
# Cisco vWAAS on Microsoft Hyper-V

This chapter describes how to install and use Cisco vWAAS on the Microsoft Hyper-V hypervisor.

This chapter contains the following sections:


- About Cisco vWAAS on Microsoft Hyper-V
- vWAAS on Hyper-V Deployments
- vWAAS on Hyper-V Requirements
- Installing vWAAS on Hyper-V
- Activating and Registering vWAAS on Hyper-V
- Operating Considerations for vWAAS on Hyper-V


## About Cisco vWAAS on Microsoft Hyper-V

Cisco vWAAS on Microsoft Hyper-V extends Cisco networking benefits to Microsoft Windows Server Hyper-V deployments.

Hyper-V is a hypervisor-based server-virtualization product that improves utilization, consolidates server workloads, and reduces costs. To achieve this, vWAAS on Hyper-V uses hardware virtualization to enable multiple operating systems to run on a single host, and allows the operating systems to share the same underlying physical hardware.

vWAAS on Hyper-V supports all the WAN-optimization functionality that is supported by physical WAAS devices. Physical memory for vWAAS on Hyper-V is provided by a Cisco UCS server.

## vWAAS on Hyper-V Deployments

You can deploy vWAAS on Hyper-V as a standalone role or as an installable product:

- Standalone role in the Hyper-V server—Hyper-V Server 2012 or Hyper-V Server 2012 R2
- In installable product in the Windows server—Windows Server 2008 R2, Windows Server 2012 or Windows Server 2012 R2.

This section contains the following topics:


- Operating Guidelines for vWAAS on Hyper-V
- Platforms Supported for vWAAS on Hyper-V
- Interoperability Support

# Operating Guidelines for vWAAS on Hyper-V

**Caution** Multiple deployments of vWAAS on the same Hyper-V host *in parallel* may cause unexpected results, due to availability of free space when creating VHDs. We recommend that you do *not* deploy multiple vWAAS on Hyper-V in parallel, unless you have verified that you have enough free disk space required for the respective vWAAS models.

To ensure reliable throughput with the following configuration—**vWAAS on Windows Server 2012 R2 Hyper-V in Cisco UCS-E Series 160S-M3**—we recommend that you do the following:

- Upgrade to the latest UCS-E firmware (Version 3.1.2), available on the Cisco Download Software Page for UCS E-Series Software, UCS E160S M3 Software.
- Verify that you have installed the critical Windows Server updates, available on the Microsoft Windows RT 8.1, Windows 8.1, and Windows Server 2012 R2 Update Rollup page. You can also obtain the standalone update package through the Microsoft Download Center by searching for **KB2887595**.

# Platforms Supported for vWAAS on Hyper-V

This section shows the platforms supported for vWAAS on Microsoft Hyper-V:

- Table 5-1 shows vWAAS support for the Microsoft Hyper-V servers and SCVMM (System Center Virtual Machine Manager).
- Table 5-2 shows platforms supported for vWAAS on Microsoft Hyper-V, deployed as a standalone or installable product.

***Table 5-1 vWAAS Support for Microsoft Hyper-V Servers and SCVMM***

| Microsoft Hyper-V Server | Microsoft SCVMM | vWAAS Supported |
|---|---|---|
| Microsoft Hyper-V Server 2008 | SCVMM 2008 | No |
| Microsoft Hyper-V Server 2008 R2 | SCVMM 2008 R2 | No |
| Microsoft Hyper-V Server 2008 R2 | Microsoft System Center 2012 VMM 2012 or VMM 2012 R2 | Yes |
| Microsoft Hyper-V Server 2012 | Microsoft System Center 2012 VMM 2012 or VMM 2012 R2 | Yes |
| Microsoft Hyper-V Server 2012 R2 | Microsoft System Center 2012 VMM 2012 or VMM 2012 R2 | Yes |

**Note** If you want to install SCVMM in Windows 2008 R2, you must first register it to Windows 2012 or Windows 2012 R2.

***Table 5-2 Platforms Supported for vWAAS on Microsoft Hyper-V, as a Standalone or Installable Product***

| *Standalone Product in Hyper-V Server* | | *Installable Product in Windows Server* |
|---|---|---|
| **Hyper-V Server 2008 R2** | **Hyper-V Server 20012 or 2012 R2** | **Windows Server 2012 or 2012 R2** |
| UCS E-Series and UCS servers | UCS E-Series and UCS servers | UCS E-Series and UCS servers |
| vCM-100 | vCM-100 | vCM-100 |
| vCM-500 | vCM-500 | vCM-500 |
| vCM-1000 | vCM-1000 | vCM-1000 |
| vCM-2000 | vCM-2000 | vCM-2000 |
| vWAAS-150<br>(For WAAS Version 6.2.1 and later, supported on Cisco EHWIC and NIM.) | vWAAS-150<br>(For WAAS Version 6.2.1 and later, supported on Cisco EHWIC and NIM.) | vWAAS-150<br>(For WAAS Version 6.2.1 and later, supported on Cisco EHWIC and NIM.) |
| vWAAS-200 | vWAAS-200 | vWAAS-200 |
| vWAAS-750 | vWAAS-750 | vWAAS-750 |
| vWAAS-1300 | vWAAS-1300 | vWAAS-1300 |
| vWAAS-2500 | vWAAS-2500 | vWAAS-2500 |
| vWAAS-6000 | vWAAS-6000 | vWAAS-6000 |
| vWAAS-12000 | vWAAS-12000 | vWAAS-12000 |
| | vWAAS-50000 | vWAAS-50000 |

## Interoperability Support

You can configure VM on Hyper-V as virtual WAAS Central Manager (vCM) or as vWAAS:

- The Hyper-V device configured as vCM has the same functionality as WAAS Central Manager, and can manage any other device managed by WAAS Central Manager.
- The Hyper-V device configured as vWAAS has the same functionality as the non-Hyper-V vWAAS. Physical memory for vWAAS on Hyper-V is provided by the UCS server.

# vWAAS on Hyper-V Requirements

This section contains the following topics:

- System Infrastructure Requirements
- Hardware Virtualization

## System Infrastructure Requirements

Your WAAS system must have the following to deploy vWAAS on Hyper-V:

- Microsoft Hyper-V Hypervisor—Hypervisor enables multiple operating systems to run on a single host. vWAAS runs as a guest on any host running Hyper-V 2008 R2 or greater.

- Hyper-V Virtual Switch—The Hyper-V Virtual Switch is a software-based Layer 2 switch that connects virtual machines to both virtual networks and the physical network. It provides policy enforcement for security, isolation, and service levels, and includes features for tenant isolation, traffic shaping, simplified troubleshooting, and protection against malicious virtual machines.

  Hyper-V Virtual Switch is available in Hyper-V Manager when you install the Hyper-V server role.

## Hardware Virtualization

This section describes vWAAS on Hyper-V hardware virtualization requirements for CPU, disk, CD-ROM, and flash.

- CPU—vWAAS on Hyper-V supports 2, 4, and 8 CPU configurations. vWAAS on Hyper-V does require a minimum CPU limit.

  **Note** vWAAS VM (Virtual Machine) with different CPU configurations works, but is not recommended.

- Disk sizes for vWAAS on Hyper-V— Disk sizes for vWAAS on Hyper-V are the same as those for ESXi, for each model. For more information on disk sizes for WAAS versions up to v6.x, see Table 1-18, "vCPUs, ESXi Server Datastore Memory, and Disk Space by vWAAS Model".
- CD-ROM—vWAAS on Hyper-V supports standard ISO image file for its CD-ROM device.
- Flash—Unlike physical WAAS devices, vWAAS on Hyper-V does not have access to a separate flash device. Instead, vWAAS flash is installed on the first hard disk, and also uses this first disk for booting. A separate larger disk hosts the DRE/CIFS caches, etc. Other flash functionalities are supported as in ESXi.

# Installing vWAAS on Hyper-V

vWAAS on Hyper-V is installed using the Microsoft Virtual Machine Manager (VMM), with the Virtual Hard Disk (VHD) file. During installation, there is an option to import pre-configured and pre-installed vWAAS images to Hyper-V. After you have completed installation, complete the activation and registration process with the procedures described in Activating and Registering vWAAS on Hyper-V.

This section contains the following topic:

- Installing vWAAS on Hyper-V with a VHD Template
- Configuring GPT Disk Format for vWAAS-50000 on Hyper-V with Akamai Connect

## Installing vWAAS on Hyper-V with a VHD Template

There are seven VHD templates available for vWAAS, and four VHD templates available for vCM.

You can import a pre-configured, model-based VHD file for your deployment. For more information on installing Hyper-V with a VHD template, contact your Cisco account representative.

To install vWAAS on Hyper-V with a VHD template, follow these steps:

**Step 1** Download the vWAAS package to the computer where the SCVMM2012 or the 2012 R2 console is installed.

**Step 2** Unzip the vWAAS package.

**Step 3** Login to the SCVMM console.

**Step 4** Launch the PowerShell window that is displayed in the SCVMM.

**Step 5** Navigate to the PowerShell script in the uncompressed vWAAS package: ".\Cisco-vWAAS-model-name-6.0.0-ISO\Cisco-vWAAS-model-name-6.0.0-ISO"

**Step 6** Run the PowerShell script: "deploy-vwaas-*model-name*"

**Step 7** Follow the procedure that is requested by the deployment script.

**Step 8** If your deployment uses a vWAAS-12000 or vWAAS-50000 model, you must enter a maximum amount of memory in NUMA (Non-Uniform Memory Access) configuration of at least RAM size or higher, in MB, otherwise the device will not be able to boot up.

**Note** Entering the maximum memory amounts as shown in Step 9 should be completed *only after* you have deployed vWAAS in Hyper-V (as shown in Step 1 through Step 7).

**Step 9** To enter the maximum amount of memory, follow these steps:

a. From the SC VMM console, navigate to **Hardware > Processor > NUMA**.

b. The NUMA Configuration screen is displayed.

c. At the **Maximum amount of memory (MB)** field, enter an amount, in MB:

- For vWAAS-12000, enter an amount of at least 12288 MB.
- For vWAAS-50000, enter an amount of at least 49152 MB.

# Configuring GPT Disk Format for vWAAS-50000 on Hyper-V with Akamai Connect

The following list shows the disk requirements for vWAAS on Hyper-V for vWAAS-50000 with Akamai Connect:

- 4 GB Flash
- 48 GB Kdump
- 1500 GB
- 850 GB for disk (for Akamai Connect)

The Windows server does not detect disk size more than 2 TB in partition **C:** because it is in MBR format. Therefore, in order to have a disk size more than 2 TB, you need to create partition **D:** in GPT (GUID Partition Table) format.

To convert the HDD from MBR format to GPT format, follow these steps:

**Step 1** Install windows in one partition of the HDD.

**Step 2** After installation is complete, create a new volume to create a new disk partition:

a. Right-click the Windows command prompt and then click **Run as Administrator**.

b. Enter the **diskpart** command to enter DiskPart command mode.

c. At the DISKPART prompt, enter the **create volume** command to create a new volume on the disk.

**Step 3** At the DISKPART prompt, enter the **list disk** command to display a list of disks and associated information (including size, available free space, whether the disk is basic or dynamic).

**Step 4** Note the disk number of the disk for which you want to convert formats.

**Step 5** At the DISKPART prompt, enter the **select disk** *disk-number* command.

**Step 6** At the DISKPART prompt, enter the **clean** command to specify that all sectors on the disk are set to zero.

**Note** The **clean** command deletes all data on the disk.

**Step 7** At the DISKPART prompt, enter the **convert gpt** command to convert the disk format to GPT format.

**Step 8** With the GPT format, you can configure RAID capabilities for the HDD, including logical disk handling with RAID-5, logical disk handling with RAID-1, and disk hot-swap support. For more information on RAID support for Cisco WAAS, see the *Cisco Wide Area Application Services Configuration Guide*.

# Activating and Registering vWAAS on Hyper-V

You manage vWAAS on Hyper-V through the WAAS Central Manager (CM). vWAAS on Hyper-V supports all the functionality that is supported by WAAS devices.

This section describes how to activate and register vWAAS on Hyper-V. For information on installing vWAAS on Hyper-V, see Installing vWAAS on Hyper-V.

When a Hyper-V vWAAS virtual machine (VM) is started on the Hyper-V, it boots up and prompts you to enter basic boot configuration information, including configuring a Hyper-V interface and WAAS CM IP address.

To activate and register vWAAS on Hyper-V, following these steps:

**Step 1** Configure the IP address/gateway on the vWAAS interface. As needed, also configure *name-server*, *domain-name*, and any other static routes.

**Step 2** If necessary, configure WCCP interception. For more information on configuring WCCP interception, see WCCP Interception. No configuration is necessary for appnav-controller interception.

**Step 3** Configure the WAAS Central Manager IP address so that vWAAS can be registered with the WAAS Central Manager.

**Step 4** Hyper-V vWAAS connects with the WAAS CM and registers itself. Hyper-V vWAAS is considered in service after it is registered successfully and it optimizes the connections.

**Step 5** The following are scenarios when a vWAAS cannot not successfully register with the WAAS CM:

- If Hyper-V vWAAS cannot register with the WAAS CM, it generates an alarm and does not optimize connections. Contact Cisco Technical Support (TAC) if you need assistance to resolve this situation.

- Hyper-V vWAAS may register successfully with the WAAS CM, but lose connectivity due to a shutdown or power off. If it remains functional, vWAAS will continue to optimize connections in the offline state.
- If you de-register the Hyper-V vWAAS (with the **cms deregister** EXEC command), it is removed from service.

**Step 6** After vWAAS on Hyper-V is operational on a device, the WAAS CM displays the following information for the device:

- The Hyper-V device is displayed in the **Devices** > **All Devices** listing under Device Type as **OE-VWAAS**.
- The Hyper-V device is displayed in the **Devices** > *device-name* > **Dashboard** as **OE-VWAAS-HYPER-V**.

# Operating Considerations for vWAAS on Hyper-V

This section has the following topics:


- Configuring NTP Settings for vWAAS on Hyper-V
- Traffic Interception Methods for vWAAS on Hyper-V
- Hyper-V High Availability Features


## Configuring NTP Settings for vWAAS on Hyper-V

The Network Time Protocol (NTP) allows synchronization of time and date settings for the different geographical locations of the devices in your WAAS network, which is important for proper system operation and monitoring. When you configure NTP on vWAAS with Hyper-V, the time gets updated from the NTP server.

**Caution** To ensure that the vWAAS on Hyper-V system clock remains in synchronization with the system clocks of other WAAS devices, especially after a reload of vWAAS on Hyper-V, you must *uncheck* the **Time synchronization** option. This option must be unchecked in the system that you are using for vWAAS on Hyper-V: System Center Virtual Machine Manager (SC VMM) or the Hyper-V Manager.

To uncheck the Time Synchronization option for NTP configuration, follow these steps:

**Step 1** Uncheck the Time Synchronization option in either the SC VMM or the Hyper-V Manager:

*From the SC VMM:*

a. Select **vWAAS VM**.
b. Choose **Settings** > **Management** > **Integration Services**.
c. Verify that the **Time synchronization** option is unchecked.
d. Click **OK**.

*From the Hyper-V Manager:*

a. Select **vWAAS VM**.

b. Choose **Properties** > **Hardware Configuration** > **Advanced** > **Integration Services**.

c. Verify that the **Time synchronization** option is unchecked.

d. Click **OK**.

# Traffic Interception Methods for vWAAS on Hyper-V

This section has the following topics:


- About Traffic Interception for vWAAS on Hyper-V
- WCCP Interception
- AppNav Controller Interception


## About Traffic Interception for vWAAS on Hyper-V

When vWAAS is deployed in Hyper-V hosts, the WAE device is replaced by the Hyper-V host. No change is required in the WAAS traffic interception mechanism in the switches or routers. The WCCP protocol also works like the vWAAS VMware ESXi deployment in the vWAAS Hyper-V deployment.

vWAAS on Hyper-V provides the same WAN acceleration functionality provided by the physical WAN acceleration WAE device. You can also deploy multiple vWAAS in one or more Hyper-V hosts to form a WAAS farm in either the Edge or the Core.

## WCCP Interception

WCCP interception, WCCP GRE or WCCP L2, is supported for all vWAAS on Hyper-V deployments.

To select WCCP as the interception method for a WAE, follow these overview steps. For a full description of each step, see the *Cisco Wide Area Application Services Configuration Guide*.

**Note** Before you do the following procedure, you should have already configured your router for basic WCCP, as described in the *Cisco Wide Area Application Services Configuration Guide*.

**Step 1** From the WAAS Central Manager menu, choose Devices > *device-name*.

**Step 2** Choose **Configure** > **Interception** > **Interception Configuration**. The Interception Configuration window appears.

#### Interception Method Settings area

**Step 3** From the Interception Method drop-down list, choose **WCCP** to enable the WCCP interception on the vWAAS device.

#### WCCP Settings area

**Step 4** To enable WCCP on the device, check the **Enable WCCP Service** check box.

**Step 5** With WCCP selected, the **Service Type** field displays TCP Promiscuous.

**Step 6** In the Service ID1 field, specify the first service ID of the WCCP service pair, with an ID number of 1 to 99. After you submit, the Service ID2 field is filled in with the second service ID of the pair, which is one greater than Service ID1, with an ID number of 2 to 100.

**Step 7** To use the default gateway of the WAE as the router to associate with the WCCP TCP promiscuous service, check the **Use Default Gateway as WCCP Router** check box.

If you leave this box unchecked, you can use the **WCCP Routers** field to specify a list of one or more routers by their IP addresses, separated by spaces.

#### WCCP Assignment Settings for Load Balancing area

**Step 8** (Optional) From the **Assignment Method** drop-down list, choose the type of WAE load-balancing assignment method to use (**Mask** or **Hash**).

- Mask assignment method selected—To use a custom mask, enter a value for the source ID mask in the **Source IP Mask** field. The range, in hexadecimal, is 00000000–FE000000. The default is F00. Enter a value for the destination IP mask in the **Destination IP Mask** field. The range, in hexadecimal, is 0000000–FE000000. The default is 0.
- Hash assignment method selected—To specify the hash assignment method for the source IP address, check **Hash on Source IP:** either **Service ID1** or **Service ID2**. After you check a source IP, the complementary destinantion IP is automatically selected, **Hash on Destination IP:** check box either **Service ID2** or **Service ID1**.

#### WCCP Redirect and Egress Settings area

**Step 9** From the **Redirect Method** drop-down list, choose **WCCP GRE** or **WCCP L2**.

**Step 10** From the Egress Method drop-down list, choose **L2** or **IP Forwarding**.

#### Advanced WCCP Settings area

**Step 11** Check the **Enable Flow Protection** check box to keep the TCP flow intact and to avoid overwhelming the device when it comes up or is reassigned new traffic. For more information on flow redirection, see the Information about WCCP Flow Redirection on WAEs" section of the *Cisco Wide Area Application Services Configuration Guide*.

**Step 12** In the **Flow Protection Timeout** field, specify the amount of time (in seconds) that flow protection should be enabled. The default is 0, which means flow protection stays enabled with no timeout.

**Step 13** In the **Shutdown Delay** field, enter a maximum amount of time (in seconds) that the chosen device waits to perform a clean shutdown of WCCP. The range is 0 to 86400 seconds. The default is 120 seconds.

**Step 14** From the **Failure Detection Timeout** drop-down list, choose a failure detection timeout value: 30, 15, or 9 seconds. The default is 30 seconds. The failure detection timeout determines the length of time for the router to detect a WAE failure.

**Step 15** In the **Weight** field, specify the weight to be used for load balancing. The weight value range is 0 to 10000.

- If the total of all the weight values of the WAEs in a service group is less than or equal to 100, the weight value represents a literal percentage of the total load redirected to the device for load-balancing purposes.
- If the total of all the weight values of the WAEs in a service group is between 101 and 10000, the weight value is treated as a fraction of the total weight of all the active WAEs in the service group.

**Step 16** In the **Password** field, specify the password to be used for secure traffic between the WAEs within a cluster and the router for a specified service. Be sure to enable all other WAEs and routers within the cluster with the same password. Passwords must not exceed eight characters in length. Do not use the

following characters: space, backwards single quote (‘), double quote (“”), pipe (|), or question mark (?).

Re-enter the password in the **Confirm Password** field.

**Step 17** Click **Submit** to save the settings.

## AppNav Controller Interception

AppNav interception is supported for all vWAAS on Hyper-V deployments, and works as in the current ESXi vWAAS models.

AppNav interception enables a vWAAS node to receive traffic optimization from an AppNav controller (ANC) in an AppNav deployment. If vWAAS VMs are part of an AppNav deployment and are configured as WAAS nodes (WNs) in an AppNav cluster, you must configure the AppNav-controller interception method. These WNs receive traffic only from the ANCs; they do not receive traffic directly from routers.

To select AppNav as the interception method, follow these steps:

**Step 1** From the WAAS Central Manager menu, choose Devices > *device-name*.

**Step 2** Choose **Configure** > **Interception** > **Interception Configuration**. The Interception Configuration window appears.

**Step 3** From the Interception Method drop-down list, choose **appnav-controller** to enable appnav-controller interception on hte vWAAS device.

**Step 4** Click **Submit**.

# Hyper-V High Availability Features

vWAAS on Hyper-V provides multiple high availability solutions, including:

- Live Migration
- NIC Teaming

## Live Migration

Hyper-V live migration moves running VMs with no impact on VM availability to the user. It does this by pre-copying the memory of the migrating VM to the destination physical host. The administrator, or the script, that initiates the live migration controls which computer is the destination for the live migration. There is no need for special configuration for the guest operating system, as that is not affected by the live migration.

There are three methods you can use to initiate a live migration:

- Failover Cluster console
- Virtual Machine Manager Administration console (if Virtual Machine Manager is managing physical hosts that are configured to support live migration)
- A PowerShell or WMI script

The following is a workflow for initiating and completing a live migration:

- **Create a connection between hosts**—The source physical host creates a TCP connection with the destination physical host, which is used to transfer the VM configuration data to the destination physical host. A skeleton VM is set up on the destination physical host, and memory is allocated to the destination VM.
- **Copy the working set to the destination host**—The memory assigned to the migrating VM, called the working set, is copied to the destination physical host. This memory is referred to as the working set of the migrating VM. A page of memory is 4 kB in size.
- **Mark modified memory pages**—The utilized pages within the working set are copied to the destination Hyper-V physical host. In addition to copying the working set to the destination physical host, Hyper-V on the source physical host monitors the pages in the working set. As the migrating VM modified the memory pages during live migration, Hyper-V tracks and marks them as modified.
- **Copy modified memory pages**—During live migration, Hyper-V iterates the memory copy process several times. Each time, a smaller number of modified pages need to be copies to the destination physical host. A final memory copy process copies the remaining modified memory pages to the destination physical host.

  The source physical host transfers the register and device state of the VM to the destination physical host. During this stage of live migration, the network bandwidth available between the source and physical host is critical to the speed of the migration. Therefore, 1 Gigabit Ethernet is recommended.

  **Note** The number of pages to be transferred in this stage is dictated by how actively the VM is accessing and modifying memory pages. More modified pages means a longer VM migration time, to allow for all memory pages to be transferred to the destination physical host.

- **Complete the live migration**—After the modified memory pages have been completely copied to the destination physical host, the destination physical host has an up-to-date working set of the migrated VM: the working set for the migrated VM is present on the destination physical host in the exact state it was in when the migrated VM began the live migration process.

  **Note** You can cancel the live migration process at any point before this phase of the process.

- **Transfer control of the migrated VM memory and storage**—Control of storage associated with the migrated VM, such as VHD files or pass-through disks, and control of memory (working set) are transferred to the destination physical host.
- **Bring migrated VM online**—The migrated VM is brought online on the destination physical host.

## NIC Teaming

The failure of an individual Hyper-V port or virtual network adapter can cause a loss of connectivity for a virtual machine. To prevent this, multiple virtual network adapter are used in a NIC (Network Interface Card) teaming configuration, which provides both high availability and load balancing across multiple physical network interfaces. NIC teaming is also known as network adapter teaming technology and LBFO (Load Balancing Failover).

For vWAAS on Hyper-V, NIC teaming, in Windows Server 2012, enables a virtual machine to have virtual network adapters that are connected to more than one virtual switch, and will still have connectivity even if the network adapter under that virtual switch is disconnected. NIC teaming on Windows Server 2012 supports up to 32 network adapters in a team.

With NIC teaming, you can set up two virtual switches, each connected to its own SR-IOV-capable network adapter. NIC teaming then works in one of two ways:

- Each virtual machine can install a virtual function from one or both SR-IOV network adapters. If a adapter disconnection occurs, the traffic can fail over from the primary virtual function to the backup virtual function without losing connectivity.
- Each virtual machine can have a virtual function from one network adapter and a non-virtual functional interface to the other switch. If the network adapter associated with the virtual function becomes disconnected, the traffic can fail over to the other switch without losing connectivity.