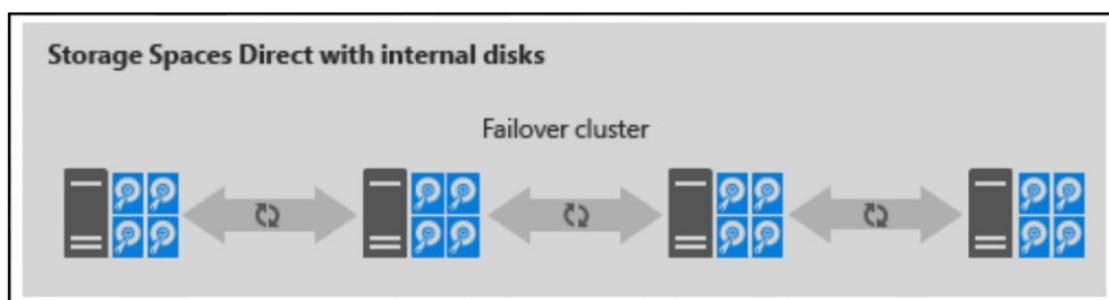
CONTENTS

Abstract	1
REVISIONS	2
CONTENTS	3
Windows Server 2016/2019 Storage Spaces Direct (S2D)	4
Introduction.....	4
Storage Space Direct (S2D) Key benefits	5
S2D Use Cases.....	6
How it works.....	8
QCT Deployment Lab	12
1. Hardware for this Guide	12
2. Server Nodes / Network Diagram	13
3. Prepare the QCT physical server node	14
4. Install Windows Server 2019 OS	21
5. Install Hardware Drivers and Windows Roles/ Features	26
6. Create Failover Cluster and Enable S2D	30
Scale-out storage solution	34
1. Deploy Windows Admin Center and QCT Management Suits	34
2. Integrated with hybrid storage management by WAC	45
3. Config Azure Backup in Windows Admin Center	49
4. Config Azure File Sync in Windows Admin Center	52
About QCT	53

Windows Server 2016/2019 Storage Spaces Direct (S2D)

Introduction

Windows Server 2016 introduced Storage Spaces Direct (S2D), which enables building highly available storage systems which is virtual shared storage across the servers using local disks. This is a significant step forward in Microsoft's Windows Server Software-defined Storage (SDS) as it simplifies the deployment and management of storage and also unlocks the use of new classes of disk devices, such as SATA and NVMe disk devices, that were previously not possible with clustered Storage Spaces with shared disks. With Windows Server 2016/2019 S2D, you can now build HA Storage Systems using storage nodes with only local storage that consists of internal disk devices.



Storage Spaces Direct uses industry-standard servers with local-attached drives to create highly available, highly scalable software-defined storage at a fraction of the cost of traditional SAN or NAS arrays. Its converged or hyper-converged architecture radically simplifies procurement and deployment, while features like caching, storage tiers, and erasure coding, together with the latest hardware innovation like RDMA networking and NVMe drives, deliver unrivaled efficiency and performance. Storage Spaces Direct is included in Windows Server 2016/2019 Datacenter Edition.

Storage Space Direct (S2D) Key benefits



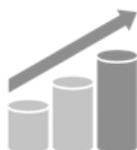
Simplicity. Go from industry-standard servers running Windows Server 2016 to your first Storage Spaces Direct cluster in under 15 minutes. For System Center users, deployment is just [one checkbox](#).



Unrivalled Performance. Whether it being all-flash or hybrid, Storage Spaces Direct easily exceeds [150,000 mixed 4k random IOPS per server](#) with consistency and low latency thanks to its hypervisor-embedded architecture, its built-in read/write cache, and support for cutting-edge NVMe drives mounted directly on the PCIe bus.



Fault Tolerance. Built-in resiliency handles drive, server, or component failures with continuous availability. Larger deployments can also be configured for [chassis and rack fault tolerance](#). When hardware fails, just swap it out; the software heals itself, with no complicated management steps.



Resource Efficiency. Erasure coding delivers up to 2.4x greater storage efficiency, with unique innovations like [Local Reconstruction Codes](#) and [real-time tiering](#) to extend these gains to hard disk drives and mixed hot/cold workloads, all while minimizing CPU consumption to give resources back to where they're needed most - the VMs.



Manageability. Use [Storage QoS Controls](#) to keep overly busy VMs in check with minimum and maximum per-VM IOPS limits. The [Health Service](#) provides continuous built-in monitoring and alerting, and new APIs make it easy to collect rich, cluster-wide performance and capacity metrics.



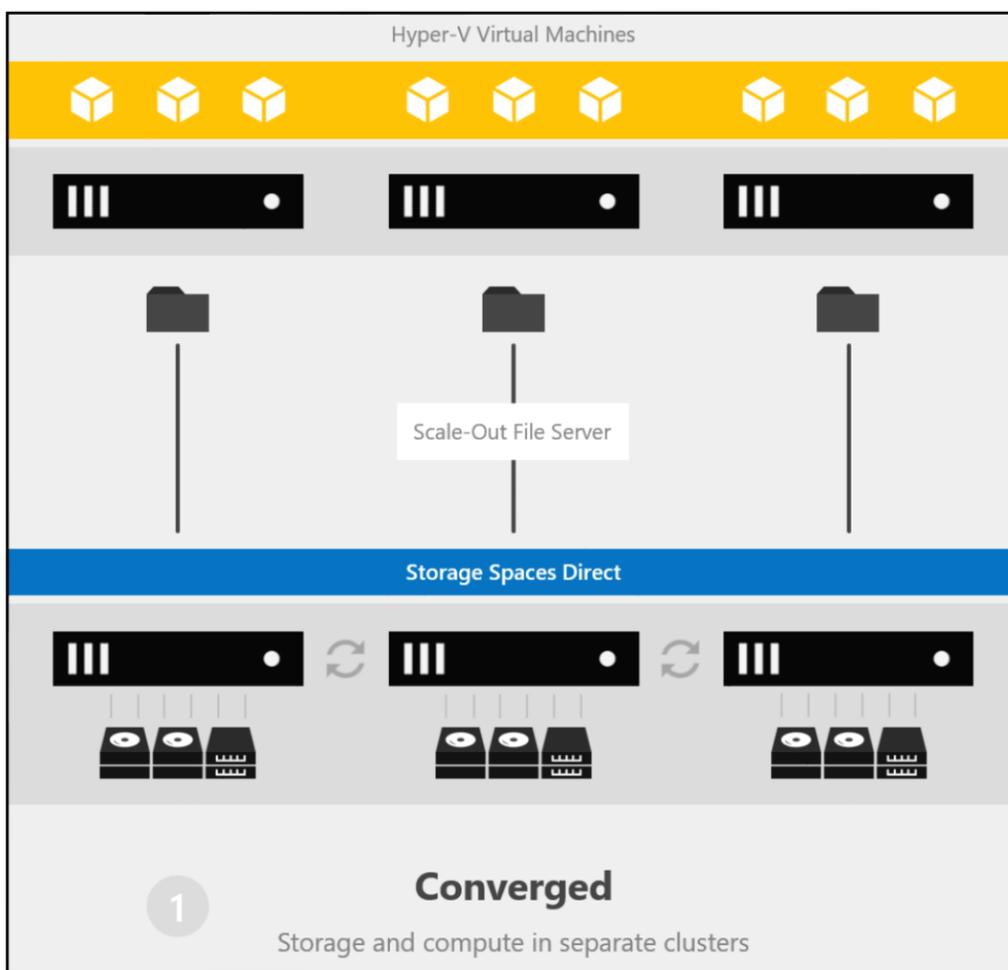
Scalability. Go up to 16 servers and over 400 drives, for multiple petabytes of storage per cluster. To scale out, simply add drives or add more servers; Storage Spaces Direct will automatically onboard new drives and begin using them. Storage efficiency and performance improve predictably at scale.

S2D Use Cases

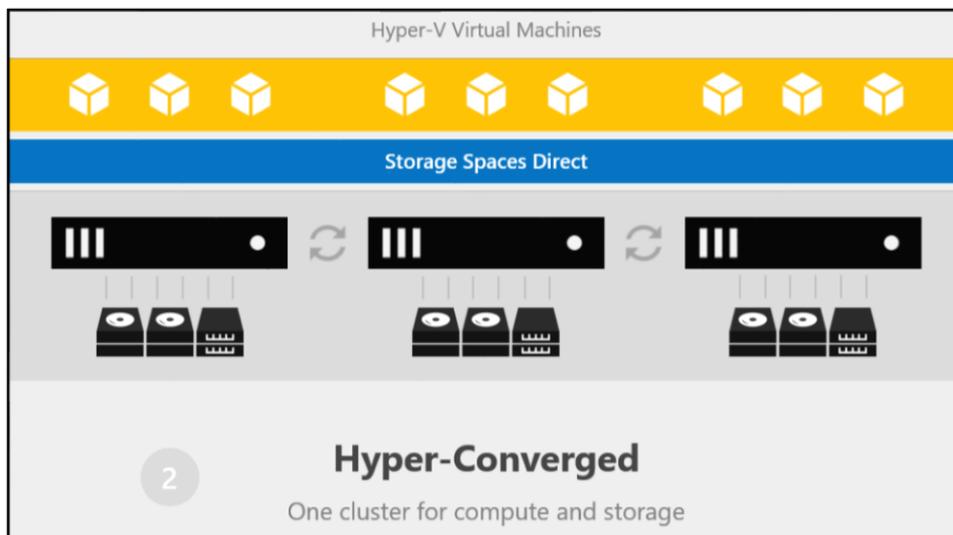
There are two targeted deployment scenarios for Windows Server 2016/2019 Storage Spaces Direct: **Hyper-converged & Disaggregated**. Both cases provide storage for Hyper-V, specifically focusing on Hyper-V IaaS (Infrastructure as a Service) for Service Providers and Enterprises.

1. The **converged** deployment option, also known as **disaggregated**, layers a Scale-out File Server (SoFS) atop Storage Spaces Direct to provide network-attached storage over SMB3 file shares. This allows for scaling compute/workload independently from the storage cluster, essential for larger-scale deployments such as Hyper-V IaaS (Infrastructure as a Service) for service providers and enterprises.

The **disaggregated** deployment scenario has the Hyper-V servers (compute component) in a separate cluster from the Storage Spaces Direct servers (storage component).



2. The **hyper-converged** deployment option runs Hyper-V virtual machines or SQL Server databases directly on the servers providing the storage, storing their files on the local volumes. This eliminates the need to configure file server access and permissions, and reduces hardware costs for small-to-medium business or remote office/branch office deployments.



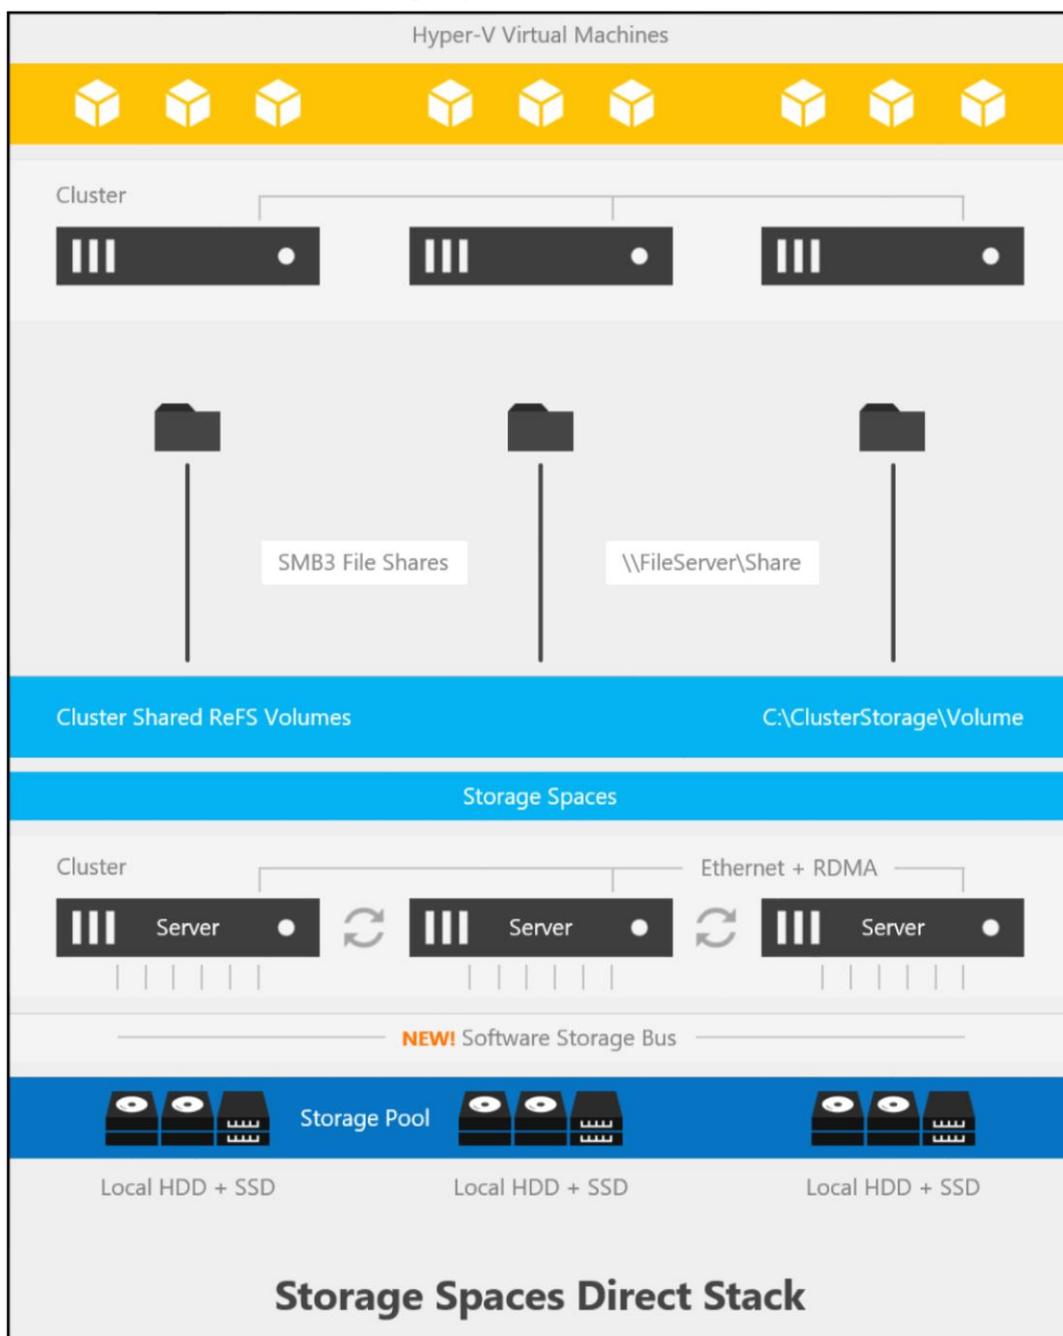
When deciding which option works best for your situation, remember that hyper-converged has less complexity because of the unified compute and storage, while converged provides more flexibility to scale storage and compute independent of one another.

Large scale infrastructures will typically use converged to allow for the independent scaling, whereas smaller systems can benefit from the simplicity of hyper-converged.

How it works

Storage Spaces Direct is the evolution of Storage Spaces, firstly introduced in Windows Server 2012. It leverages many of the features you have known today in Windows Server, such as Failover Clustering, the Cluster Shared Volume (CSV) file system, Server Message Block (SMB) 3, and of course Storage Spaces. It also introduces new technology, most notably the Software Storage Bus.

Here's an overview of the Storage Spaces Direct stack:





Networking Hardware. Storage Spaces Direct uses SMB3, including SMB Direct and SMB Multichannel, over Ethernet to communicate between servers. We strongly recommend 10+ GbE with remote-direct memory access (RDMA), either iWARP or RoCE.

Storage Hardware. From 2 to 16 servers with local-attached SATA, SAS, or NVMe drives. Each server must have at least 2 solid-state drives, and at least 4 additional drives. The SATA and SAS devices should be behind a host-bus adapter (HBA) and SAS expander.

Failover Clustering. The built-in clustering feature of Windows Server is used to connect the servers.

Software Storage Bus. The Software Storage Bus is new in Storage Spaces Direct. It spans the cluster and establishes a software-defined storage fabric whereby all the servers can see all of each other's local drives.

Storage Bus Layer Cache. The Software Storage Bus dynamically binds the fastest drives present (e.g. SSD) to slower drives (e.g. HDDs) to provide server-side read/write caching that accelerates IO and boosts throughput.

Storage Pool. The collection of drives that will form the basis of Storage Spaces is called the storage pool. It is automatically created, and all eligible drives are automatically discovered and added to it. We strongly recommend you use one pool per cluster, with the default settings.

Storage Spaces. Storage Spaces provides fault tolerance to virtual "disks" using [mirroring, erasure coding, or both](#). You can think of it as distributed, software-defined RAID using the drives in the pool. In Storage Spaces Direct, these virtual disks typically have resiliency to two simultaneous drive or server failures.

Resilient File System (ReFS). ReFS is the premier filesystem purposely-built for virtualization. It includes dramatic accelerations for .vhdx file operations such as creation, expansion, and checkpoint merging, and built-in checksums to detect and correct bit errors. It also introduces real-time tiers that rotate data between so-called "hot" and "cold" storage tiers in real-time based on usage.

Cluster Shared Volumes. The CSV file system unifies all the ReFS volumes into a single namespace accessible through any server, so that to each server, every volume looks and acts like it's mounted locally.

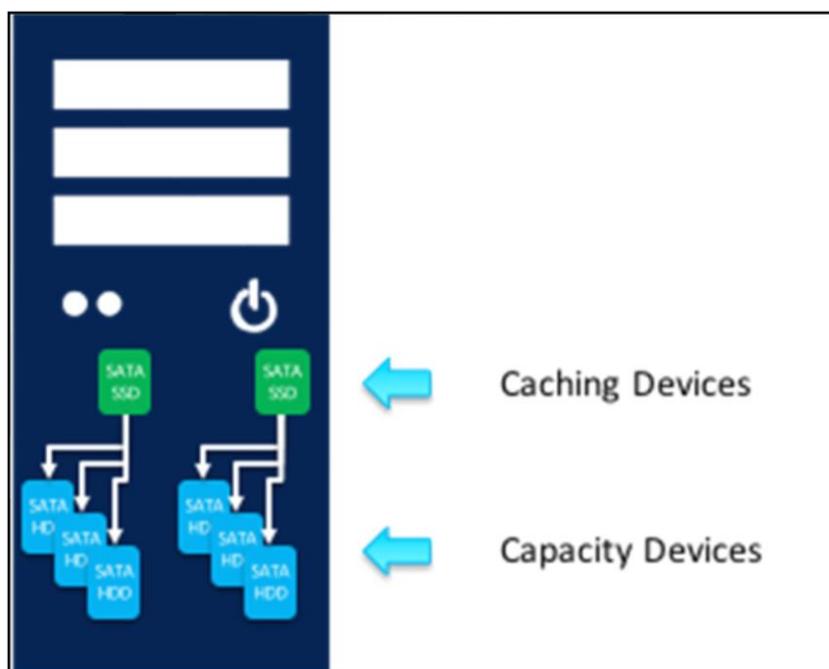
Scale-Out File Server. This final layer is necessary for converged deployments only. It provides remote file access using the SMB3 access protocol to clients, such as another cluster running Hyper-V, over the network, effectively turning Storage Spaces Direct into network-attached storage (NAS).

Software Storage Bus Bandwidth Management

Software Storage Bus also implements a fair access algorithm that ensures fair device access from any server to protect against one server starving out other servers. It also implements an algorithm for IO prioritization that prioritizes Application IO, which usually is IO from virtual machines, over system IO, which usually would be rebalance or repair operations. However, at the same time it ensures that rebalance and repair operations can make forward the progress. Finally, it implements an algorithm that de-randomizes IO going to rotational disk devices to drive a more sequential IO pattern on these devices, despite the IO coming from the application (virtual machines) being a random IO pattern.

Software Storage Bus Cache

Software Storage Bus implements a caching mechanism, which is known as Storage Bus Cache. The Storage Bus Cache is scoped to each server (per node cache) and is agnostic to the storage pools and virtual disks defined in the system. Storage Bus Cache is resilient to failures as it sits underneath the virtual disk, which provides resiliency by writing data copies to different fault domains, which by default are storage nodes. When Storage Spaces Direct is enabled in a cluster, Software Bus Cache identifies which devices to use as caching devices and which devices are capacity devices. Caching devices cache data for the capacity devices, essentially creating hybrid disks. Once it has been determined whether a device is a caching device or a capacity device, the capacity devices are bound to a caching device in a round robin manner, as shown in the following diagram. Rebinding will occur if there is a topology change, such as if a caching device fails.



The behavior of the caching devices is determined by the actual disk configuration of the system and outlined in the table below:

Storage Configuration	Caching Devices	Capacity Devices	Caching Behavior
SATA SSD + SATA HDD	All SATA SSD	All SATA HDD	Read + Write
NVMe SSD + SATA HDD	All NVMe SSD	All SATA HDD	Read + Write
NVMe SSD + SATA SSD	All NVMe SSD	All SATA SSD	Write only
NVMe SSD + SATA SSD + SATA HDD	All NVMe SSD	All SATA SSD and SATA HDD	Write only for SATA SSD and Read + Write for SATA HDD

In systems with rotational capacity devices (HDD), Software Storage Cache will act as both a read and write cache. This is because there is a seek penalty on rotational disk devices. In systems with all flash devices (NVMe SSD + SATA SSD), SBC will only act as a write cache. Because the NVMe devices will absorb most of the writes in the system, it is possible to use mixed-use or even read-intensive SATA SSD devices, which can lower the overall cost of flash in the system. In systems with only a single tier of devices, such as an all NVMe system or all SATA SSD system, SBC will need to be disabled.

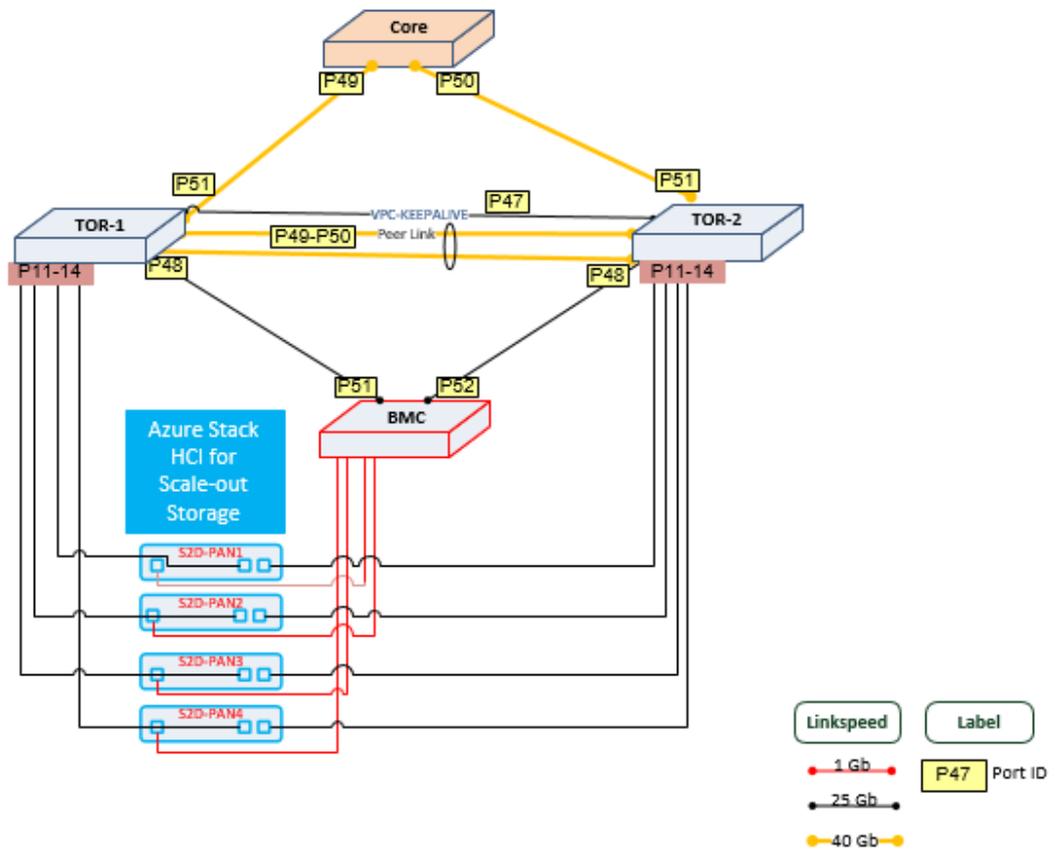
Software Bus Cache creates a special partition on each caching device that, by default, consumes all available capacity except 32GB. The 32GB is used for storage pool and virtual disk metadata. Software Bus Cache uses memory for runtime data structures, about 5GB of memory per TB of caching devices in the node. For instance, a system with 4x 800GB caching devices requires about 16GB of memory to manage the cache in addition to the base operating systems and any hosted hyper-converged virtual machines.

QCT Deployment Lab

1. Hardware for this Guide

T21P-4U (4 nodes) (Alias name: S2PA)				
SKU	Description	Qty per unit	Total	Version
Server Platform	T21P-4U 3.5" Storage server	1	4	BIOS: S2P_3B10 BMC: 3.33
CPU	Intel E5-2698 v4	2	8	
Memory	Samsung 32G DDR4 RDIMM 2133Mhz	16	64	
Cache	Samsung 1.92TB 2.5" SATA SSD	8	32	104Q
Storage	Seagate 8TB 3.5" SATA HDD (ST8000NM0055)	70	280	PN01
Boot Drive	Intel 480G 2.5" SATA SSD	1	4	G2010140
HBA Card	QCT LSI 3008- IT firmware mode	1	4	FW: 1B03
NIC Card	Mellanox ConnectX-4 Lx Ethernet Adapter	1	4	FW: 14.24.1000

2. Server Nodes / Network Diagram



3. Prepare the QCT physical server node

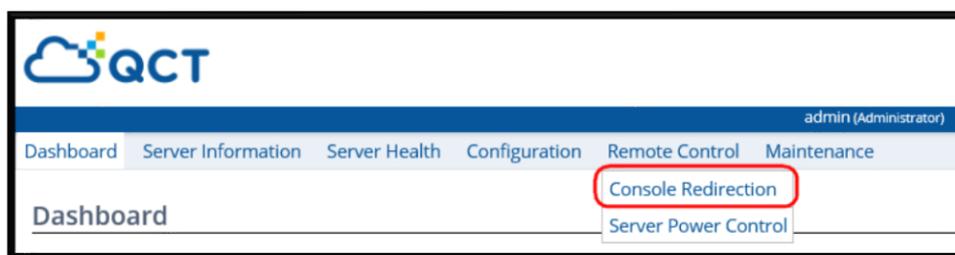
Best practices dictate that with a new server deployment, the first task is to review the system firmware and drivers relevant to the incoming operating system. If the system has the latest firmware and drivers installed it will expedite tech support calls, and may reduce the need for such calls.

<https://qct.io/product/index/Storage/Storage-Server/4U-Storage-Server/QuantaPlex-T21P-4U#download>

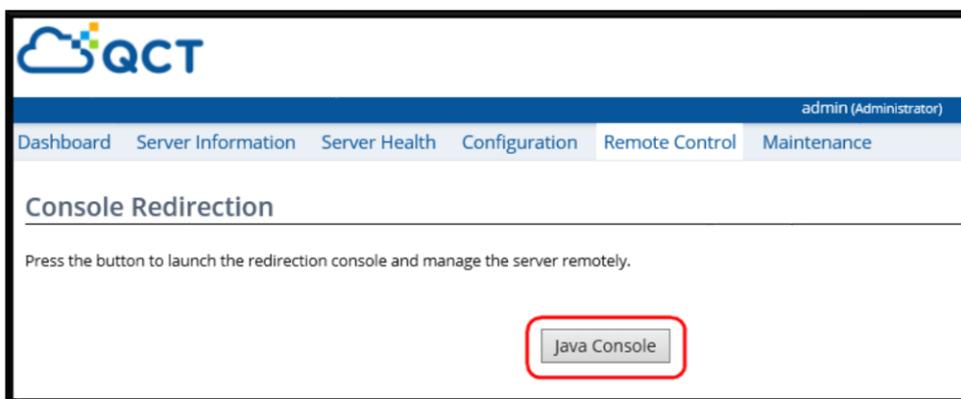
Windows Server 2019 contains all the drivers necessary for this solution with the exception of the Qlogic and Mellanox CX-3 driver, which was updated by Mellanox and Qlogic after the final Release to Manufacturing (RTM) build of the OS was released.

All of the QCT servers feature an BMC web console to provide remote out-of-band management, including remote control and remote media. The following steps will shed some light on the installation:

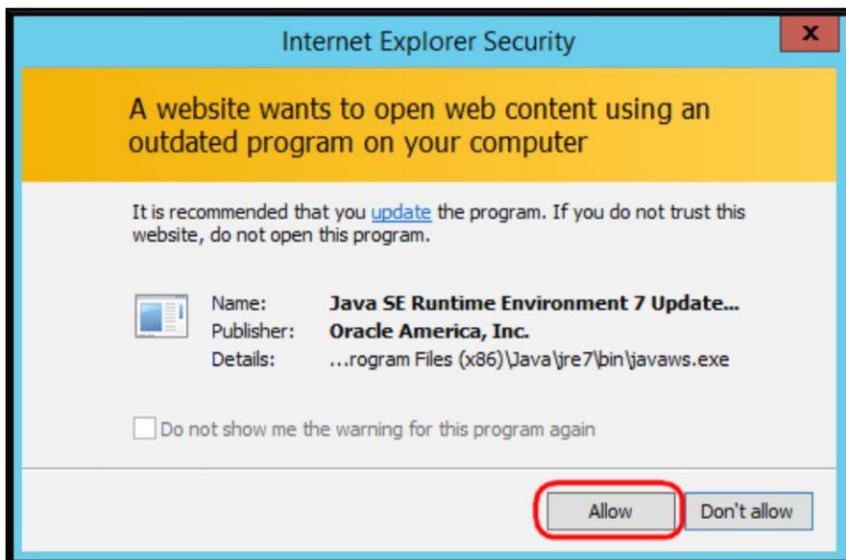
1. Login the BMC web console with administrator account (Default ID/Password is **admin/admin**). Select the item **[Console Redirection]** under **[Remote Control]**.



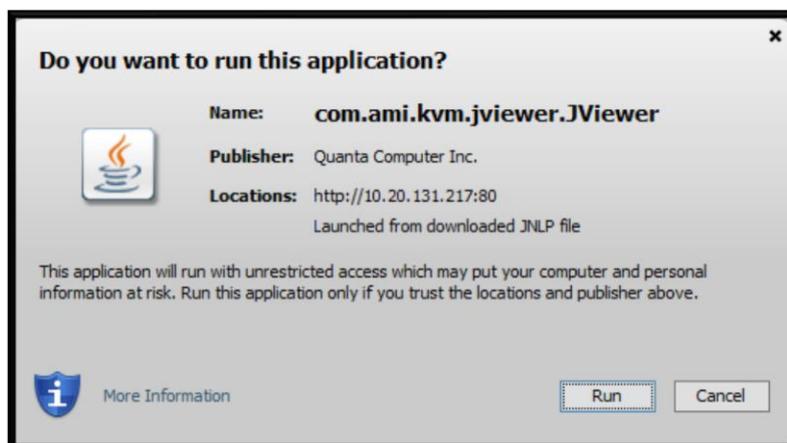
2. Click on **[Java Console]** button. (Please [install Java](#) before running the console).



3. If the [Internet Explorer Security] warning appears, please choose **[Allow]**.



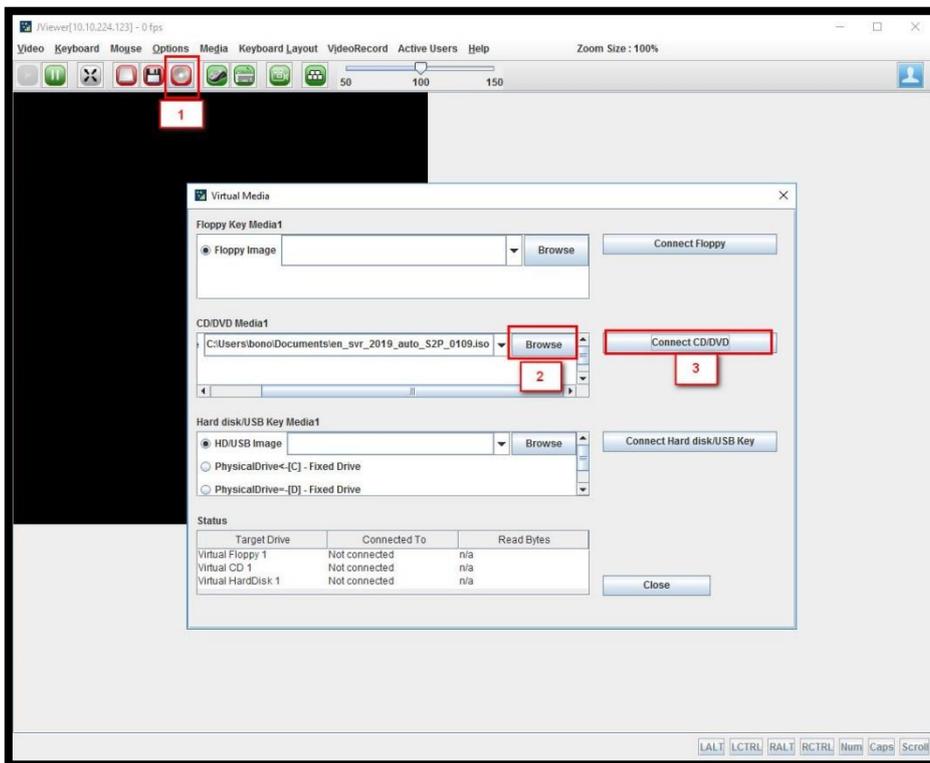
4. Confirm to Run the JViewer application.



5. At the JViewer Console, click on the CD-ROM icon.



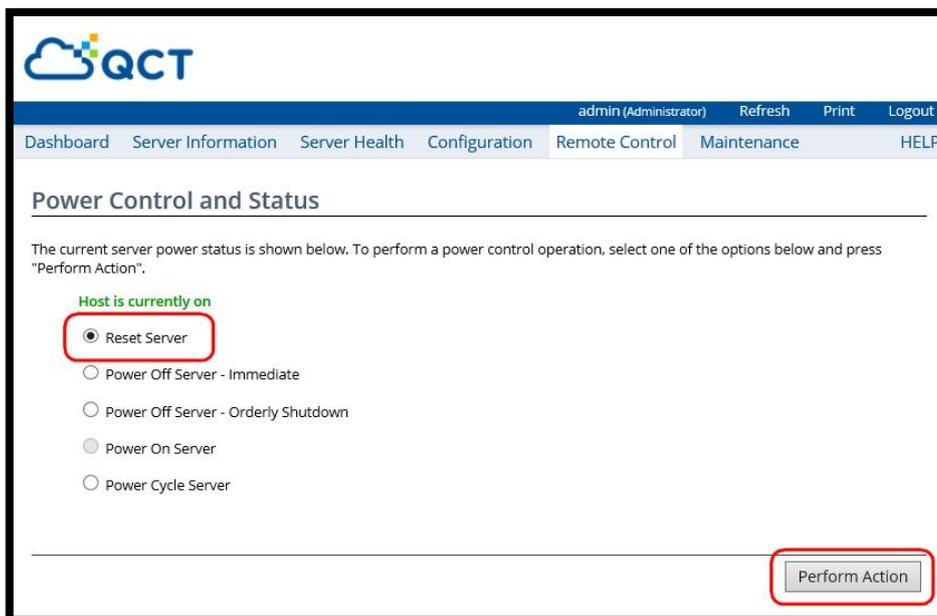
6. Browse and Select the Windows Server 2019 iso file then connected.



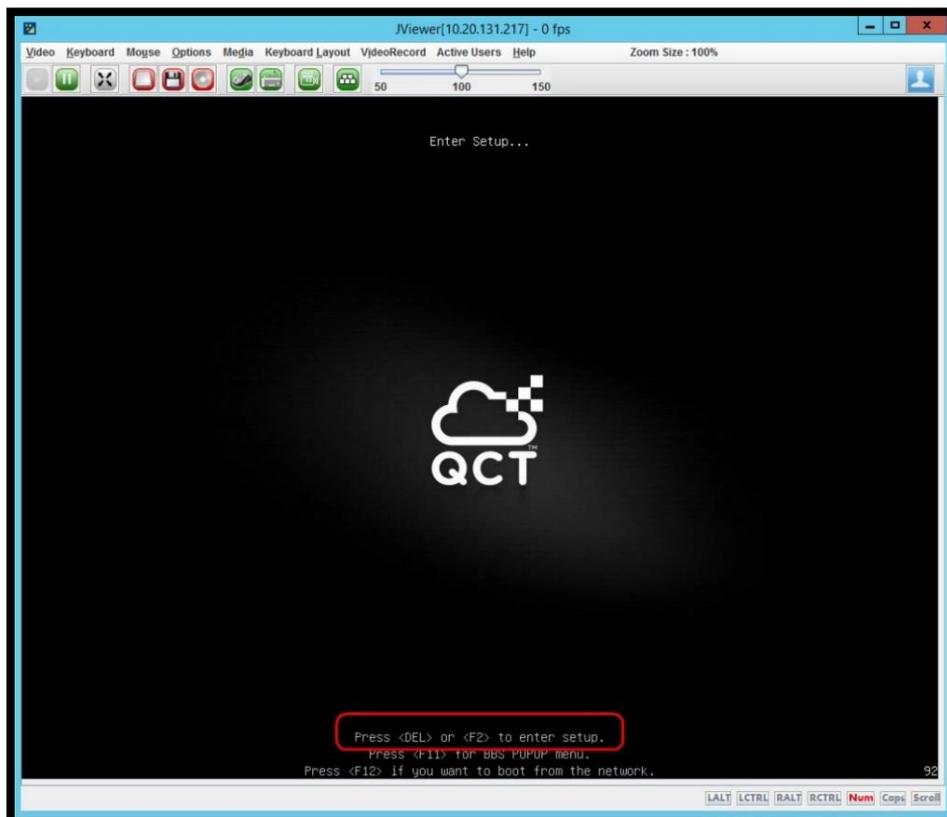
7. Back to the BMC web console, and select the item [Server Power Control] under [Remote Control].



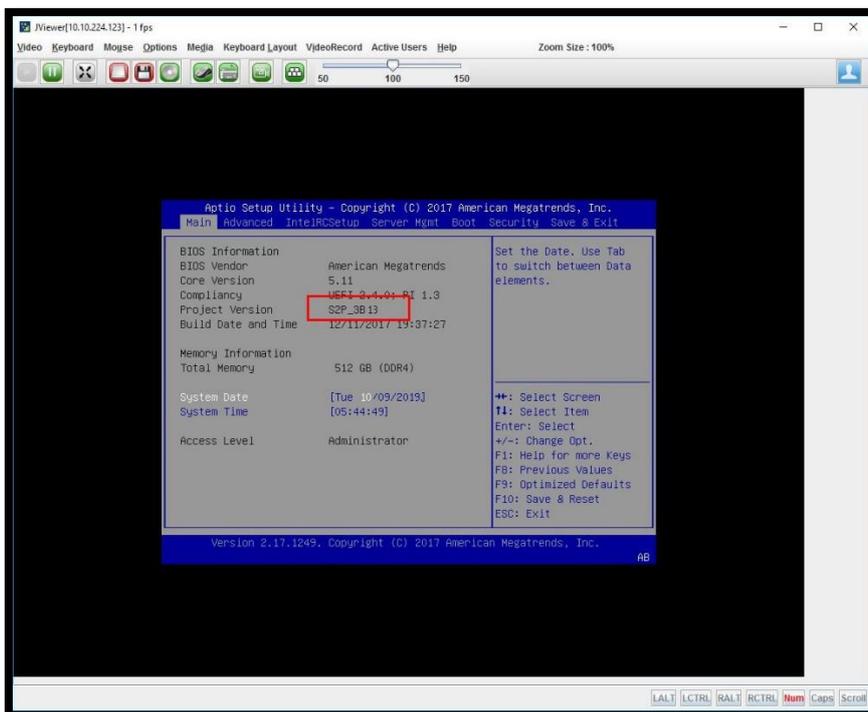
8. Select [Reset Server] or [Power On Server], then press **[Perform Action]** button.



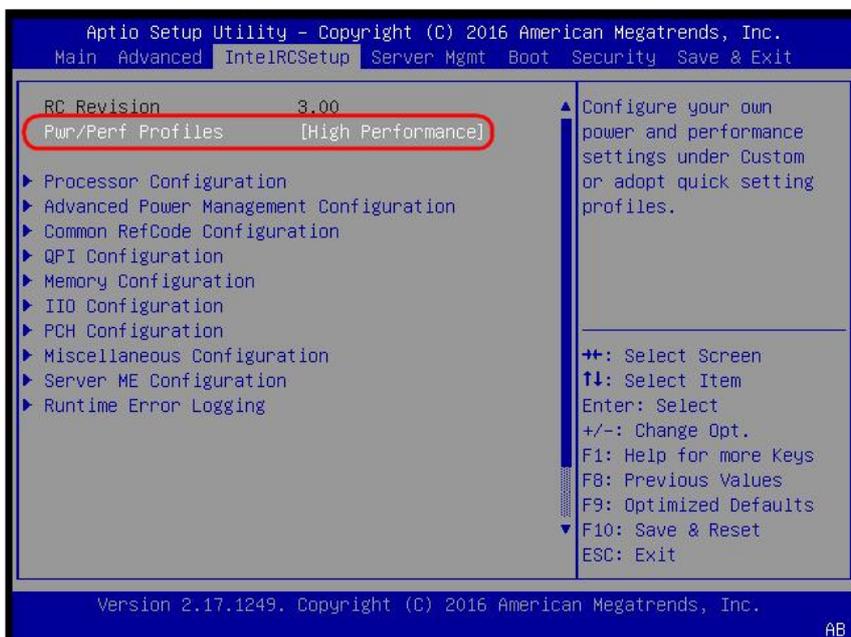
9. At the startup screen, press **[DEL]** key or **[F2]** key to enter BIOS setup.



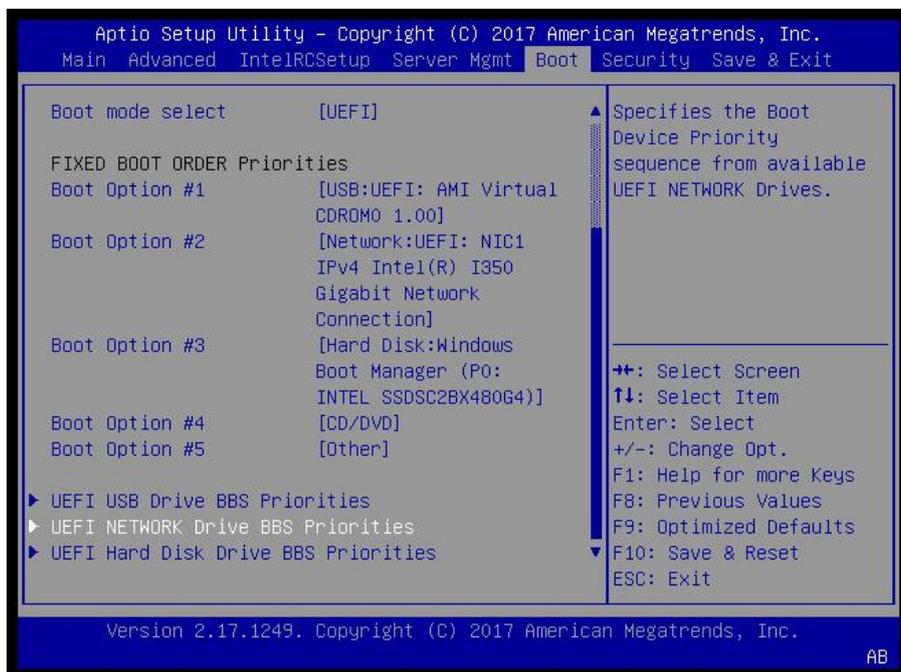
10. Review the Server Model and BIOS version.



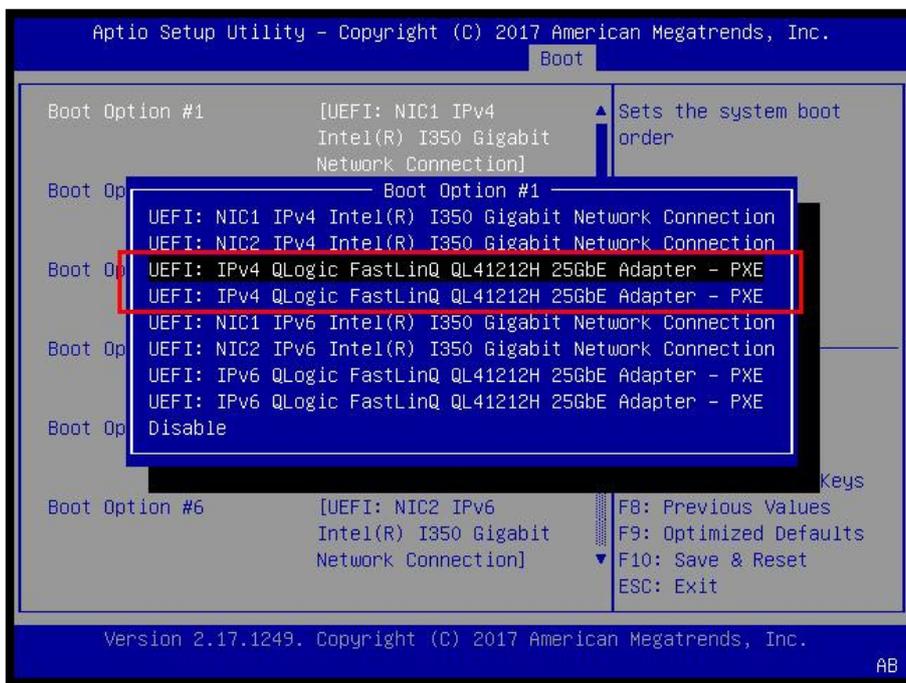
11. Select the [IntelRCSetup] Tab, setting the [Pwr/Perf Profiles] to [High Performance].



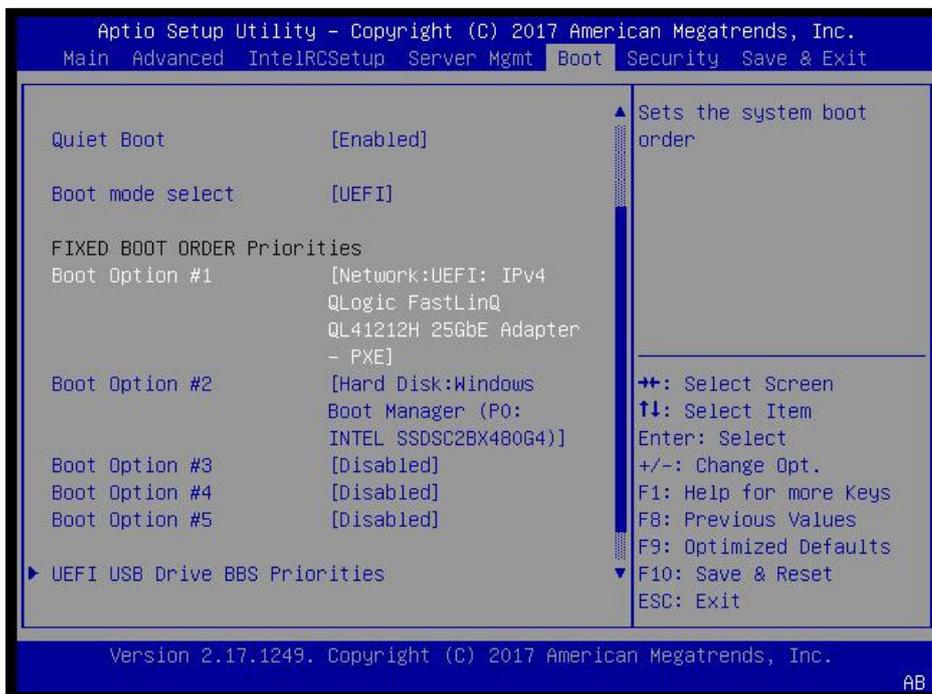
12. Back to the main menu, select the [Boot] tab, config the **[Boot mode select]** as **[UEFI]**. And select the **[UEFI NETWORK BBS Priorities]**.



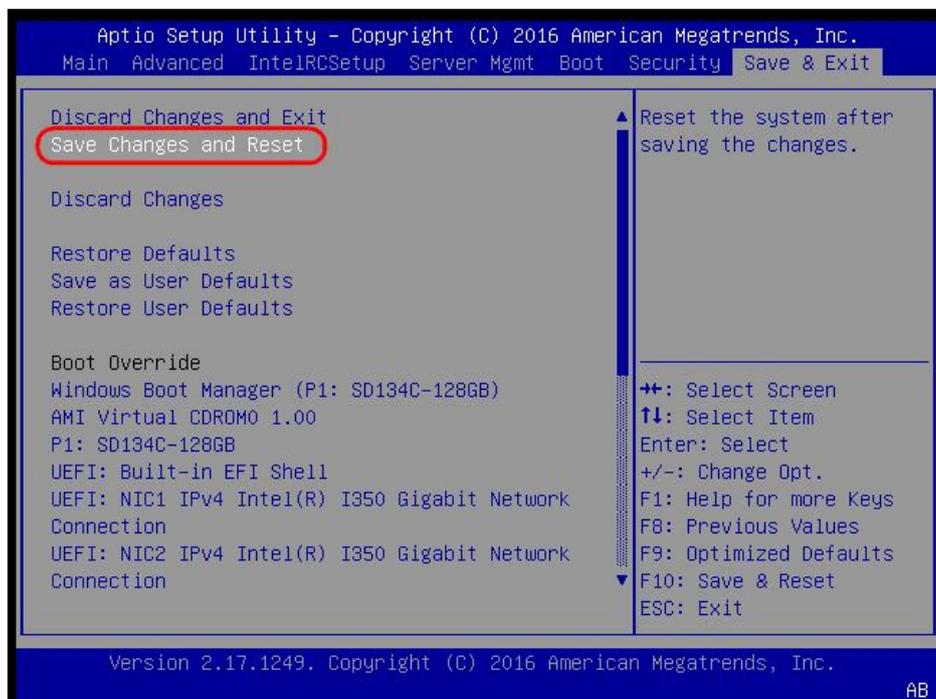
13. On the **[Boot Option #1]**, select the PXE boot NIC then back to the main menu.



14. Setting the NIC as the **[Boot Option #1]**.



15. Finally, **[Save Changes and Reset]**.



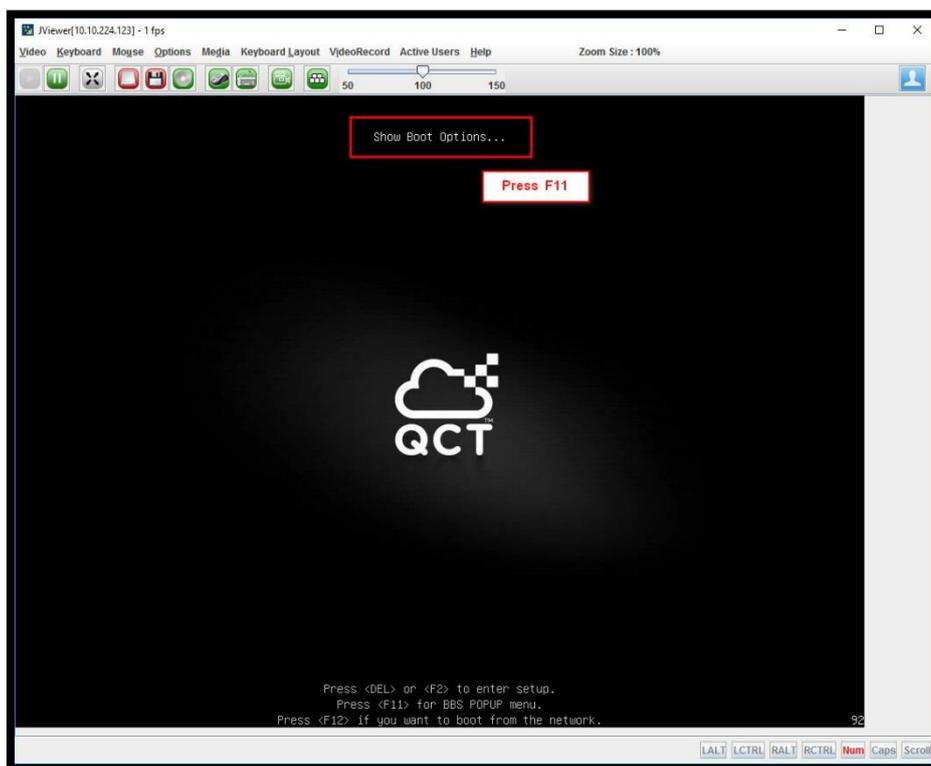
In the next section, we'll continue to install the Windows Server 2019 OS.

4. Install Windows Server 2019 OS

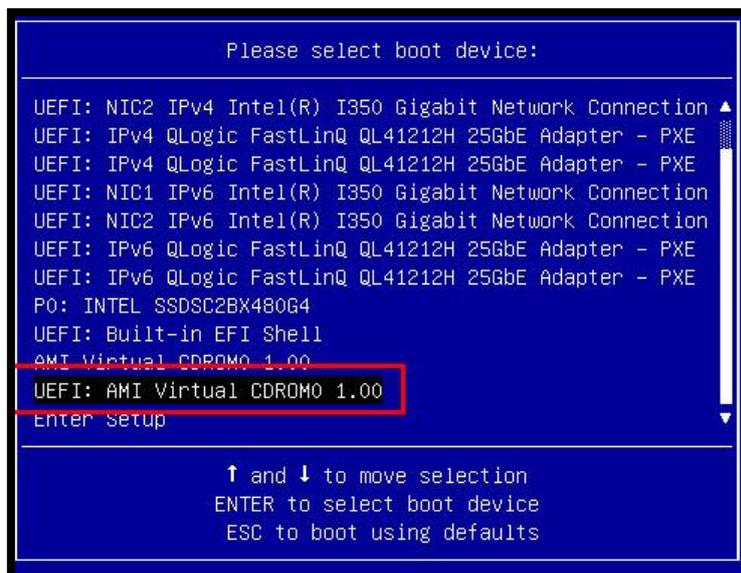
Install Windows from a variety of sources:

- Remote ISO media mount via the BMC/QSM
- Bootable USB media with the installation content

1. With the method of Windows deployment selected, power the server on and press **[F11]** to choose the installation method.



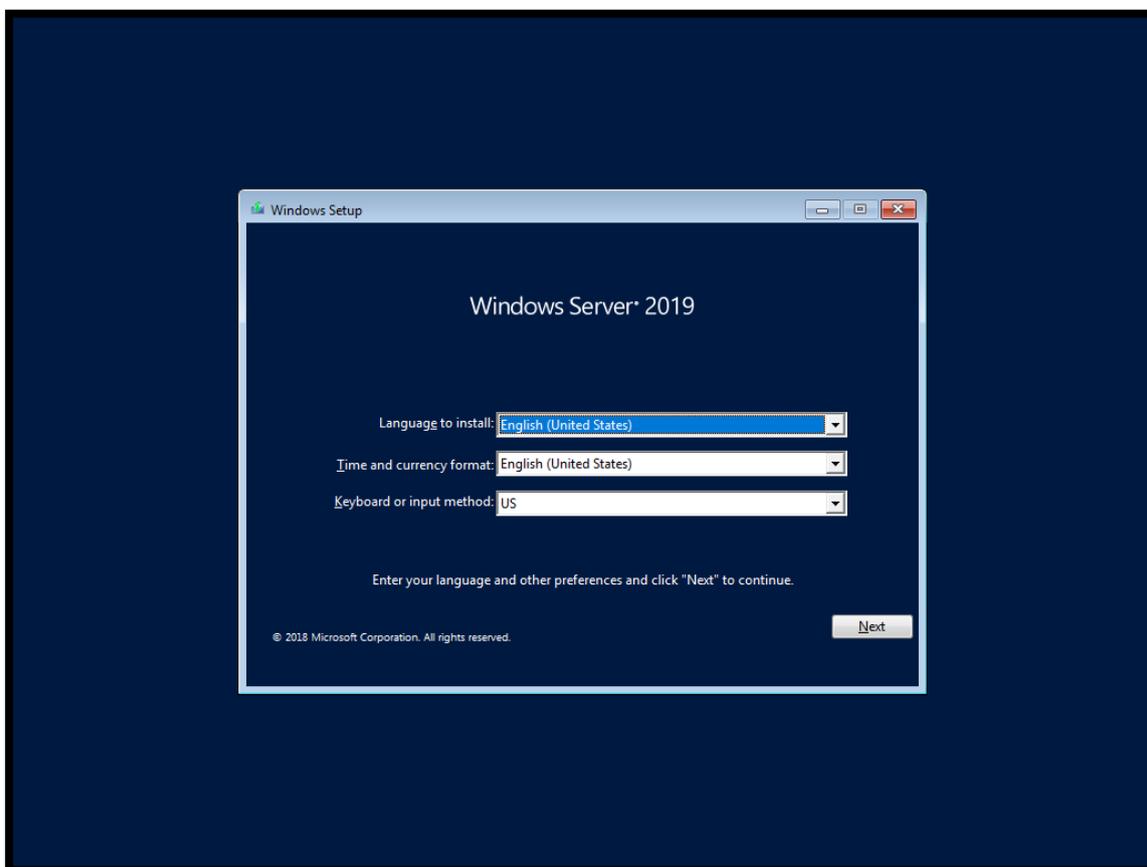
2. Select the **[UEFI: AMI Virtual CDROM0 1.00]**.



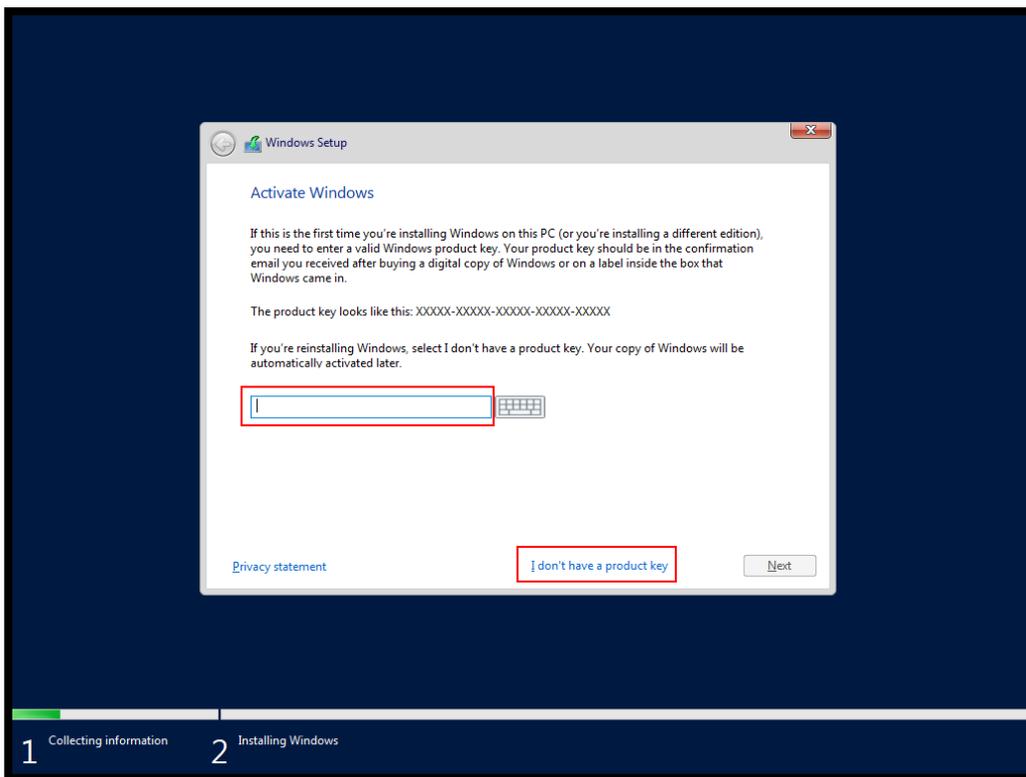
3. Press any key to boot from installation media.



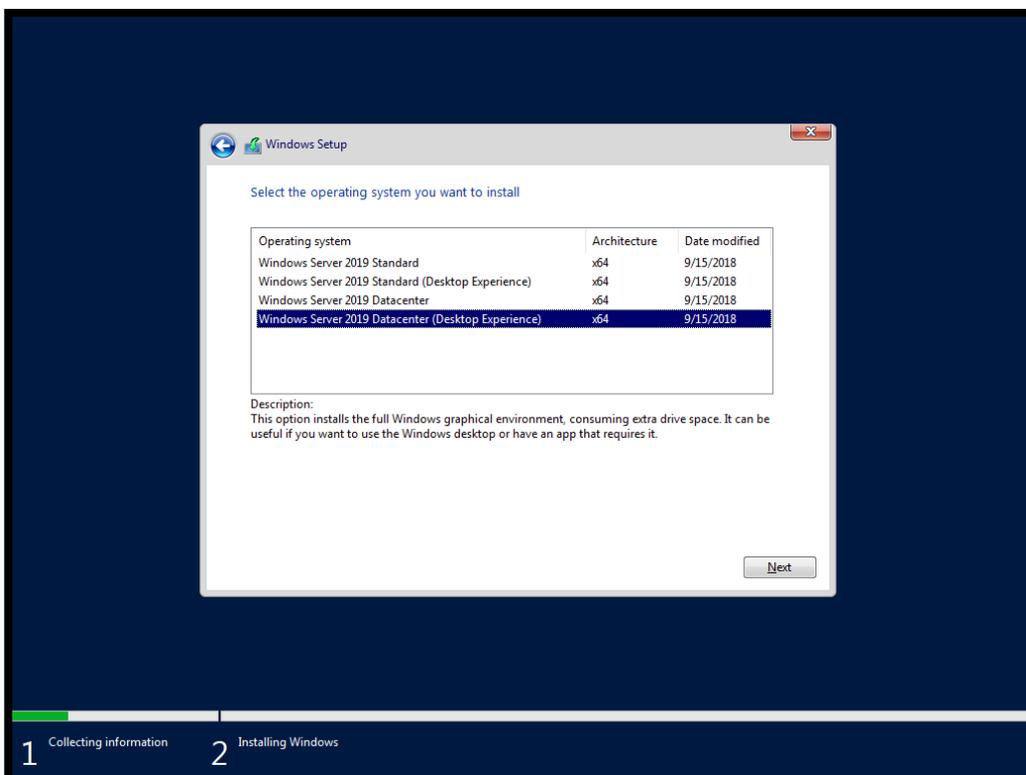
4. Select the appropriate language pack, correct input device, and determine the geography.



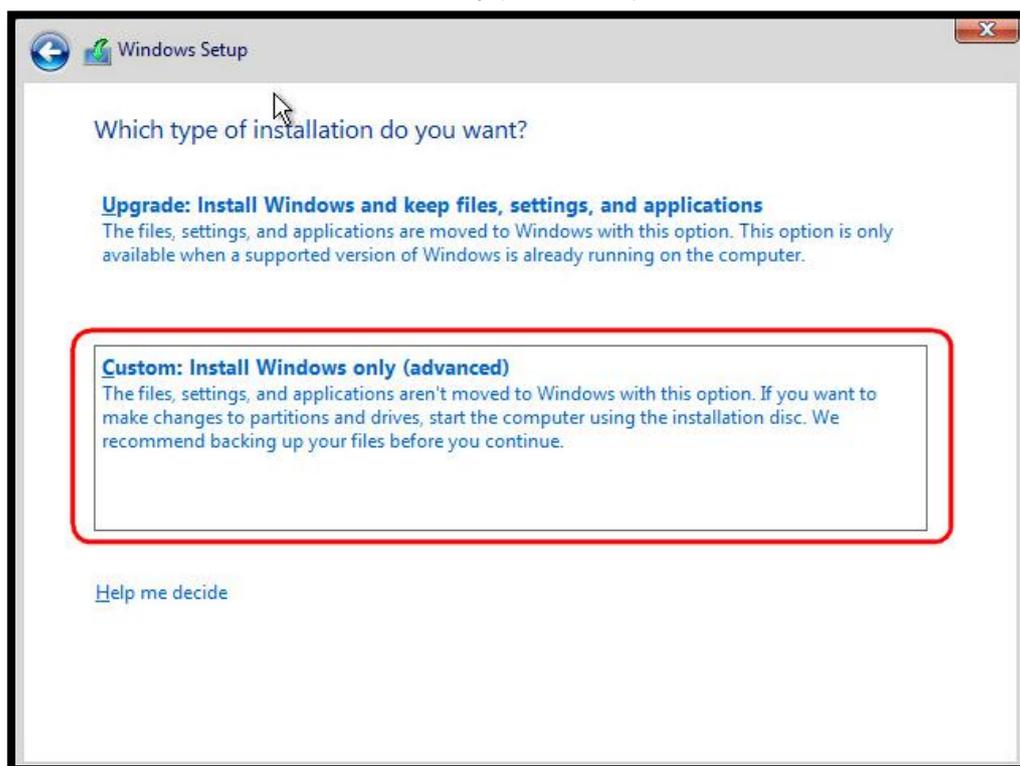
5. Input production key or activate later:



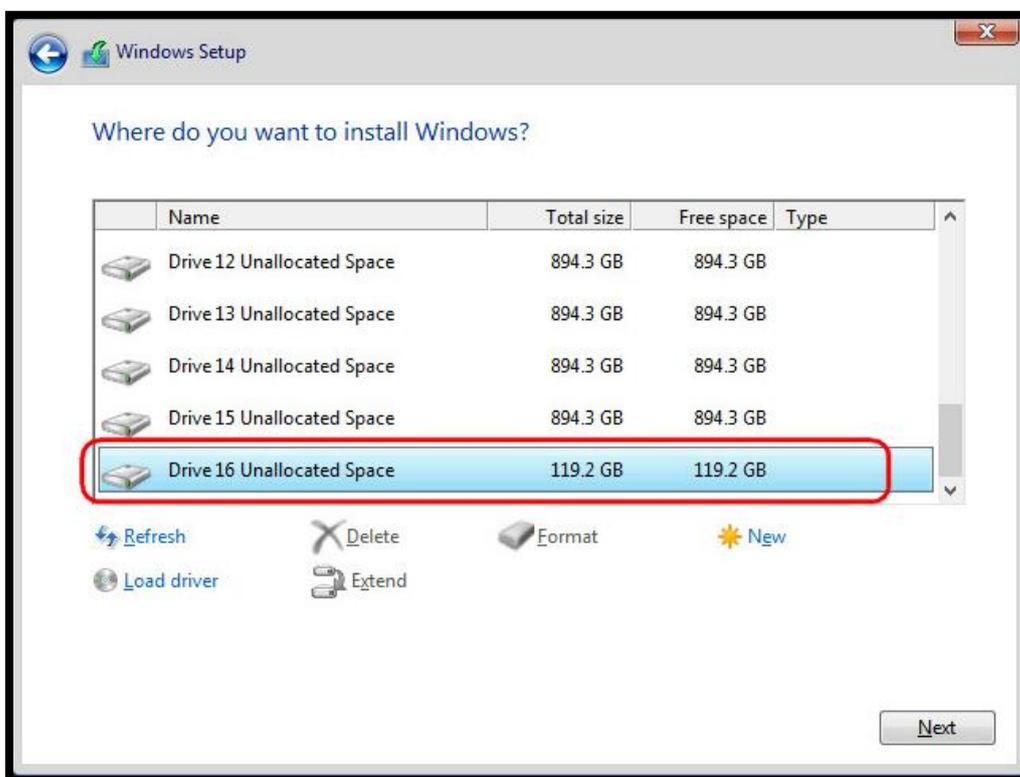
6. Select the desired OS edition (GUI or Core components only).



7. Choose [Custom: Install Windows only (advanced)].



8. Select the SSD as the target to install Windows (you might need to scroll through a list of available drives). Then start to complete the installation of the OS.





5. Install Hardware Drivers and Windows Roles/Features

This document is used in Windows Server 2019 RTM (Build number 17763), please use following **Powershell** commands to install and config :

1. Install hardware drivers

```
Start-Process -FilePath "d:\auto\IntelChipSet\SetupChipset0529.exe" -ArgumentList "-norestart -s" -WindowStyle Hidden -Wait
```

```
Start-Process -FilePath "d:\auto\Mellanox_Driver\MLNX_WinOF2-2_20_50000_All_x64" -ArgumentList '/S /v"/qn"' -Wait
```

2. Set power option ,page file ,advanced firewall ,remote desktop

```
Start-Process -FilePath "powercfg" -ArgumentList "/S SCHEME_MIN" -WindowStyle Hidden -Wait
```

```
Start-Process -FilePath "powercfg" -ArgumentList "-x -monitor-timeout-ac 0" -WindowStyle Hidden -Wait
```

```
Set-CimInstance -Query "SELECT * FROM Win32_ComputerSystem" -Property @{AutomaticManagedPagefile="False"}
```

```
Set-CimInstance -Query "SELECT * FROM Win32_PageFileSetting" -Property @{InitialSize=12800; MaximumSize=12800}
```

```
Start-Process -FilePath "netsh" -ArgumentList "advfirewall set allprofiles state off" -Wait
```

```
Set-ItemProperty -Path "HKLM:SYSTEM\CurrentControlSet\Control\Terminal Server" -name "fDenyTSConnections" -Value 0
```

3. Config administrator auto logon (optional)

```
Set-ItemProperty -Path "HKLM:\Software\Microsoft\Windows NT\CurrentVersion\Winlogon" -name "AutoAdminLogon" -Value 1
```

```
Set-ItemProperty -Path "HKLM:\Software\Microsoft\Windows NT\CurrentVersion\Winlogon" -name "DefaultDomainName" -Value "rd1demo.qct"
```

```
Set-ItemProperty -Path "HKLM:\Software\Microsoft\Windows NT\CurrentVersion\Winlogon" -name "DefaultUserName" -Value "Administrator"
```

```
Set-ItemProperty -Path "HKLM:\Software\Microsoft\Windows NT\CurrentVersion\Winlogon" -name "DefaultPassword" -Value "s2d@QCTRD"
```

```
Set-ItemProperty -Path "HKLM:\Software\Microsoft\Windows NT\CurrentVersion\Winlogon" -name "autologoncount" -value 8
```

4. Install Windows Roles and Features, rename computer and reboot

```
Install-WindowsFeature -Name Hyper-V, File-Services, Data-Center-Bridging, Failover-Clustering -IncludeManagementTools
```

```
Get-WindowsFeature -Name Hyper-V*,File*,Data*,Failover*
```

```
Rename-Computer -NewName PA-01
```

```
Restart-Computer
```

5. Rename NetAdapters ,setup DCB/PFC/ETS functions

Please type the following command sequence respectively to enable DCB functionality for each host and set S2D PFC / ETS:

```
Rename-NetAdapter -InterfaceDescription "mell*ter" -NewName RDMA01
```

```
Rename-NetAdapter -InterfaceDescription "mell*ter #2" -NewName RDMA02
```

```
New-NetQosPolicy "SMB" -NetDirectPortMatchCondition 445 -PriorityValue8021Action 3
```

```
Enable-NetQosFlowControl -Priority 3
```

```
Disable-NetQosFlowControl 0,1,2,4,5,6,7
```

```
Enable-NetAdapterQos -InterfaceAlias RDMA01, RDMA02
```

```
New-NetQoSSTrafficClass "SMB" -Priority 3 -BandwidthPercentage 50 -Algorithm ETS
```

```
Get-NetAdapterQos
```

```
PS C:\> New-NetQoSSTrafficClass SMB -Priority 3 -BandwidthPercentage 50 -Algorithm ETS
```

Name	Algorithm	Bandwidth(%)	Priority	PolicySet
SMB	ETS	50	3	Global

```
PS C:\> Get-NetAdapterQos
```

```
Name : RDMA02
Enabled : True
Capabilities :
Hardware :
Current :
MacSecBypass : NotSupported
DcbxSupport : None
NumTCs (Max/ETS/PFC) : 8/8/8
OperationalTrafficClasses : TC TSA Bandwidth Priorities
0 ETS 50% 0-2,4-7
1 ETS 50% 3
OperationalFlowControl : Priority 3 Enabled
OperationalClassifications : Protocol Port/Type Priority
NetDirect 445 3
```

```
Name : RDMA01
Enabled : True
Capabilities :
Hardware :
Current :
MacSecBypass : NotSupported
DcbxSupport : None
NumTCs (Max/ETS/PFC) : 8/8/8
OperationalTrafficClasses : TC TSA Bandwidth Priorities
0 ETS 50% 0-2,4-7
1 ETS 50% 3
```

6. Create SET (Switch Embedded Teaming) and config network

```
New-VMSwitch -Name S2D-vSwitch -NetAdapterName RDMA01, RDMA02 -  
EnableEmbeddedTeaming $True -AllowManagementOS $false
```

```
Add-VMNetworkAdapter -SwitchName S2D-vSwitch -Name SMB01 -ManagementOS  
Add-VMNetworkAdapter -SwitchName S2D-vSwitch -Name SMB02 -ManagementOS  
Add-VMNetworkAdapter -SwitchName S2D-vSwitch -Name MGMT -ManagementOS
```

```
$netAdapter = Get-NetAdapter -Name "vEthernet (MGMT)"  
$netAdapter | Set-NetIPInterface -Dhcp Disabled  
$netAdapter | New-NetIPAddress -IPAddress 10.106.5.137 -PrefixLength 24 -  
DefaultGateway 10.106.5.1
```

```
Set-DnsClientServerAddress -InterfaceAlias "vEthernet (MGMT)" -ServerAddresses  
10.106.48.100
```

```
$netAdapter = Get-NetAdapter -Name "vEthernet (SMB01)"  
$netAdapter | Set-NetIPInterface -Dhcp Disabled  
$netAdapter | New-NetIPAddress -IPAddress 10.106.7.137 -PrefixLength 24
```

```
$netAdapter = Get-NetAdapter -Name "vEthernet (SMB02)"  
$netAdapter | Set-NetIPInterface -Dhcp Disabled  
$netAdapter | New-NetIPAddress -IPAddress 10.106.8.137 -PrefixLength 24
```

6-1. Set vNIC Vlan (if necessary!!)

```
Set-VMNetworkAdapterVlan -VMNetworkAdapterName MGMT -Access -VlanId 1605  
-ManagementOS  
Set-VMNetworkAdapterVlan -VMNetworkAdapterName SMB01 -Access -VlanId 1607  
-ManagementOS  
Set-VMNetworkAdapterVlan -VMNetworkAdapterName SMB02 -Access -VlanId 1608  
-ManagementOS
```

7. Re-Enable NetAdapters and check RDMA function

```
Disable-NetAdapter "vEthernet (SMB01)" -Confirm:0  
Enable-NetAdapter "vEthernet (SMB01)"  
Disable-NetAdapter "vEthernet (SMB02)" -Confirm:0  
Enable-NetAdapter "vEthernet (SMB02)"  
Enable-NetAdapterRdma "vEthernet (SMB01)", "vEthernet (SMB02)"  
Set-VMNetworkAdapterTeamMapping -VMNetworkAdapterName 'SMB01' -  
ManagementOS -PhysicalNetAdapterName 'RDMA01'  
Set-VMNetworkAdapterTeamMapping -VMNetworkAdapterName 'SMB02' -  
ManagementOS -PhysicalNetAdapterName 'RDMA02'  
  
Get-NetAdapter  
Get-NetAdapterRdma  
Get-SmbClientNetworkInterface
```

```

PS C:\Users\Administrator.RD1DEMO> Get-NetAdapterRdma

Name                InterfaceDescription          Enabled
-----
vEthernet (MGMT)    Hyper-V Virtual Ethernet Adapter #3  False
vEthernet (SMB02)   Hyper-V Virtual Ethernet Adapter #2  True
vEthernet (SMB01)   Hyper-V Virtual Ethernet Adapter     True
RDMA01              Mellanox ConnectX-3 Pro Ethernet Adapter True
RDMA02              Mellanox ConnectX-3 Pro Ethernet Adap... True

PS C:\Users\Administrator.RD1DEMO> Get-SmbClientNetworkInterface

Interface Index  RSS Capable  RDMA Capable  Speed  IpAddresses  Friendly Name
-----
10           False        False         10 Gbps {fe80::81c0:4430:273c:c22f} Local Area Connection* 2
11           True         True          20 Gbps {fe80::e566:96c0:c190:2954, 10.20.135.117} vEthernet (SMB01)
3            True         True          20 Gbps {fe80::557f:cd94:bee2:24a9, 10.20.135.217} vEthernet (SMB02)
7            True         False         20 Gbps {fe80::25a5:4ef:19fc:1914, 10.20.134.217} vEthernet (MGMT)

```

8. Join Domain

```

$DomainName = ' rd1demo.qct '
$DomainUserName = 'Administrator'
$DomainPasswordSec = "s2d@QCTRD" | ConvertTo-SecureString -asPlainText -Force
$DomainCredential = New-Object
System.Management.Automation.PSCredential("$DomainName\$DomainUserName"
,$DomainPasswordSec)

Add-Computer -DomainName $DomainName -Credential $DomainCredential -Force -
Restart

```

Note!

All four nodes must join the same AD Domain!

6. Create Failover Cluster and Enable S2D

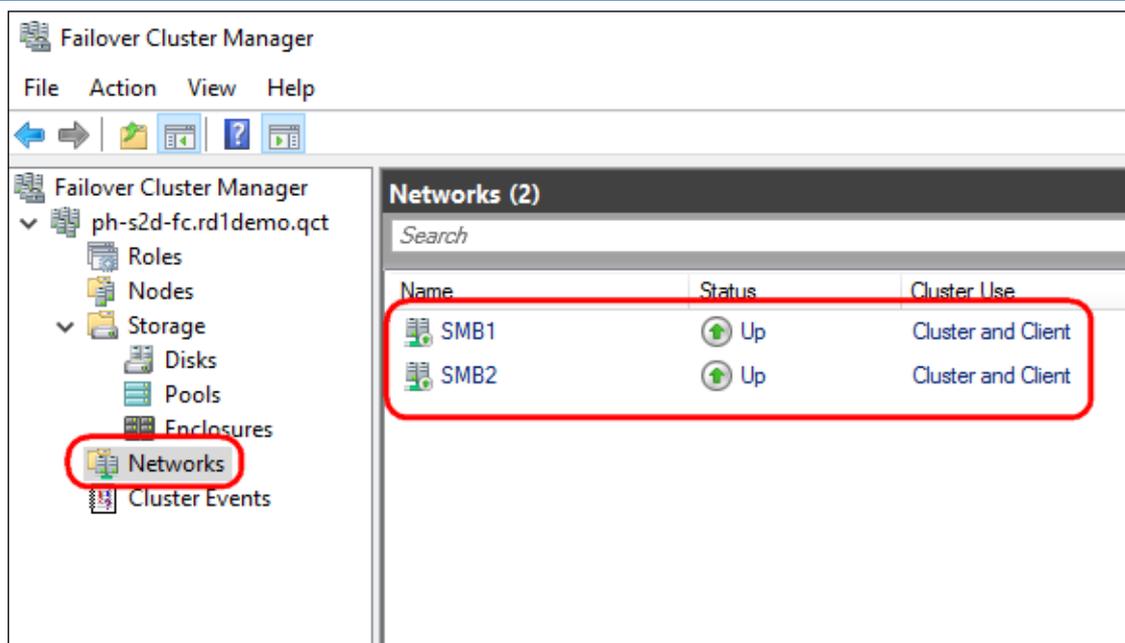
1.Pre-Validation: create the failover cluster

```
Test-Cluster -Node PA-01, PA-02, PA-03, PA-04 -Include "Storage Spaces
Direct",inventory,Network,"System Configuration"
New-Cluster -Name pa-s2d-fc -Node PA-01, PA-02, PA-03, PA-04 -NoStorage -
StaticAddress 10.106.5.140
```

```
# Change the cluster network names
(Get-ClusterNetwork -Name "Cluster Network 1").Name = "SMB1"
(Get-ClusterNetwork -Name "Cluster Network 2").Name = "SMB2"
```

```
# Enable Client traffic on the cluster network
(Get-ClusterNetwork -Name "SMB1").Role = 3
(Get-ClusterNetwork -Name "SMB2").Role = 3
```

```
# Check to make sure the cluster network names and roles are set properly
Get-ClusterNetwork | ft Name, Role, Address
```



The screenshot shows the Failover Cluster Manager interface. The left-hand navigation pane is expanded to show the 'Networks' folder under the cluster 'ph-s2d-fc.rd1demo.qct'. The main pane displays a table of network configurations:

Name	Status	Cluster Use
SMB1	Up	Cluster and Client
SMB2	Up	Cluster and Client

A red rectangle highlights the two network entries (SMB1 and SMB2) in the table.

2. Create S2D Storage Pool

Please confirm the Cluster members host hard drive total number is correct, we have S2PA * 4 nodes in this environment.

```
Get-StorageSubSystem *Cluster*
(Get-StorageSubSystem -Name *pa-s2d-fc* | Get-PhysicalDisk).Count
```

```
Enable-ClusterS2D
```

```
Get-StoragePool *pa-s2d-fc*
Get-PhysicalDisk | ft FriendlyName,HealthStatus,Size,Usage,MediaType -AutoSize
```

```
PS C:\> Get-PhysicalDisk | ft FriendlyName,HealthStatus,Size,Usage,MediaType -AutoSize
```

FriendlyName	HealthStatus	Size	Usage	MediaType
32G MLC SATADOM	Healthy	31016853504	Auto-Select	SSD
ATA INTEL SSDSC2BB01	Healthy	1200174923776	Journal	SSD
ATA INTEL SSDSC2BB01	Healthy	1200174923776	Journal	SSD
ATA TOSHIBA MG03ACA3	Healthy	3000571527168	Auto-Select	HDD
ATA INTEL SSDSC2BB01	Healthy	1200174923776	Journal	SSD
ATA TOSHIBA MG03ACA3	Healthy	3000571527168	Auto-Select	HDD
ATA TOSHIBA MG03ACA3	Healthy	3000571527168	Auto-Select	HDD
ATA TOSHIBA MG03ACA3	Healthy	3000571527168	Auto-Select	HDD
ATA TOSHIBA MG03ACA3	Healthy	3000571527168	Auto-Select	HDD

Get-StorageTier Performance, Capacity

```
PS C:\> Get-StorageTier Performance, Capacity
```

```
ObjectID           : {1}\S2PHFC\root\Microsoft\Windows\Storage\Providers_v2\SPACES_StorageTier.ObjectId="{0de6160e-822c-4a34-ad6f-8345e8d7cb06}
PassThroughClass   :
PassThroughIds     :
PassThroughNamesp :
PassThroughServer  :
UniqueId           : {0039a7bf-25b1-4aa6-8e0b-66ef97070630}
Description        :
FriendlyName       : Capacity
Interleave         : 262144
MediaType          : HDD
NumberOfColumns   : Auto
NumberOfGroups    : Auto
PhysicalDiskRedundancy : 2
ResiliencySettingName : Parity
Size               : 0
PSComputerName     :
```

```
ObjectID           : {1}\S2PHFC\root\Microsoft\Windows\Storage\Providers_v2\SPACES_StorageTier.ObjectId="{0de6160e-822c-4a34-ad6f-8345e8d7cb06}
PassThroughClass   :
PassThroughIds     :
PassThroughNamesp :
PassThroughServer  :
UniqueId           : {88ee3491-05e7-404a-8f30-80bca346a60d}
Description        :
FriendlyName       : Performance
Interleave         : 262144
MediaType          : SSD
NumberOfColumns   : Auto
NumberOfGroups    : 1
PhysicalDiskRedundancy : 2
ResiliencySettingName : Mirror
Size               : 0
PSComputerName     :
```

Note: From the beginning of Windows Server 2016 RTM, related actions have been simplified. Therefore, when you run Enable-ClusterS2D command, will perform the following actions:

- ✓ Automatic Storage Pool Creation
- ✓ Automatic Cache Configuration (SSD to Journal)
- ✓ Automatic Storage Tier Definitions
 - Automatic SSD (Performance) to Mirror
 - Automatic HDD (Capacity) to Parity

3. Create S2D VirtualDisk

```
$pool = Get-StoragePool S2D*
$perf = Get-StorageTier Performance
$capa = Get-StorageTier Capacity

#Create a Multi-Resilient Volume ( Performance + Capacity )
New-Volume -StoragePool $pool -FriendlyName S2D-vDisk01 -FileSystem CSVFS_ReFS -StorageTiers $perf, $capa -StorageTierSizes 500GB,5TB
```

Note: We use S2D Volume type for the Multi-Resilient (using Erasure Coding), so you must specify Mirror (SSD) / Parity (HDD) Size.

```
#Create a Mirror Volume
New-Volume -StoragePool $pool -FriendlyName "Mirror" -FileSystem CSVFS_ReFS -StorageTierfriendlyNames Performance -StorageTierSizes 5TB

#Create a Parity Volume
New-Volume -StoragePool $pool -FriendlyName "Parity" -FileSystem CSVFS_ReFS -StorageTierfriendlyNames Capacity -StorageTierSizes 12TB
```

Note: For establishing a S2D environment for **Hyper-Converged**, please create a VM in CSV. If using **Disaggregated**, please create SOFS role and share resources to build VM. If you want to re-build the S2D, you have to clean the S2D storage pool and reset the physical disks :

```
Update-StorageProviderCache -DiscoveryLevel Full

$pool = Get-StoragePool | ? IsPrimordial -eq $false
$pool | Set-StoragePool -IsReadOnly:$false
$pool | Get-VirtualDisk | Remove-VirtualDisk -Confirm:$false
$pool | Remove-StoragePool -Confirm:$false -ErrorAction Stop

Get-PhysicalDisk | Reset-PhysicalDisk # remove any spaceport partitions
Update-StorageProviderCache -DiscoveryLevel Full
Clear-ClusterNode -CleanupDisks -Force
```

4. Add / Remove Cluster Node

1. Run the [cluster validation wizard](#) using the Failover Cluster snap-in or with the **Test-Cluster** cmdlet in PowerShell (run as Administrator).

```
Test-Cluster -Node <Node>, <Node>, <Node>, <NewNode> -Include "Storage Spaces Direct", Inventory, Network, "System Configuration"
```

This confirms that the new server is running Windows Server 2019 Datacenter Edition, has joined the same Active Directory Domain Services domain as the existing servers, has all the required roles and features, and has networking properly configured.

2. Run the following cmdlet on the cluster to finish adding the server:

```
Add-ClusterNode -Name <NewNode>
```



5. Storage Spaces Optimize Pool

Windows Server 2016/2019 Storage Spaces Direct can optimize a storage pool to balance data equally across the set of physical disks that comprise the pool.

```
Optimize-StoragePool "S2D on pa-s2d-fc"
```

You can monitor the progress of the optimize job with the following command:
(Open in a new Powershell window)

```
Get-StorageJob | ? Name -eq Optimize
```

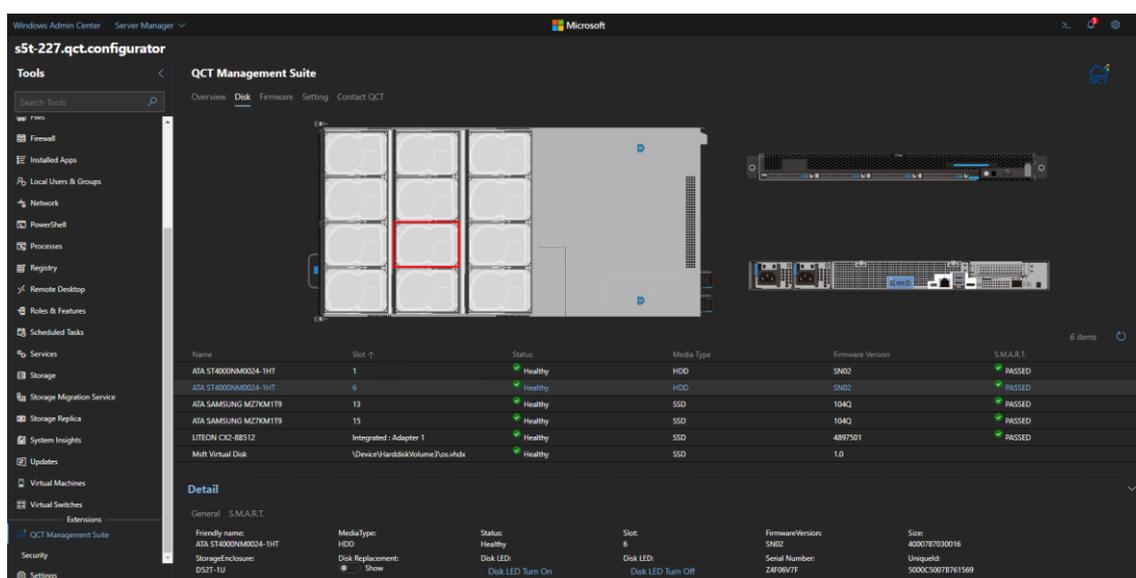
Scale-out storage solution

1. Deploy Windows Admin Center and QCT Management Suite

A simple path to server infrastructure management

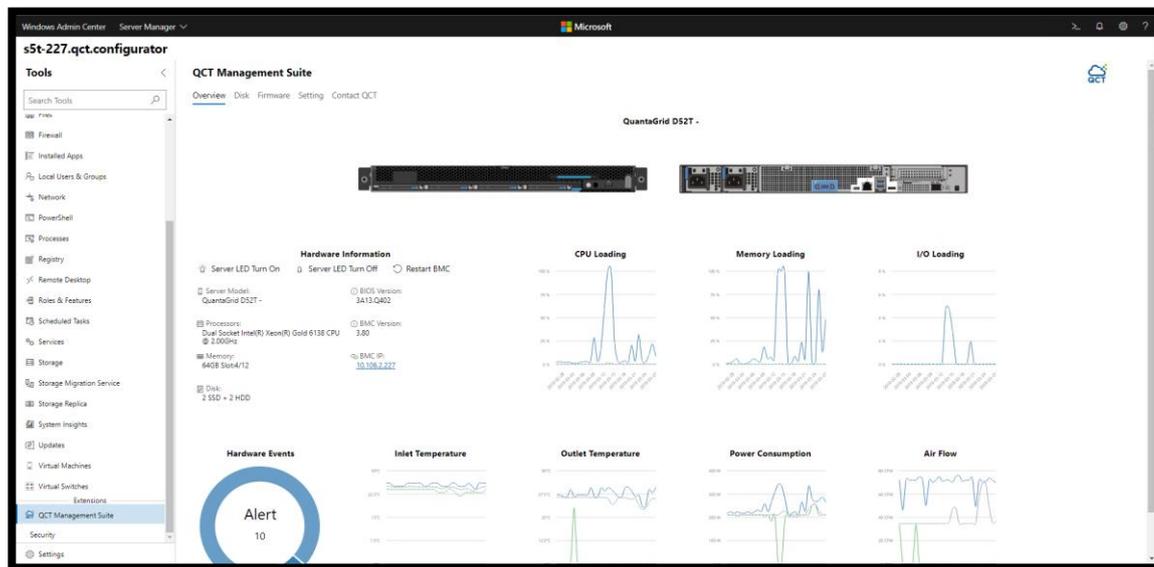
The QCT Management Suite extension for Windows Admin Center provides a single pane of glass dashboard for monitoring system configurations and managing server health of [QCT Azure Stack HCI certified systems](#) : [QuantaGrid D52BQ-2U](#), [QuantaGrid D52T-1ULH](#) and [QuantaPlex T21P-4U](#).

Driven by customer pain points around existing monitoring and management, QCT provides exclusive, complementary features and functions, which includes an overview of system event logs, monitoring drivers, and hardware component health to enhance the overall management experience.

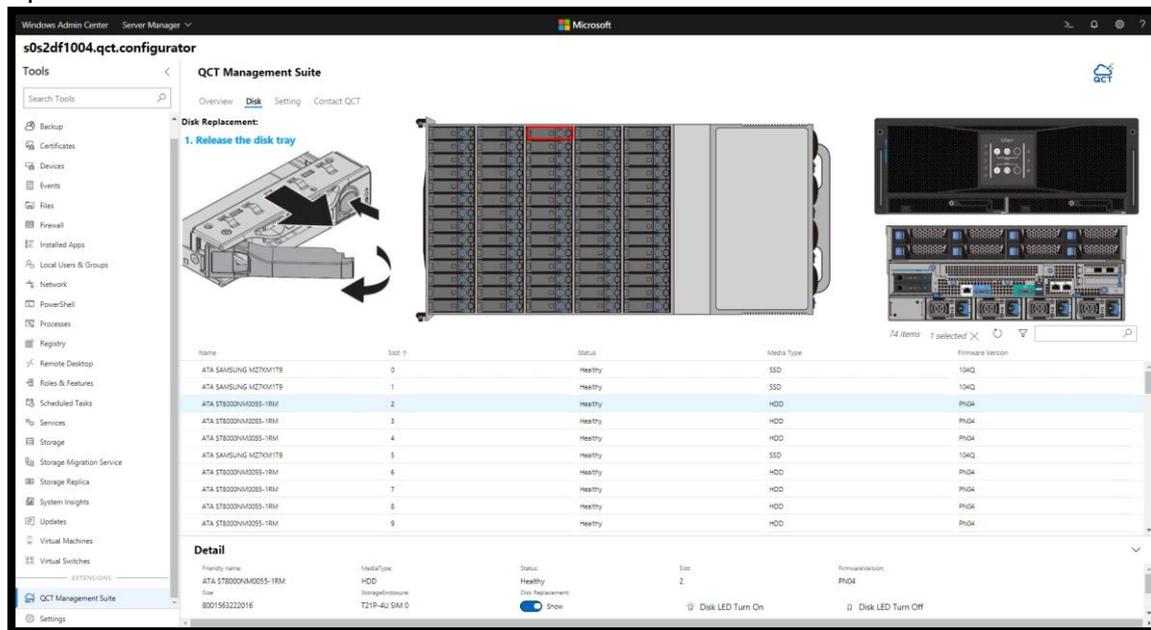


The QCT Management Suite extends the functionality of WAC with the key features below:

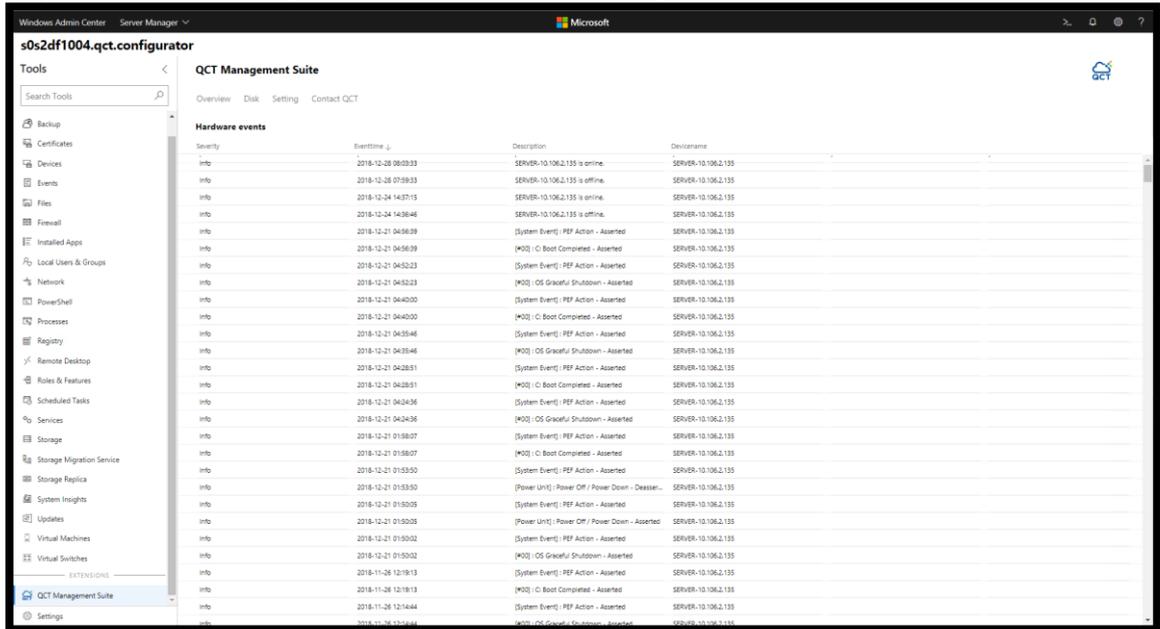
- One-click exclusive hardware management - An intuitive user interface displays hardware information, including model name, processor, memory and BIOS.



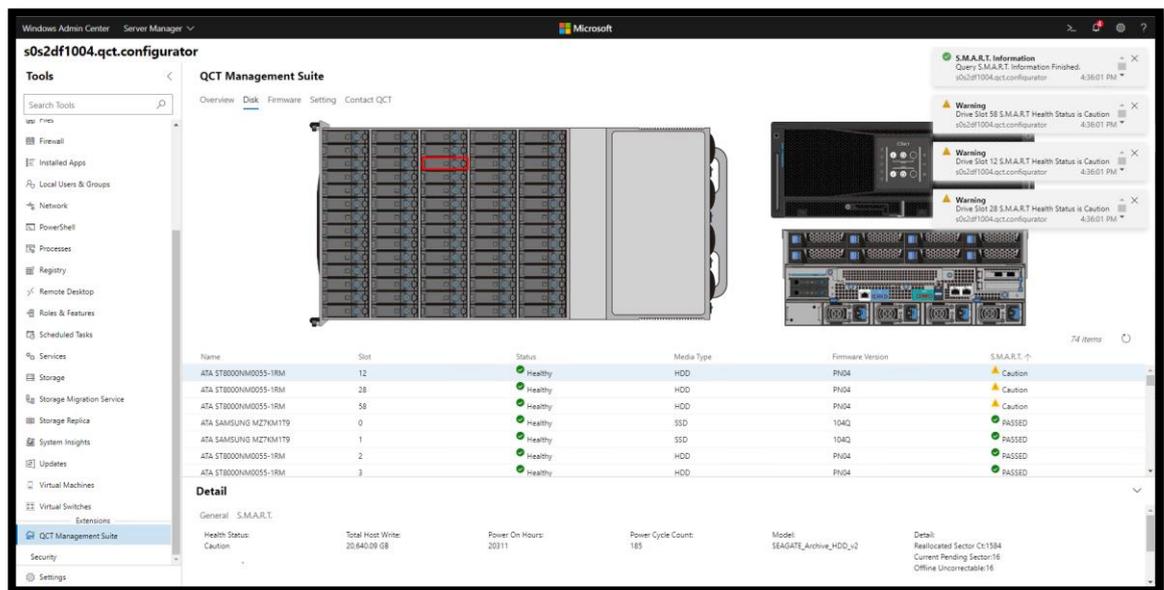
- Disk mapping and LED identification for efficient service support - QCT Management Suite wizard UI design displays the slot of each selected disk with an overview of disk profiles and LED light controls of selected disks for efficient replacement.



- Easy-to-use monitoring tool for hardware event log and health status.



- Predictive disk management - Evaluate the system condition with S.M.A.R.T information and unhealthy notifications which allow organizations to take action before total failure occurs.



Deploy Windows Admin Center

You can manage clusters of servers, Hyper-V clusters, and hosts that run on-premises or in Azure. And this tool is particularly useful when used for hybrid workloads.

Windows Admin Center integrates with Azure services. It provides services, such as Azure Active Directory, Azure Backup, Azure Site Recovery, and more.

This environment can be mixed (starting with Windows Server 2008 R2 with limited functionalities), and the tool can handle all different versions of the server OS from Microsoft.

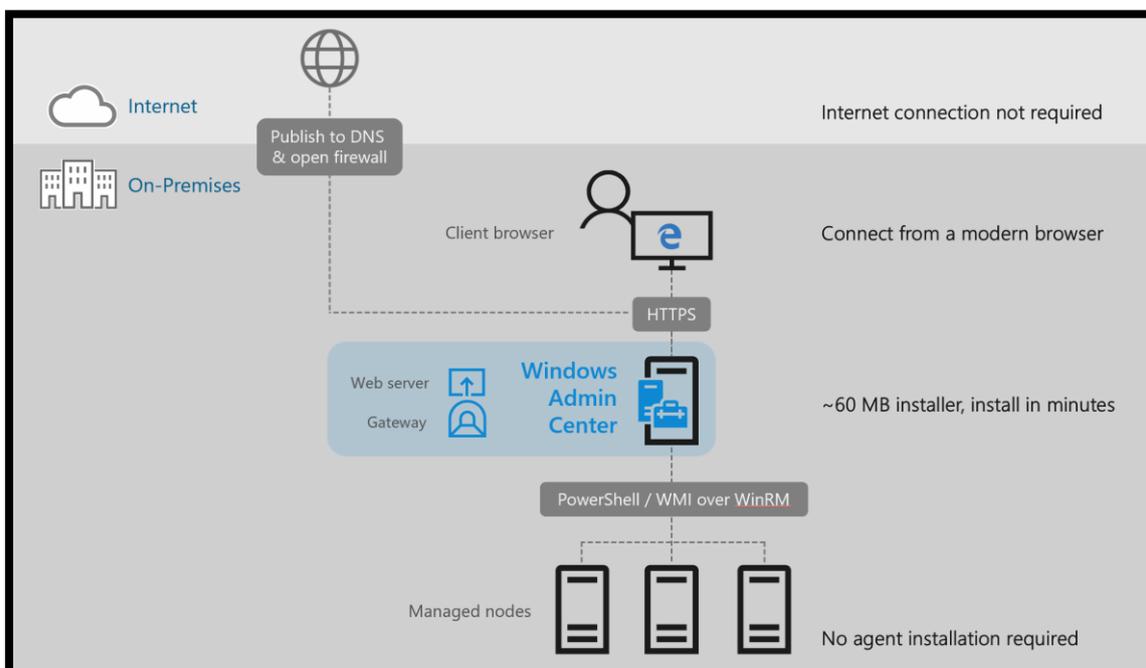
It is a web-based tool and Windows Admin Center uses PowerShell under the hood, and there is a way that you can see those scripts, too, which is pretty cool. It's supported only by the latest releases of the Microsoft Edge or Google Chrome browsers. **! Note ! Microsoft's older browser, Internet Explorer, isn't supported.**

System Requirements:

Installation is possible on Windows 10 Fall Anniversary Update (1709) or newer (**1809 is better**), or Server core (gateway), or on one of the management servers itself – Windows Server 2016 or newer.

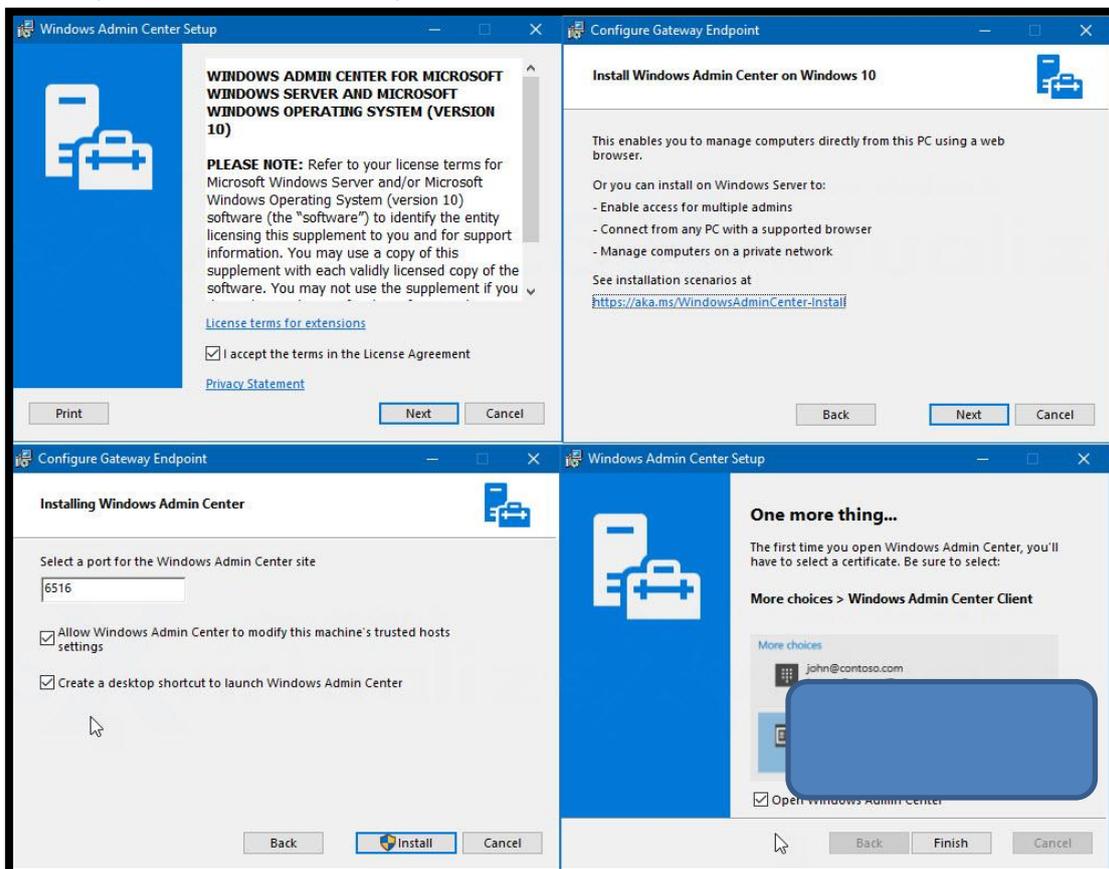
Then, you can manage Windows Server 2008 R2 (with limited functionality), 2012, or 2012 R2, 2016 and, lastly, Windows Server 2019. (Note that you can also manage Windows 10 systems.).

Installation of Windows Management Framework 5.1 is required on those servers. There are no other dependencies.



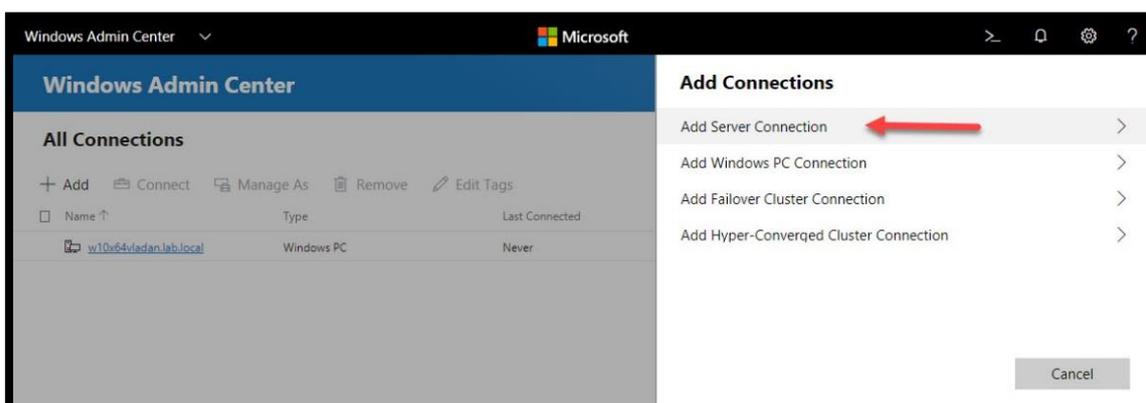
Installation and Configuration of Windows Admin Center

You can download Windows Admin Center from [here](#). You can use Windows Admin Center free of charge. After downloading, start the installation by clicking the MSI. Below, you can see the different installation screens. You have the option to add a shortcut to your desktop or change the default port through which this solution will be accessed.



Open your web browser and type "localhost" or "FQDN" of the machine where the installation has been done. Then, you can access Windows Admin Center (WAC).
https://FQDN of server:port specified in setup'
 ex. <https://wac-admin:6516>

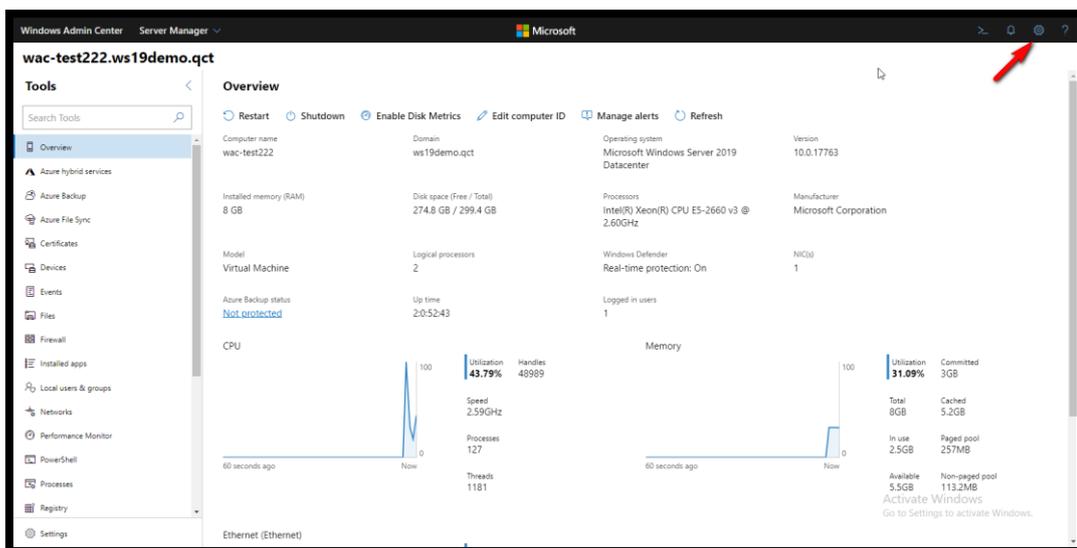
Add your first server



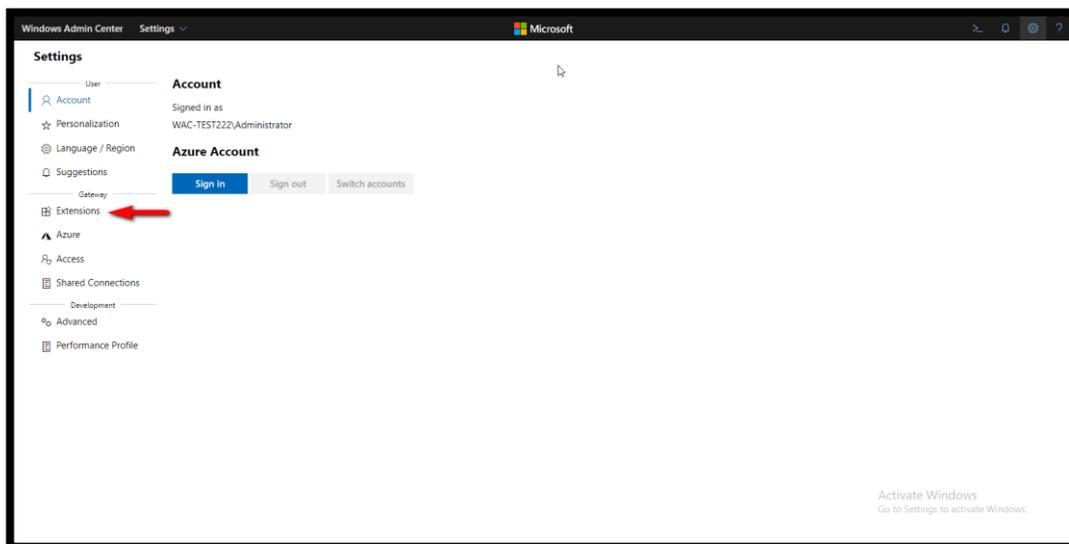
Installing an extension

Windows Admin Center will show extensions available from the specified NuGet feed. By default, Windows Admin Center points to the Microsoft official NuGet feed which hosts extensions published by Microsoft and other developers.

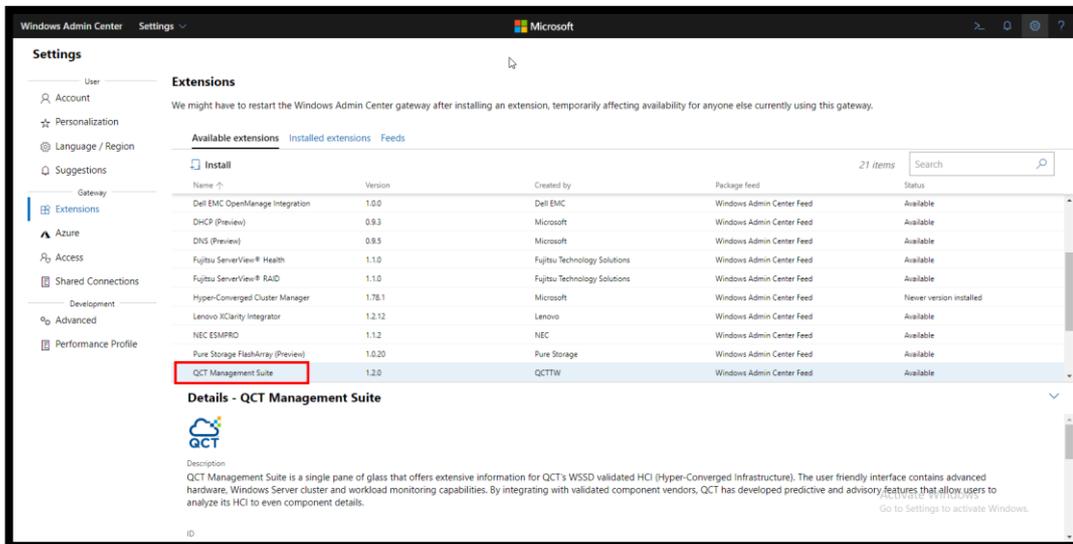
1. Click the Settings button in the top-right > In the left pane, click Extensions.
2. The Available Extensions tab will list the extensions on the feed that are available for installation.



3. Click Install to install an extension. After installation is complete, your browser will automatically be refreshed and Windows Admin Center will be reloaded with the new extension installed.



4. Install the QCT Management Suits



The screenshot shows the Windows Admin Center Settings page for Extensions. A table lists available extensions, with 'QCT Management Suite' highlighted. Below the table, the details for the QCT Management Suite are shown, including its description.

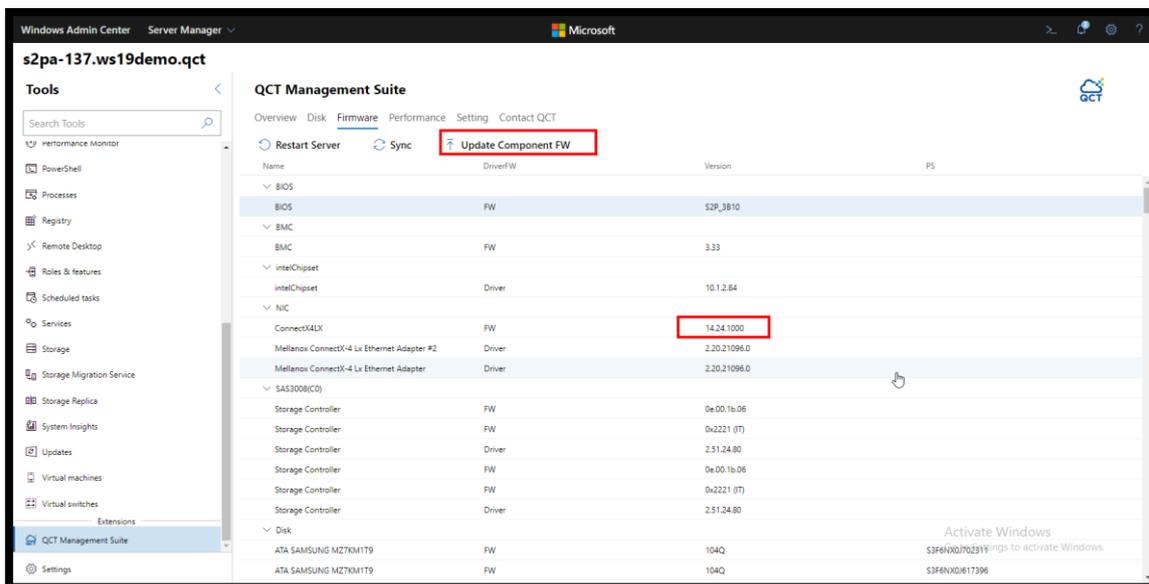
Name	Version	Created by	Package feed	Status
Dell EMC OpenManage Integration	1.0.0	Dell EMC	Windows Admin Center Feed	Available
DHCP (Preview)	0.9.3	Microsoft	Windows Admin Center Feed	Available
DNS (Preview)	0.8.5	Microsoft	Windows Admin Center Feed	Available
Fujitsu ServerView® Health	1.1.0	Fujitsu Technology Solutions	Windows Admin Center Feed	Available
Fujitsu ServerView® RAID	1.1.0	Fujitsu Technology Solutions	Windows Admin Center Feed	Available
Hyper-Converged Cluster Manager	1.78.1	Microsoft	Windows Admin Center Feed	Never version installed
Lenovo XClarity Integrator	1.2.12	Lenovo	Windows Admin Center Feed	Available
NEC ESMPRO	1.1.2	NEC	Windows Admin Center Feed	Available
Pure Storage FlashArray (Preview)	1.0.20	Pure Storage	Windows Admin Center Feed	Available
QCT Management Suite	1.2.0	QCTTW	Windows Admin Center Feed	Available

Details - QCT Management Suite

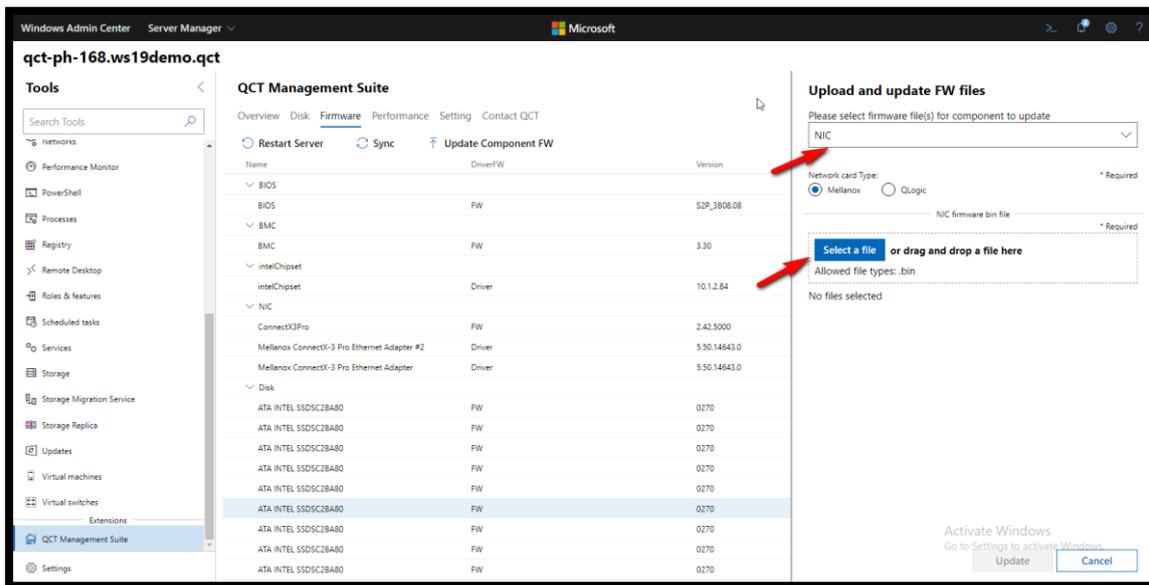
Description
 QCT Management Suite is a single pane of glass that offers extensive information for QCT's WSSD validated HCI (Hyper-Converged Infrastructure). The user friendly interface contains advanced hardware, Windows Server cluster and workload monitoring capabilities. By integrating with validated component vendors, QCT has developed predictive and advisory features that allow users to analyze its HCI to even component details.

Update Firmware and drivers

You can use QCT Management Suites to update Firmware:

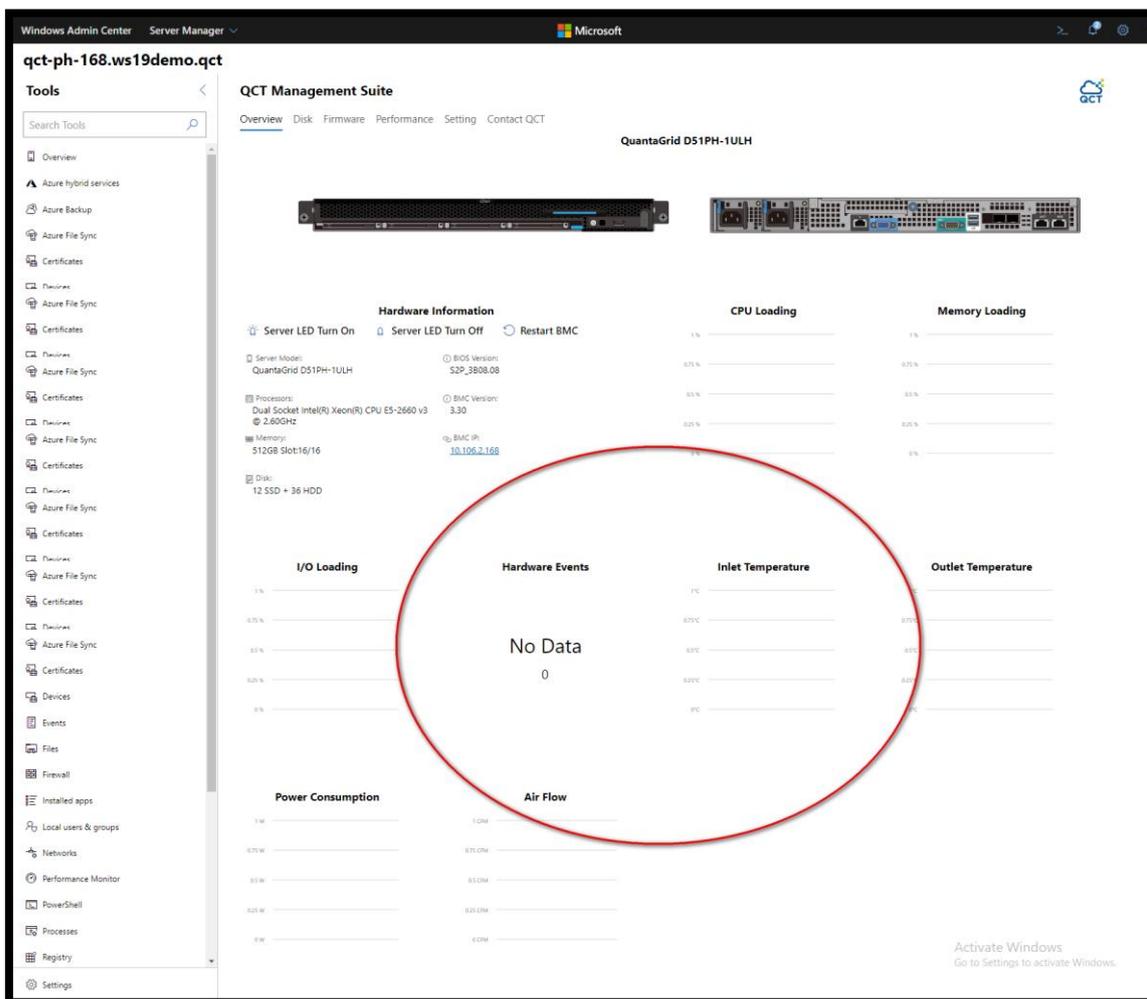
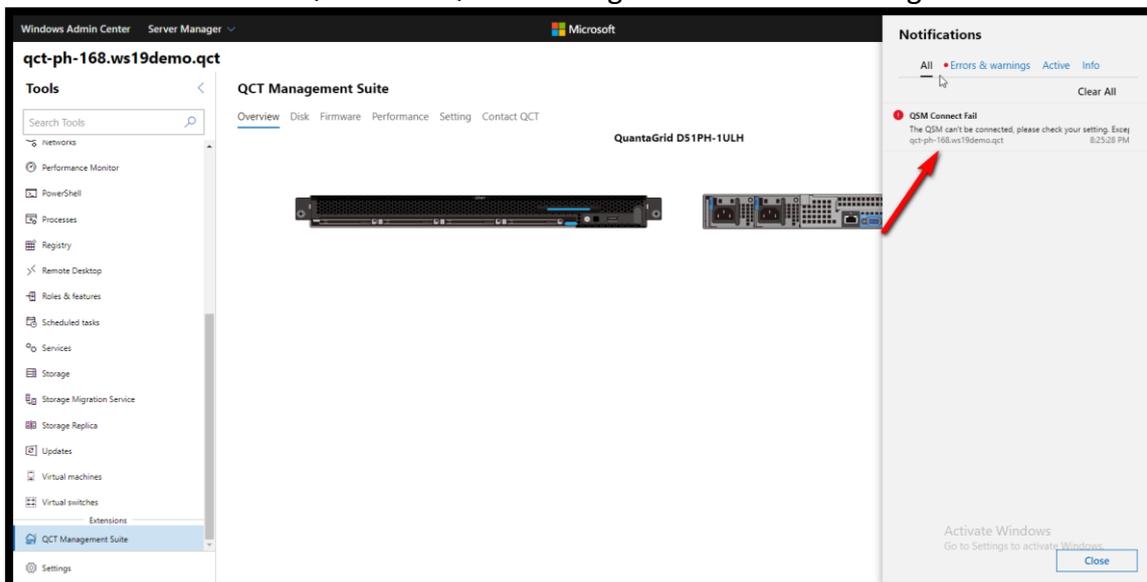


Just select NIC/HBA/Drive then provide firmware file.

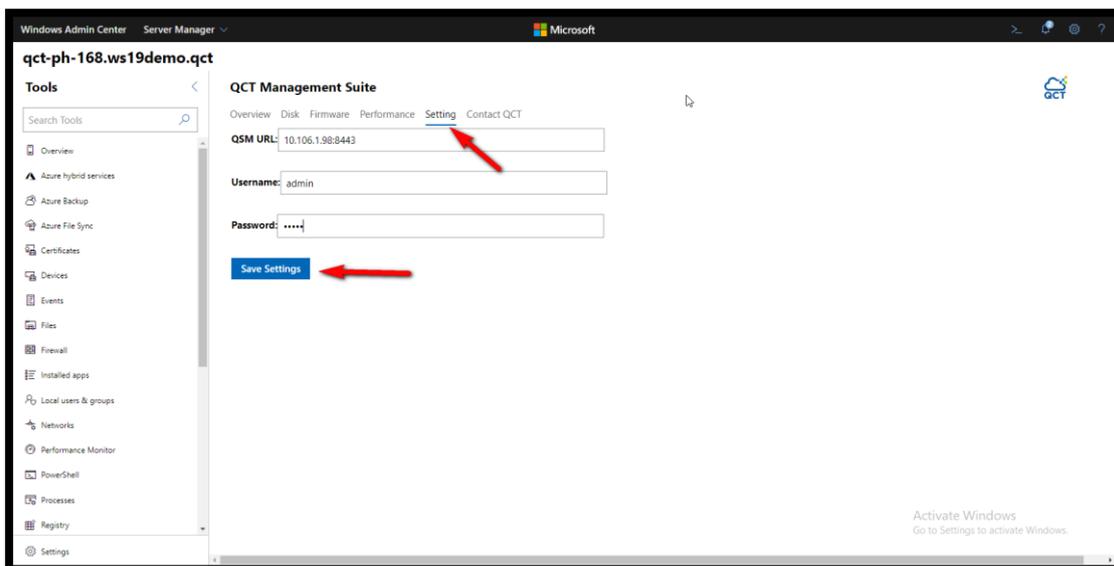


Setting Real-time System Monitor

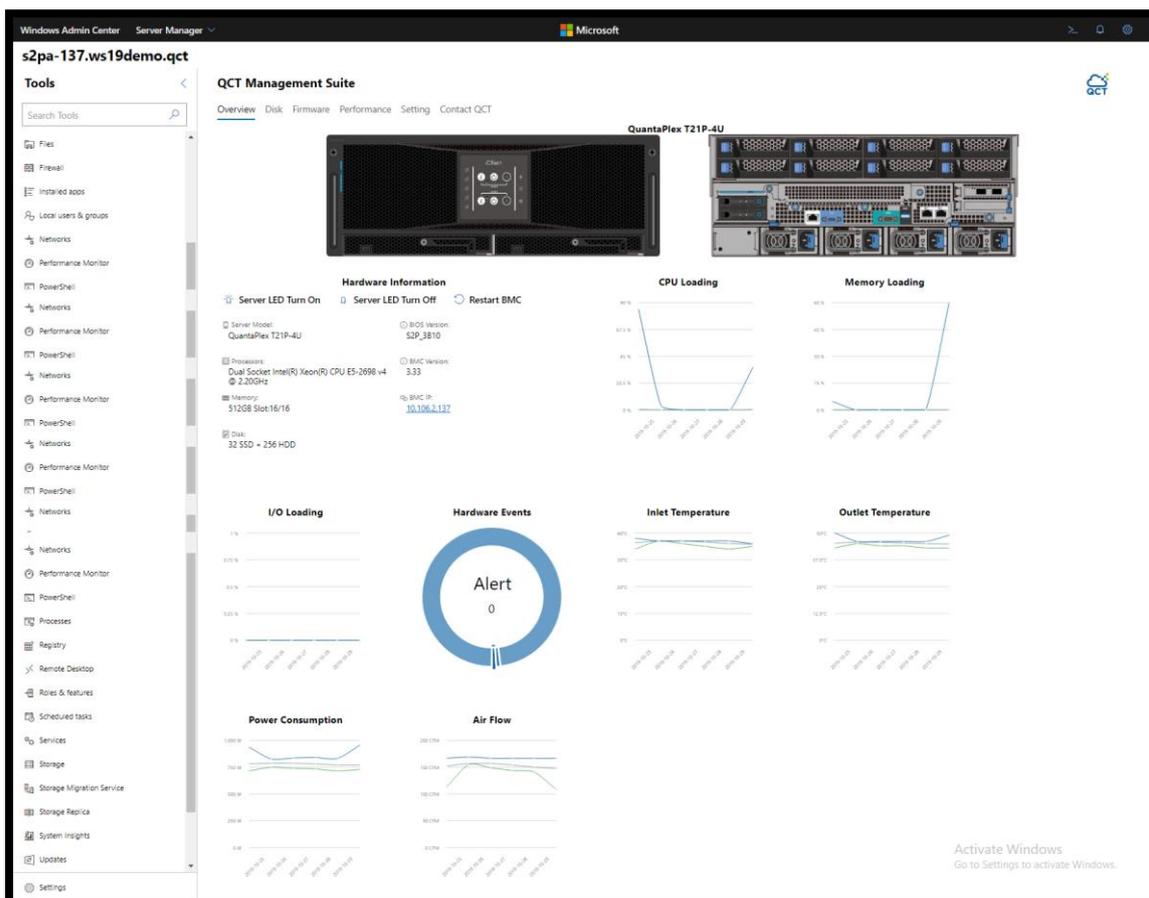
Because not connect to QSM, the real-time monitor will show no data.
You can download the QSM from QCT and integrated hardware management.



Setting the QSM URL in Setting :



If QSM connection successful, the real-time monitor will show data.

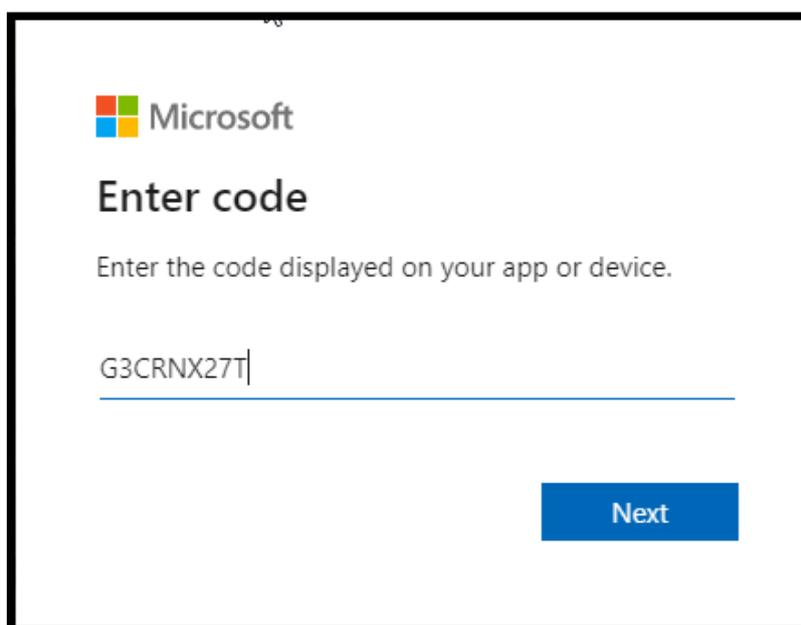
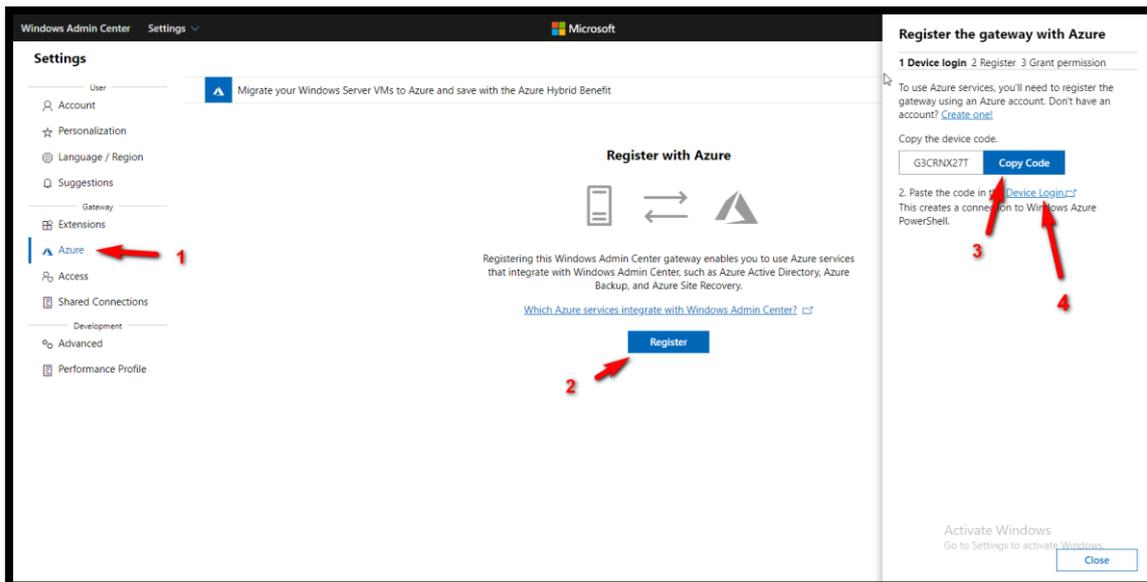


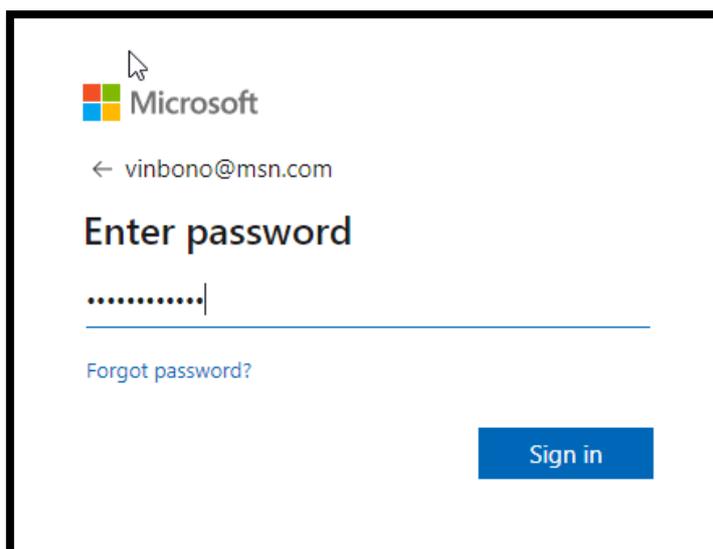
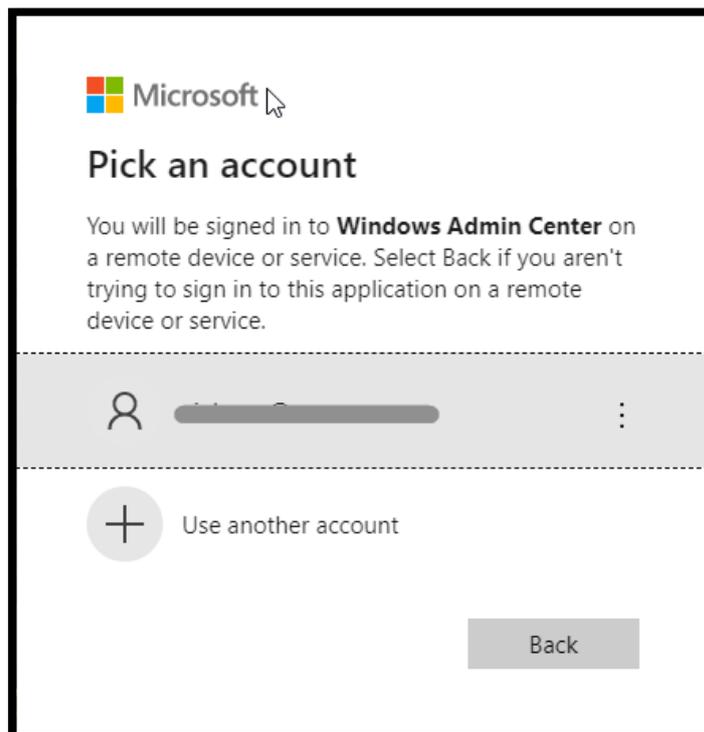


2. Integrated with hybrid storage management by WAC

To allow the Windows Admin Center gateway to communicate with Azure to leverage Azure AD authentication for gateway access, or to create Azure resources on your behalf (for example, to protect VMs managed in Windows Admin Center using Azure Site Recovery), you will need to first register your Windows Admin Center gateway with Azure. You only need to do this once for your Windows Admin Center gateway – the setting is preserved when you update your gateway to a newer version.

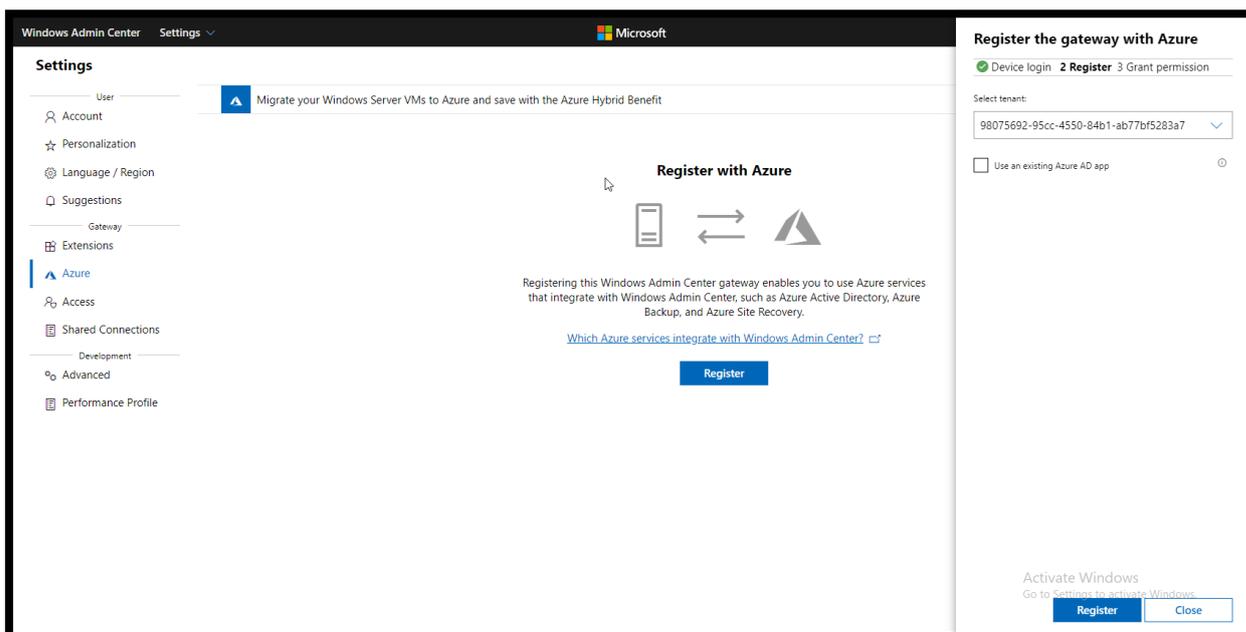
Register WAC gateway with Azure, just copy the device code to login:



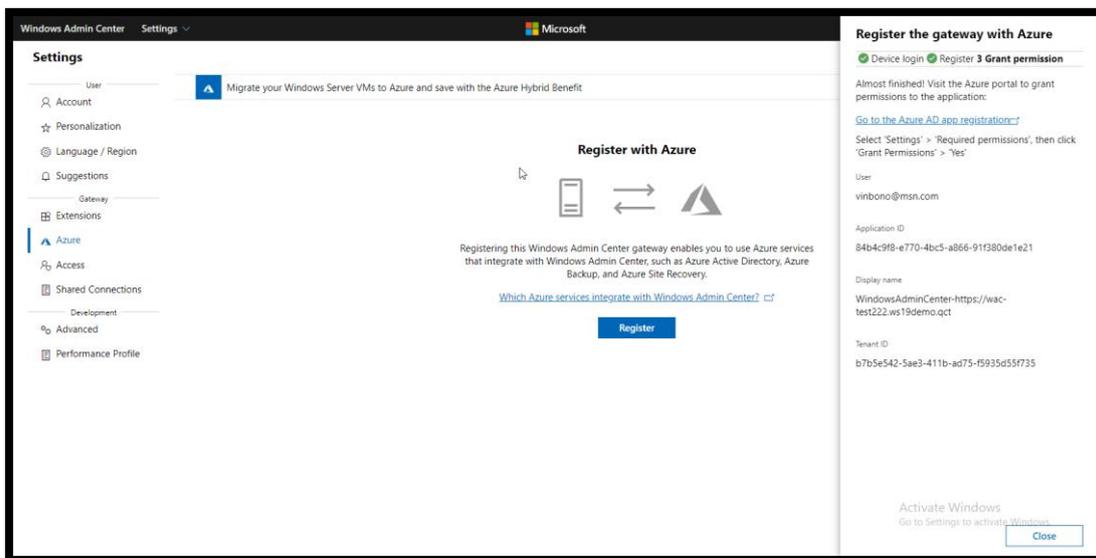




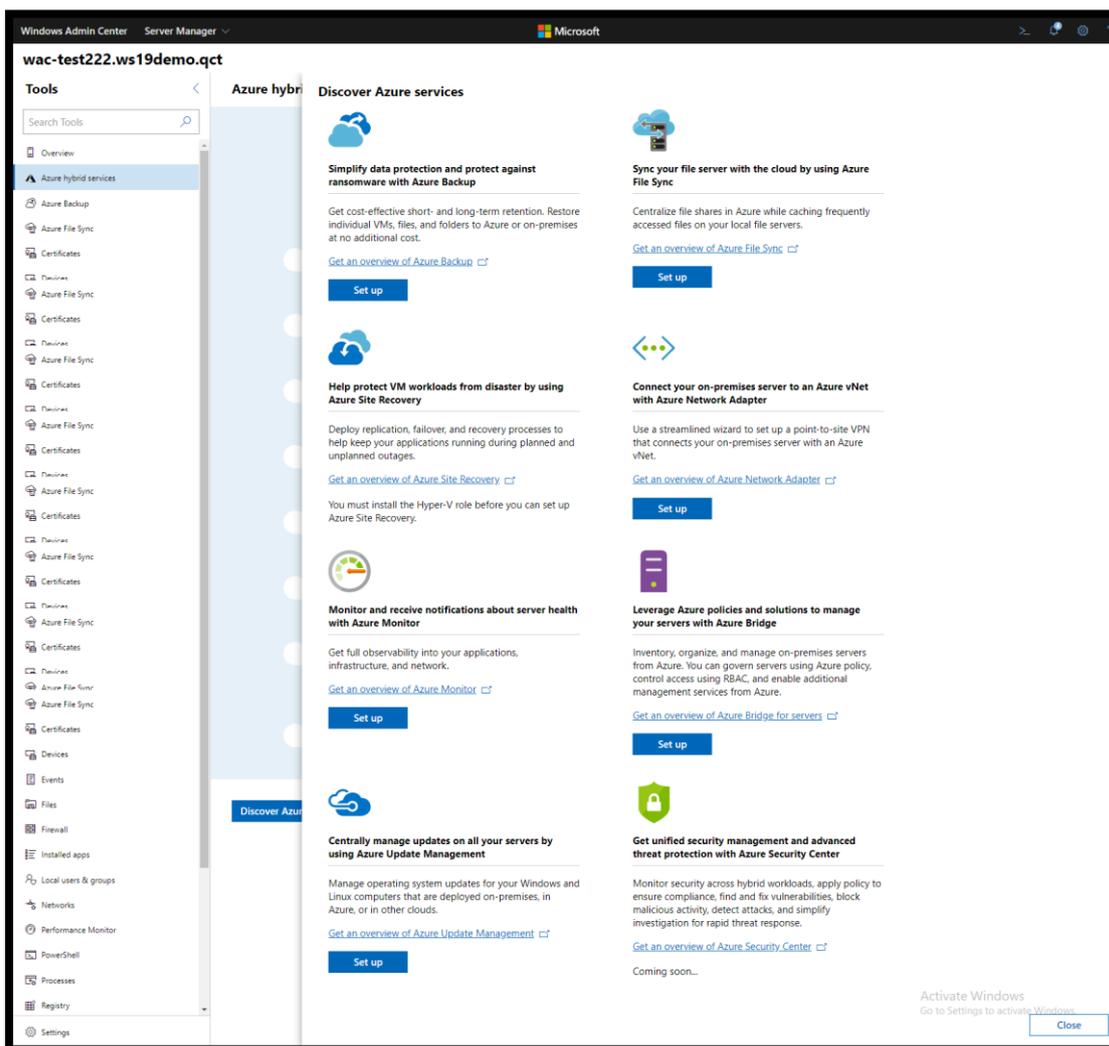
Back to WAC then select tenant, and register:



Finish the registration and grant app permission:



You can explore the Azure hybrid services now.

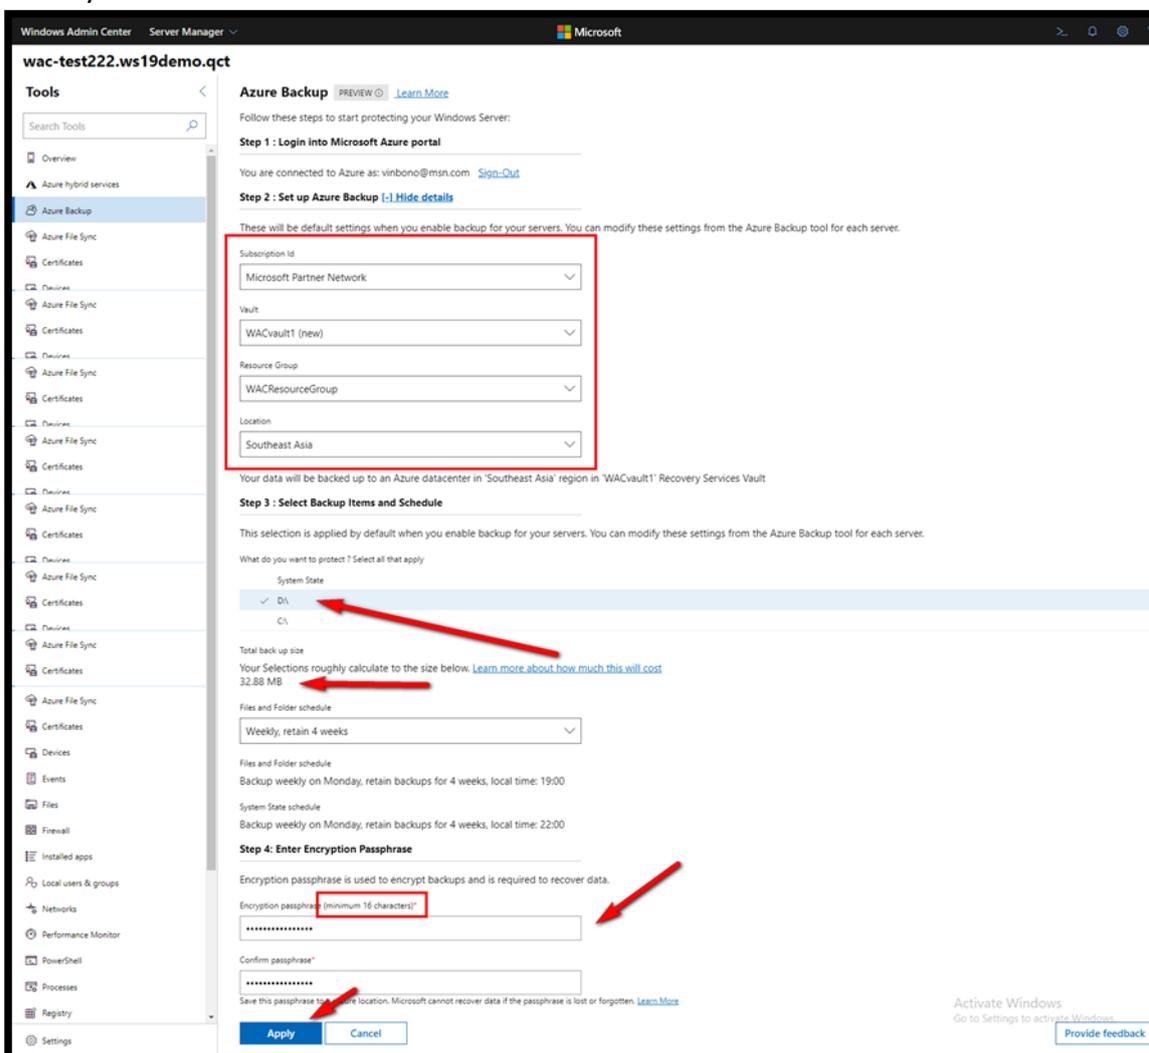


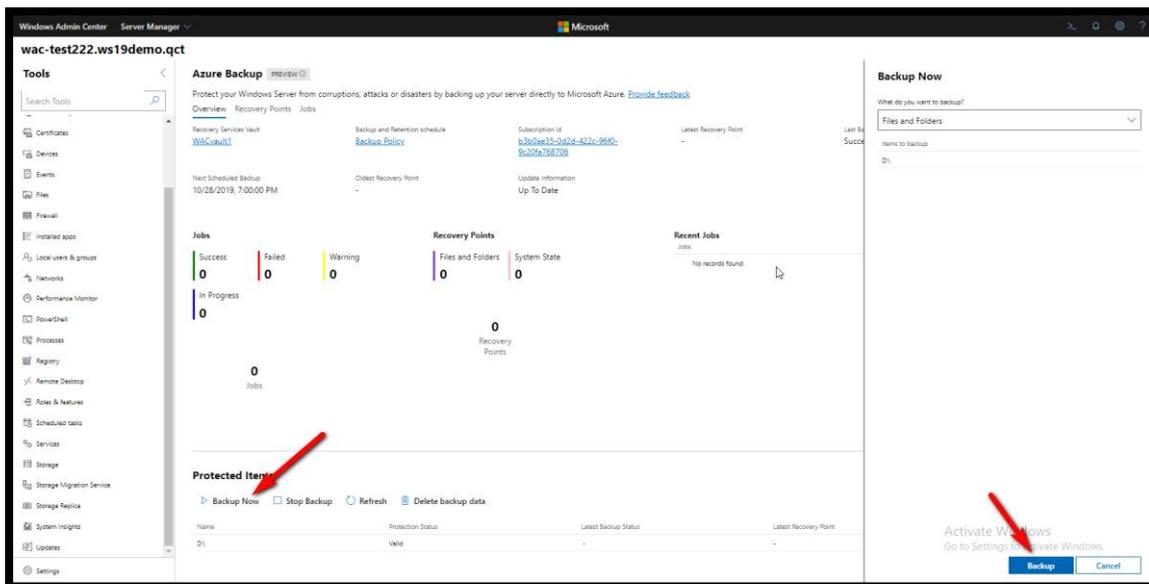
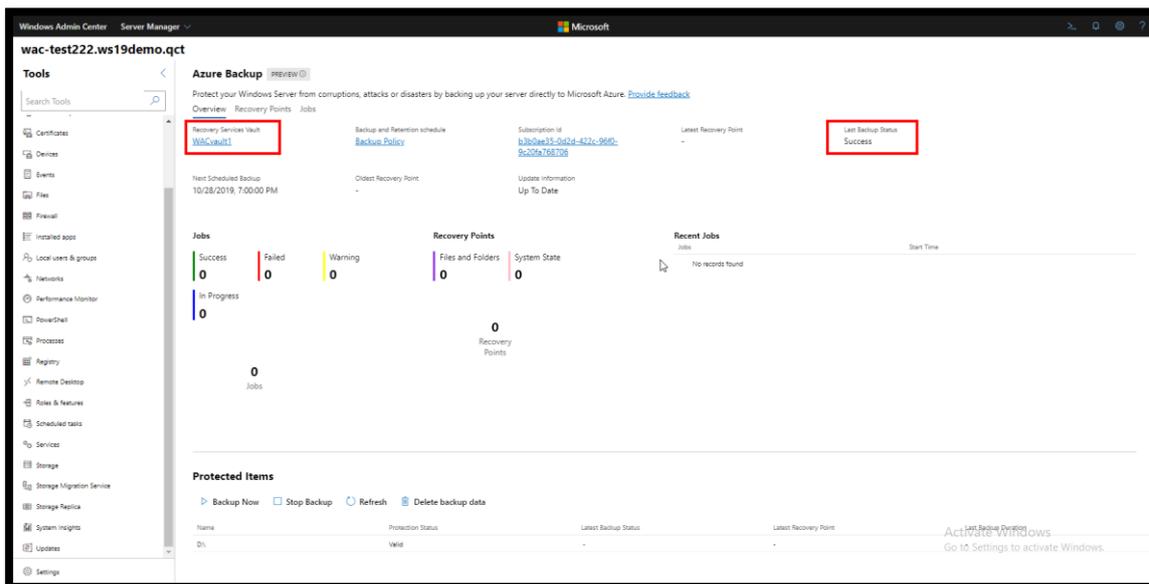
3. Config Azure Backup in Windows Admin Center

With WAC there's the possibility to protect your local machine with Azure Backup without do nothing inside itself; this means no local download, registration and job management. Everything is managed via console.

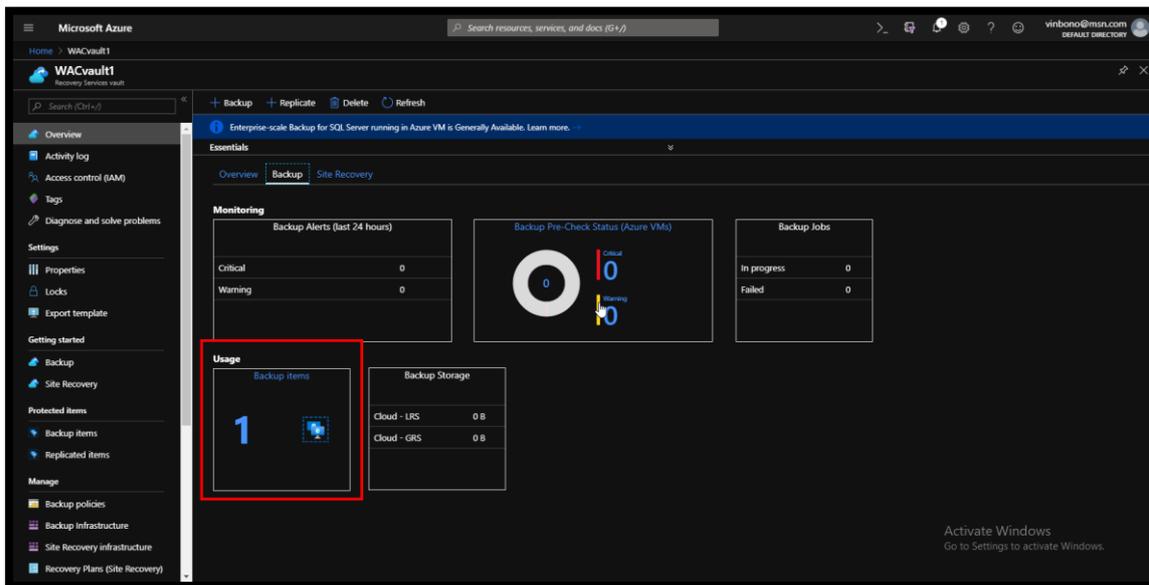
With Microsoft efforts to bring Hybrid Cloud capabilities closer to your on-premises systems, they added support for [Azure Backup](#) in Windows Admin Center. This allows you to simply configure Azure Backup for your Windows Server with a couple of clicks.

First you will need to register your Windows Admin Center to Microsoft Azure. This can be done in the settings of Windows Admin Center. If you haven't done this yet, the wizard will guide you through. After this is done you can go to the Azure Backup Extension in Windows Admin Center and sign in. You can now configure Azure Backup directly in Windows Admin Center.





After Azure Backup is fully configured, you can see the configuration, the latest recovery points and you also will be able to recover data.



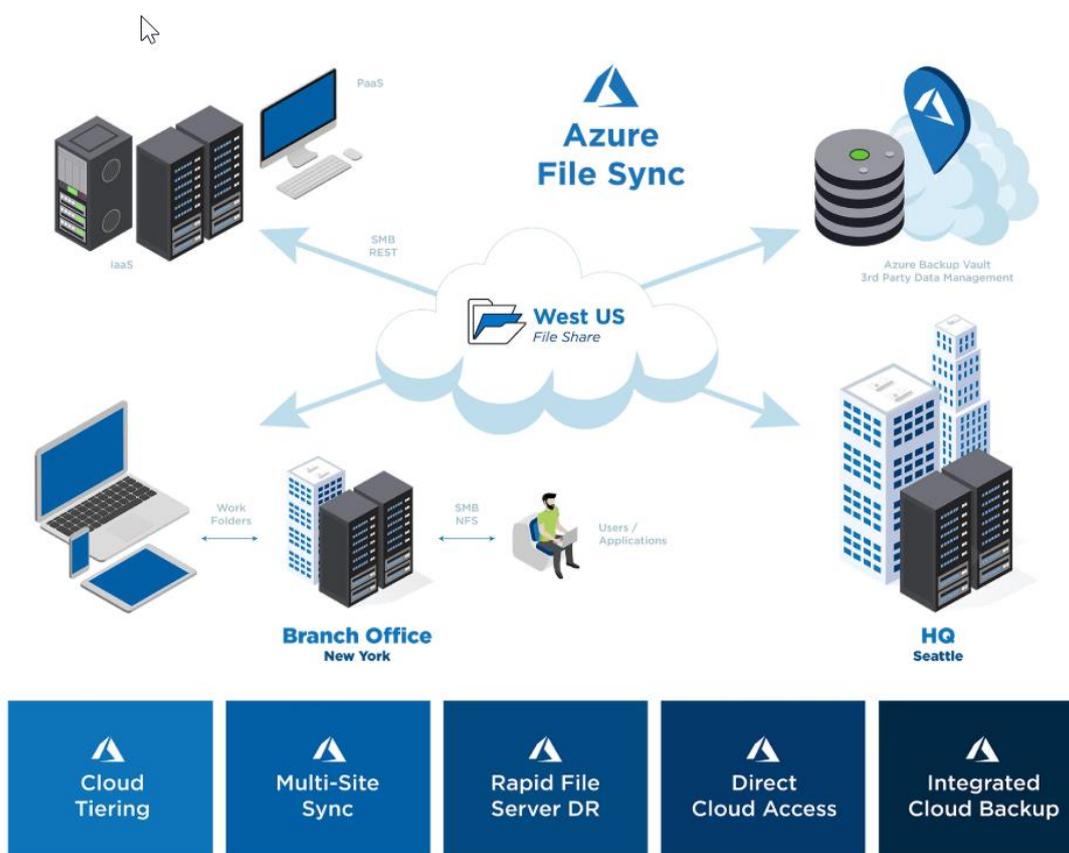
4. Config Azure File Sync in Windows Admin Center

Azure File Sync is a cloud service that allows users to keep aligned one or more folders between different servers in different locations. The Azure File service gives you a fully managed cloud file share and extends the ability of organizations to share files across on-premises and the cloud.

The powerful behind this project is the flexibility: all the information can be reached via local SMB share, via Work Folders but also via remote SMB share. All the files are stored into Azure File Share.

Azure File Sync provides:

- Multi-site access – provide write access to the same data across Windows servers and Azure Files
- Cloud tiering – store only recently accessed data on local servers
- Integrates with Azure backup – no need to back up your data on premises
- Fast disaster recovery – restore file metadata immediately and recall data as needed



Before starts, just for information, is a good idea read the limits about service:

- Storage Sync Services per region: 20 Storage Sync Services
- Sync groups per Storage Sync Service: 100 sync groups
- Registered servers per Storage Sync Service: 99 servers
- Cloud endpoints per sync group: 1 cloud endpoint
- Server endpoints per sync group: 50 server endpoints
- Server endpoints per server: 30 server endpoints
- File system objects (directories and files) per sync group: 25 million objects
- Maximum number of file system objects (directories and files) in a directory: 5 million objects
- Maximum object (directories and files) security descriptor size: 64 KiB
- File size: 100 GiB
- Minimum file size for a file to be tiered: 64 KiB

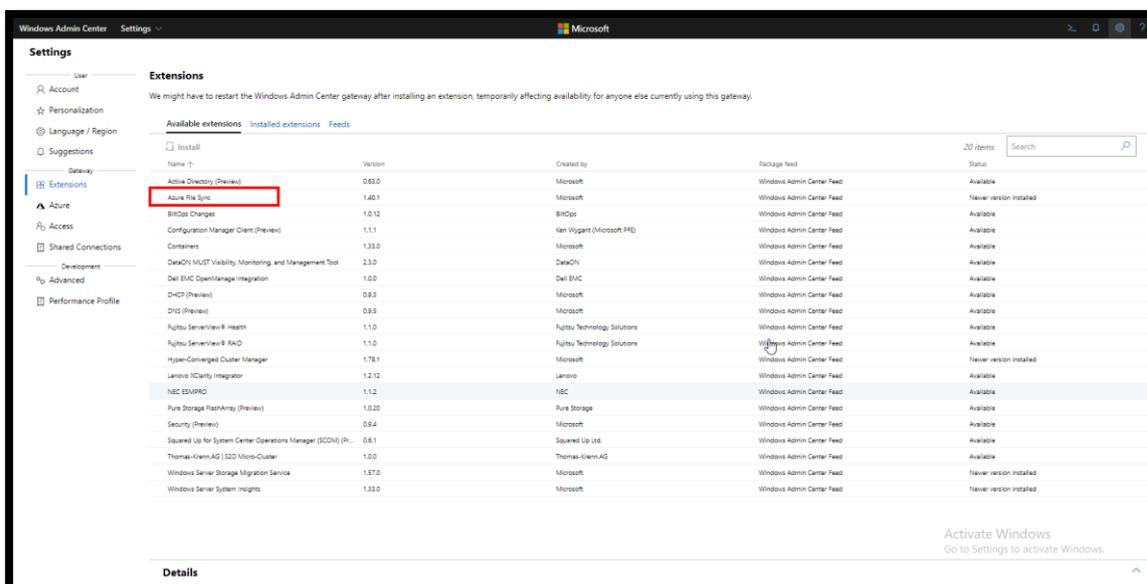
If you are planning to deploy this solution, I highly recommend to read the [planning for an Azure File Sync deployment documentation](#), to make sure you understand the topology, requirements, architecture and limitations. You can also find the documentation to [deploy Azure File Sync](#) without Windows Admin Center on the Microsoft Docs.

High-level steps:

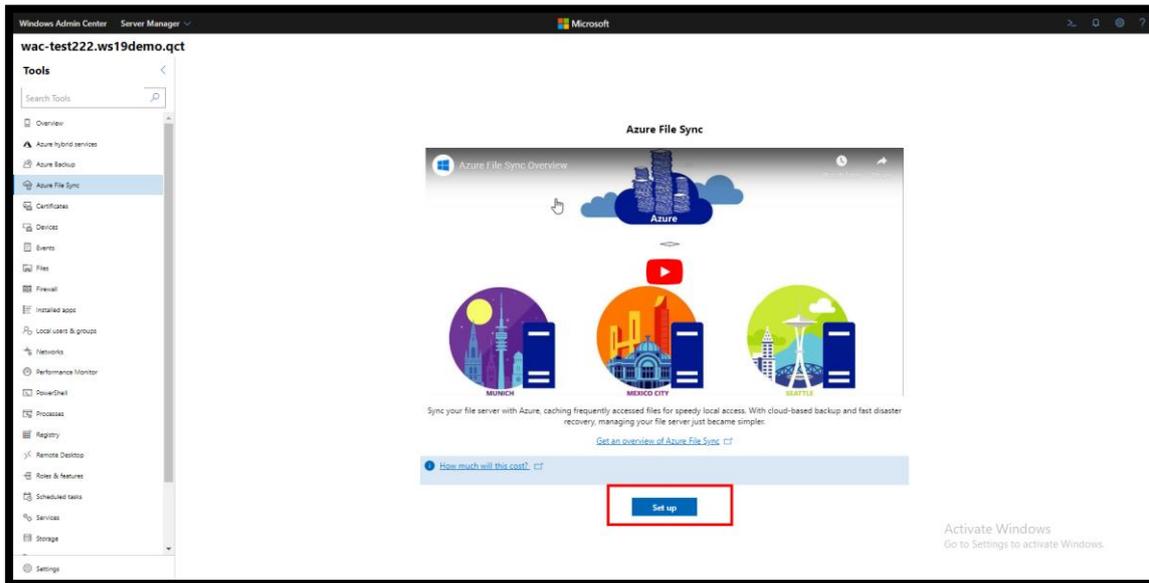
1. Deploy the Storage Sync Service (Azure)
2. Create a sync group and a cloud endpoint (Azure)
3. Install the Azure File Sync Agent (Windows Admin Center)
4. Register Windows Server with Storage Sync Service (Windows Admin Center)
5. Add a server endpoint in the sync group (Azure)

Deploying the Storage Sync Service and creating the sync group with the cloud endpoint, take the same amount of time as you would deploy the solution without Windows Admin Center, follow the steps on the [documentation page](#).

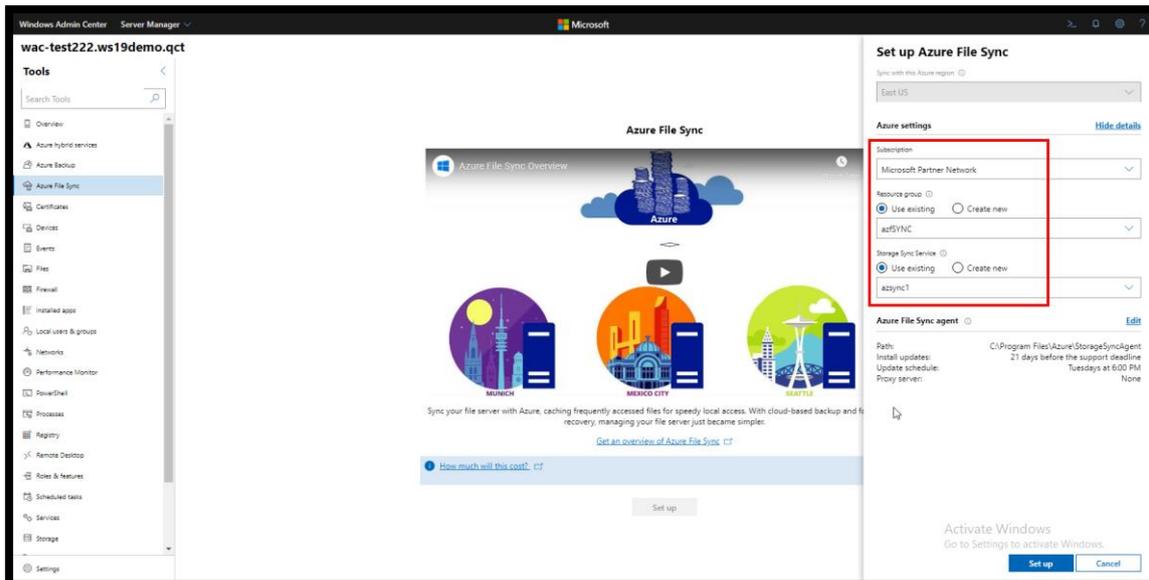
To deploy the agent, open Windows Admin Center and click on Azure File Sync.



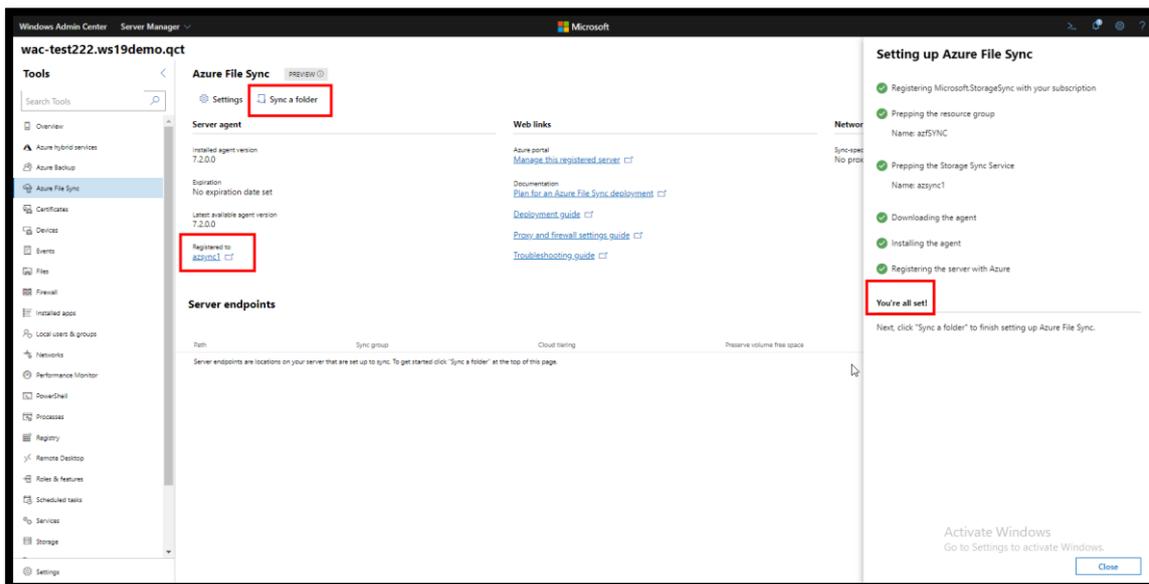
You can click on setup to start the wizard which will walk you through the setup process.



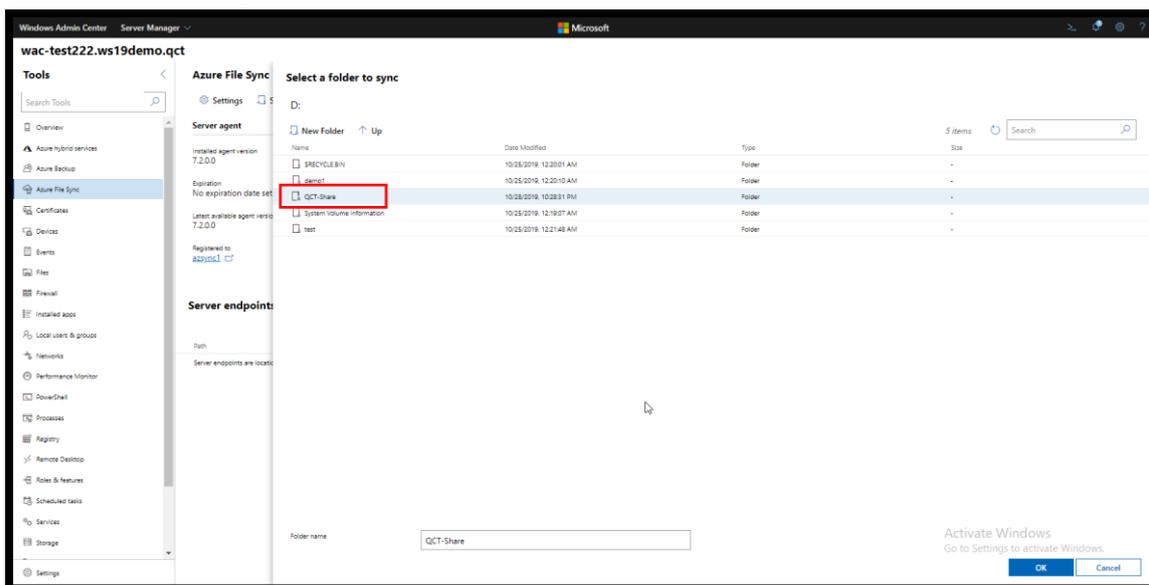
The Server Registration UI should open automatically after installation of the Azure File Sync agent. If it doesn't, you can open it manually from its file location: C:\Program Files\Azure\StorageSyncAgent\ServerRegistration.exe.

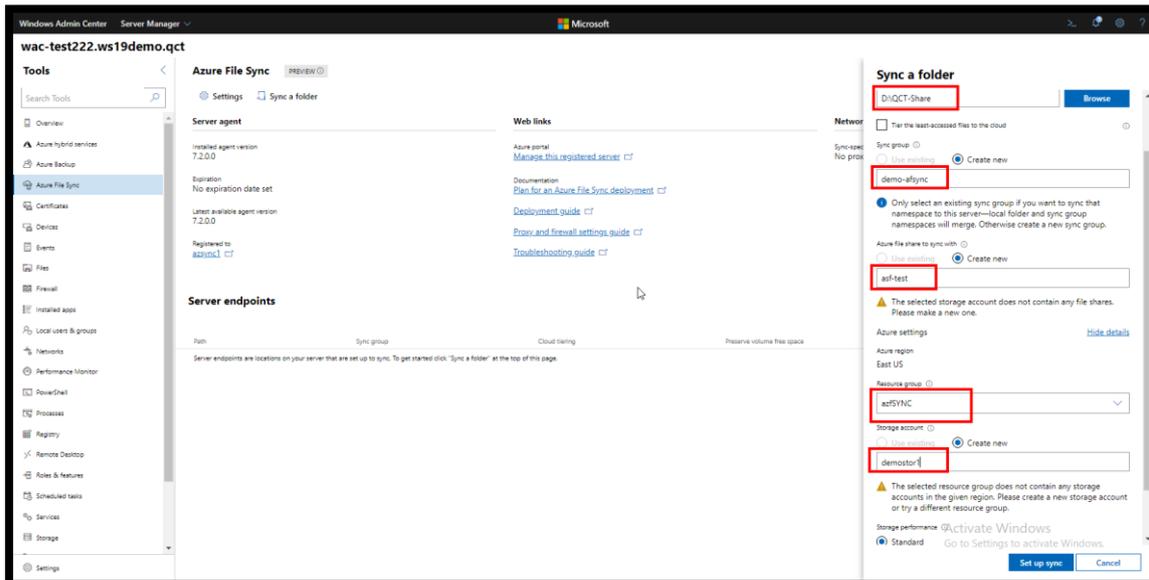


Select Sync a folder and the WAC will help you finish the steps.

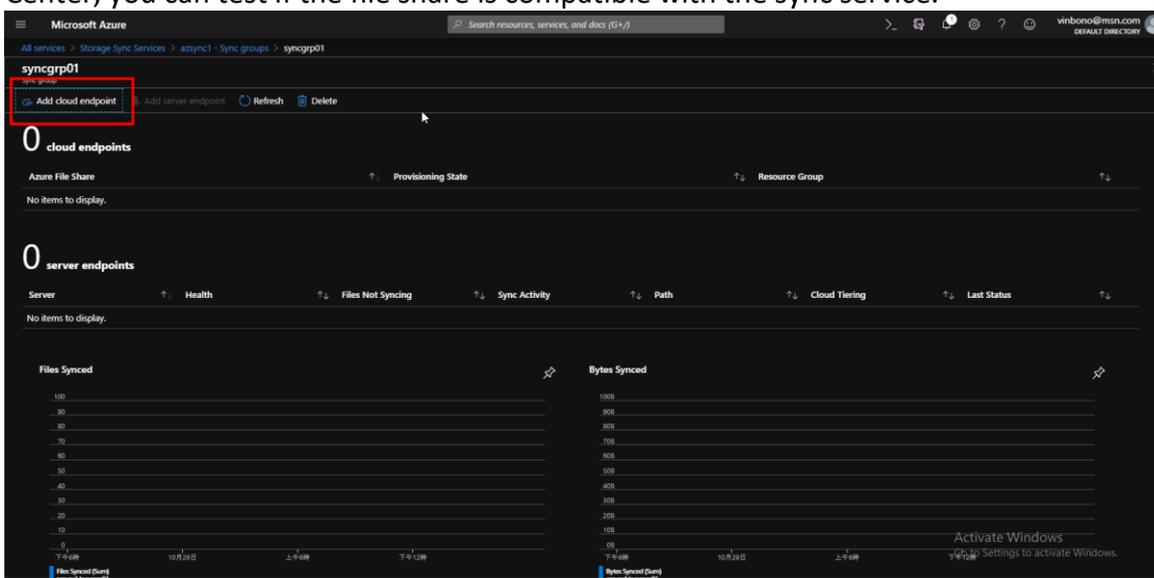


Select a folder to sync:





You can now add the server to the sync group as a server endpoint. In Windows Admin Center, you can test if the file share is compatible with the sync service.





About QCT

QCT (Quanta Cloud Technology) is a global datacenter solution provider extending the power of hyperscale datacenter design in standard and open SKUs to all datacenter customers.



Product lines include servers, storage, network switches, integrated rack systems and cloud solutions, all delivering hyperscale efficiency, scalability, reliability, manageability, serviceability and optimized performance for each workload.

QCT offers a full spectrum of datacenter products and services from engineering, integration and optimization to global supply chain support, all under one roof.

The parent of QCT is Quanta Computer Inc., a Fortune Global 500 technology engineering and manufacturing company.

<http://www.QCT.io>

United States QCT LLC., Silicon Valley office
1010 Rincon Circle, San Jose, CA 95131
TOLL-FREE: 1-855-QCT-MUST
TEL: +1-510-270-6111
FAX: +1-510-270-6161
Support: +1-510-270-6216

QCT LLC., Seattle office
13810 SE Eastgate Way, Suite 190, Building 1,
Bellevue, WA 98005
TEL: +1-425-633-1620
FAX: +1-425-633-1621

China 云达科技, 北京办公室 (Quanta Cloud Technology)
北京市朝阳区东三环中路 1 号 · 环球金融中心东楼 1508 室
Room 1508, East Tower 15F, World Financial Center
No.1, East 3rd Ring Zhong Rd., Chaoyang District, Beijing, China
TEL: +86-10-5920-7600
FAX: +86-10-5981-7958

云达科技, 杭州办公室 (Quanta Cloud Technology)
浙江省杭州市西湖区古墩路浙商财富中心 4 号楼 303 室
Room 303 · Building No.4 · ZheShang Wealth Center
No. 83 GuDun Road, Xihu District, Hangzhou, Zhejiang , China
TEL: +86-571-2819-8660

Japan Quanta Cloud Technology Japan 株式会社
日本国東京都港区芝大門二丁目五番八号
牧田ビル 3 階
Makita Building 3F, 2-5-8, Shibadai-mon ,
Minato-ku, Tokyo 105-0012, Japan
TEL: +81-3-5777-0818
FAX: +81-3-5777-0819

Taiwan 雲達科技 (Quanta Cloud Technology)
桃園市龜山區文化二路 211 號 1 樓
1F, No. 211 Wenhua 2nd Rd., Guishan Dist.,
Taoyuan City 33377, Taiwan
TEL: +886-3-286-0707
FAX: +886-3-327-0001

Other regions Quanta Cloud Technology
No. 211 Wenhua 2nd Rd., Guishan Dist.,
Taoyuan City 33377, Taiwan
TEL: +886-3-327-2345
FAX: +886-3-397-4770

All specifications and figures are subject to change without prior notice. Actual products may look different from the photos.

QCT, the QCT logo, Rackgo, Quanta, and the Quanta logo are trademarks or registered trademarks of Quanta Computer Inc.

All trademarks and logos are the properties of their representative holders.

Copyright © 2014-2015 Quanta Computer Inc. All rights reserved.