NIC Teaming

https://docs.microsoft.com/en-us/windows-server/networking/technologies/nic-teaming/nic-teaming

Applies to: Windows Server 2022, Windows Server 2019, Windows Server 2016, Azure Stack HCI, versions 21H2 and 20H2

Contents

In this topic, we give you an overview of Network Interface Card (NIC) Teaming in Windows Server. NIC Teaming allows you to group between one and 32 physical Ethernet network adapters into one or more software-based virtual network adapters. These virtual network adapters provide fast performance and fault tolerance in the event of a network adapter failure.

Important

You must install NIC Team member network adapters in the same physical host computer.

Tip

A NIC team that contains only one network adapter cannot provide load balancing and failover. However, with one network adapter, you can use NIC Teaming for separation of network traffic when you are also using virtual Local Area Networks (VLANs).

When you configure network adapters into a NIC team, they connect into the NIC teaming solution common core, which then presents one or more virtual adapters (also called team NICs [tNICs] or team interfaces) to the operating system.

Since Windows Server 2016 supports up to 32 team interfaces per team, there are a variety of algorithms that distribute outbound traffic (load) between the NICs. The following illustration depicts a NIC Team with multiple tNICs.

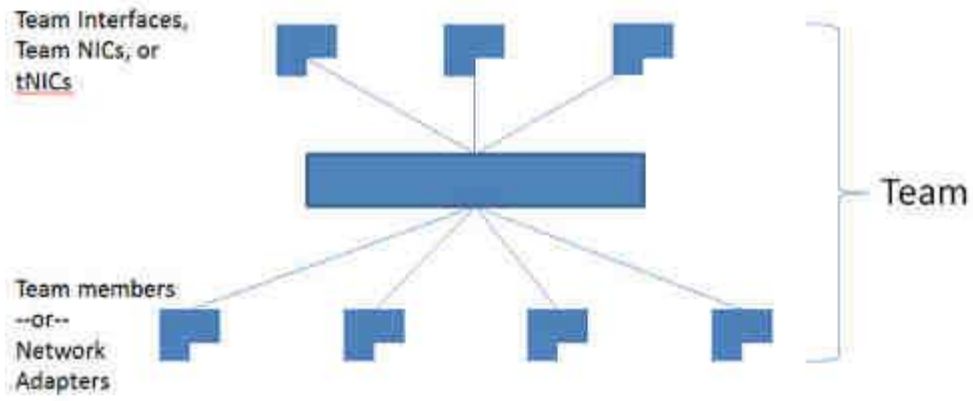

Also, you can connect your teamed NICs to the same switch or different switches. If you connect NICs to different switches, both switches must be on the same subnet.

Availability

NIC Teaming is available in all versions of Windows Server 2016. You can use a variety of tools to manage NIC Teaming from computers running a client operating system, such as:

- Windows PowerShell cmdlets
- Remote Desktop
- Remote Server Administration Tools

Supported and Unsupported NICs

You can use any Ethernet NIC that has passed the Windows Hardware Qualification and Logo test (WHQL tests) in a NIC Team in Windows Server 2016.

You can not place the following NICs in a NIC team:

- Hyper-V virtual network adapters that are Hyper-V Virtual Switch ports exposed as NICs in the host partition.

Important

Do not place Hyper-V virtual NICs exposed in the host partition (vNICs) in a team. Teaming of vNICs inside of the host partition is not supported in any configuration. Attempts to team vNICs might cause a complete loss of communication if network failures occur.

- Kernel debug network adapter (KDNIC).
- NICs used for network boot.
- NICs that use technologies other than Ethernet, such as WWAN, WLAN/Wi-Fi, Bluetooth, and Infiniband, including Internet Protocol over Infiniband (IPoIB) NICs.

Compatibility

NIC teaming is compatible with all networking technologies in Windows Server 2016 with the following exceptions.

- **Single-root I/O virtualization (SR-IOV)**. For SR-IOV, data is delivered directly to the NIC without passing it through the networking stack (in the host operating system, in the case of virtualization). Therefore, it is not possible for the NIC team to inspect or redirect the data to another path in the team.
- **Native host Quality of Service (QoS)**. When you set QoS policies on a native or host system, and those policies invoke minimum bandwidth limitations, the overall throughput for a NIC team is less than it would be without the bandwidth policies in place.
- **TCP Chimney**. TCP Chimney is not supported with NIC teaming because TCP Chimney offloads the entire networking stack directly to the NIC.
- **802.1X Authentication**. You should not use 802.1X Authentication with NIC Teaming because some switches do not permit the configuration of both 802.1X Authentication and NIC Teaming on the same port.

To learn about using NIC Teaming within virtual machines (VMs) that run on a Hyper-V host, see Create a new NIC Team on a host computer or VM.

Virtual Machine Queues (VMQs)

VMQs is a NIC feature that allocates a queue for each VM. Anytime you have Hyper-V enabled; you must also enable VMQ. In Windows Server 2016, VMQs use NIC Switch vPorts with a single queue assigned to the vPort to provide the same functionality.

Depending on the switch configuration mode and the load distribution algorithm, NIC teaming presents either the smallest number of available and supported queues by any adapter in the team (Min-Queues mode) or the total number of queues available across all team members (Sum-of-Queues mode).

If the team is in Switch-Independent teaming mode and you set the load distribution to Hyper-V Port mode or Dynamic mode, the number of queues reported is the sum of all the queues available from the team members (Sum-of-Queues mode). Otherwise, the number of queues reported is the smallest number of queues supported by any member of the team (Min-Queues mode).

Here's why:

- When the switch-independent team is in Hyper-V Port mode or Dynamic mode the inbound traffic for a Hyper-V switch port (VM) always arrives on the same team member. The host can predict/control which member receives the traffic for a particular VM so NIC Teaming can be more thoughtful about which VMQ Queues to allocate on a particular team member. NIC Teaming, working with the Hyper-V switch, sets the VMQ for a VM on precisely one team member and know that inbound traffic hits that queue.
- When the team is in any switch dependent mode (static teaming or LACP teaming), the switch that the team is connected to controls the inbound traffic distribution. The host's NIC Teaming software can't predict which team member gets the inbound traffic for a VM and it may be that the switch distributes the traffic for a VM across all team members. As a result of the NIC Teaming software, working with the Hyper-V switch, programs a queue for the VM on every team member, not just one team member.

- When the team is in switch-independent mode and uses address hash load balancing, the inbound traffic always comes in on one NIC (the primary team member) - all of it on just one team member. Since other team members aren't dealing with inbound traffic, they get programmed with the same queues as the primary member so that if the primary member fails, any other team member can be used to pick up the inbound traffic, and the queues are already in place.
- Most NICs have queues used for either Receive Side Scaling (RSS) or VMQ, but not at the same time. Some VMQ settings appear to be settings for RSS queues but are settings on the generic queues that both RSS and VMQ use depending on which feature is presently in use. Each NIC has, in its advanced properties, values for *RssBaseProcNumber and *MaxRssProcessors. Following are a few VMQ settings that provide better system performance.
- Ideally, each NIC should have the *RssBaseProcNumber set to an even number greater than or equal to two (2). The first physical processor, Core 0 (logical processors 0 and 1), typically does most of the system processing so the network processing should steer away from this physical processor. Some machine architectures don't have two logical processors per physical processor, so for such machines, the base processor should be greater than or equal to 1. If in doubt assume your host is using a 2 logical processor per physical processor architecture.
- If the team is in Sum-of-Queues mode the team members' processors should be non-overlapping. For example, in a 4-core host (8 logical processors) with a team of 2 10Gbps NICs, you could set the first one to use the base processor of 2 and to use 4 cores; the second would be set to use base processor 6 and use 2 cores.
- If the team is in Min-Queues mode the processor sets used by the team members must be identical.

Hyper-V Network Virtualization (HNV)

NIC Teaming is fully compatible with Hyper-V Network Virtualization (HNV). The HNV management system provides information to the NIC Teaming driver that allows NIC Teaming to distribute the load in a way that optimizes HNV traffic.

Live Migration

NIC Teaming in VMs does not affect Live Migration. The same rules exist for Live Migration whether or not configuring NIC Teaming in the VM.

Virtual Local Area Networks (VLANs)

When you use NIC Teaming, creating multiple team interfaces allows a host to connect to different VLANs at the same time. Configure your environment using the following guidelines:

- Before you enable NIC Teaming, configure the physical switch ports connected to the teaming host to use trunk (promiscuous) mode. The physical switch should pass all traffic to the host for filtering without modifying the traffic.
- Do not configure VLAN filters on the NICs by using the NIC advanced properties settings. Let the NIC Teaming software or the Hyper-V Virtual Switch (if present) perform VLAN filtering.

Use VLANs with NIC Teaming in a VM

When a team connects to a Hyper-V Virtual Switch, all VLAN segregation must be done in the Hyper-V Virtual Switch rather than in NIC Teaming.

Plan to use VLANs in a VM configured with a NIC Team using the following guidelines:

- The preferred method of supporting multiple VLANs in a VM is to configure the VM with multiple ports on the Hyper-V Virtual Switch and associate each port with a VLAN. Never team these ports in the VM because doing so causes network communication problems.
- If the VM has multiple SR-IOV Virtual Functions (VFs), ensure that they are on the same VLAN before teaming them in the VM. It's easily possible to configure the different VFs to be on different VLANs and doing so causes network communication problems.

Manage network interfaces and VLANs

If you must have more than one VLAN exposed into a guest operating system, consider renaming the Ethernet interfaces to clarify VLAN assigned to the interface. For example, if you associate **Ethernet** interface with VLAN 12 and the **Ethernet 2** interface with VLAN 48, rename the interface Ethernet to **EthernetVLAN12** and the other to **EthernetVLAN48**.

Rename interfaces by using the Windows PowerShell command **Rename-NetAdapter** or by performing the following procedure:

1. In Server Manager, in **Properties** for the network adapter you want to rename, click the link to the right of the network adapter name.
2. Right-click the network adapter that you want to rename, and select **Rename**.
3. Type the new name for the network adapter and press ENTER.

Virtual Machines (VMs)

If you want to use NIC Teaming in a VM, you must connect the virtual network adapters in the VM to external Hyper-V Virtual Switches only. Doing this allows the VM to sustain network connectivity even in the circumstance when one of the physical network adapters connected to one virtual switch fails or gets disconnected. Virtual network adapters connected to internal or private Hyper-V Virtual Switches are not able to connect to the switch when they are in a team, and networking fails for the VM.

NIC Teaming in Windows Server 2016 supports teams with two members in VMs. You can create larger teams, but there is no support for larger teams. Every team member must connect to a different external Hyper-V Virtual Switch, and the VM's networking interfaces must be configured to allow teaming.

If you are configuring a NIC Team in a VM, you must select a **Teaming mode** of *Switch Independent* and a **Load balancing mode** of *Address Hash*.

SR-IOV-Capable Network Adapters

A NIC team in or under the Hyper-V host cannot protect SR-IOV traffic because it doesn't go through the Hyper-V Switch. With the VM NIC Teaming option, you can configure two external Hyper-V Virtual Switches, each connected to its own SR-IOV-capable NIC.

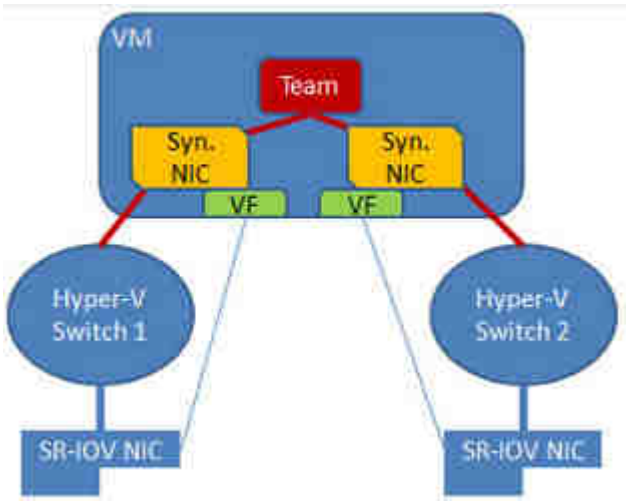

Each VM can have a virtual function (VF) from one or both SR-IOV NICs and, in the event of a NIC disconnect, failover from the primary VF to the backup adapter (VF). Alternately, the VM may have a VF from one NIC and a non-VF vmNIC connected to another virtual switch. If the NIC associated with the VF gets disconnected, the traffic can failover to the other switch without loss of connectivity.

Because failover between NICs in a VM might result in traffic sent with the MAC address of the other vmNIC, each Hyper-V Virtual Switch port associated with a VM using NIC Teaming must be set to allow teaming.

NIC Teaming settings

https://docs.microsoft.com/en-us/windows-server/networking/technologies/nic-teaming/nic-teaming-settings

In this topic, we give you an overview of the NIC Team properties such as teaming and load balancing modes. We also give you details about the Standby adapter setting and the Primary team interface property. If you have at least two network adapters in a NIC Team, you do not need to designate a Standby adapter for fault tolerance.

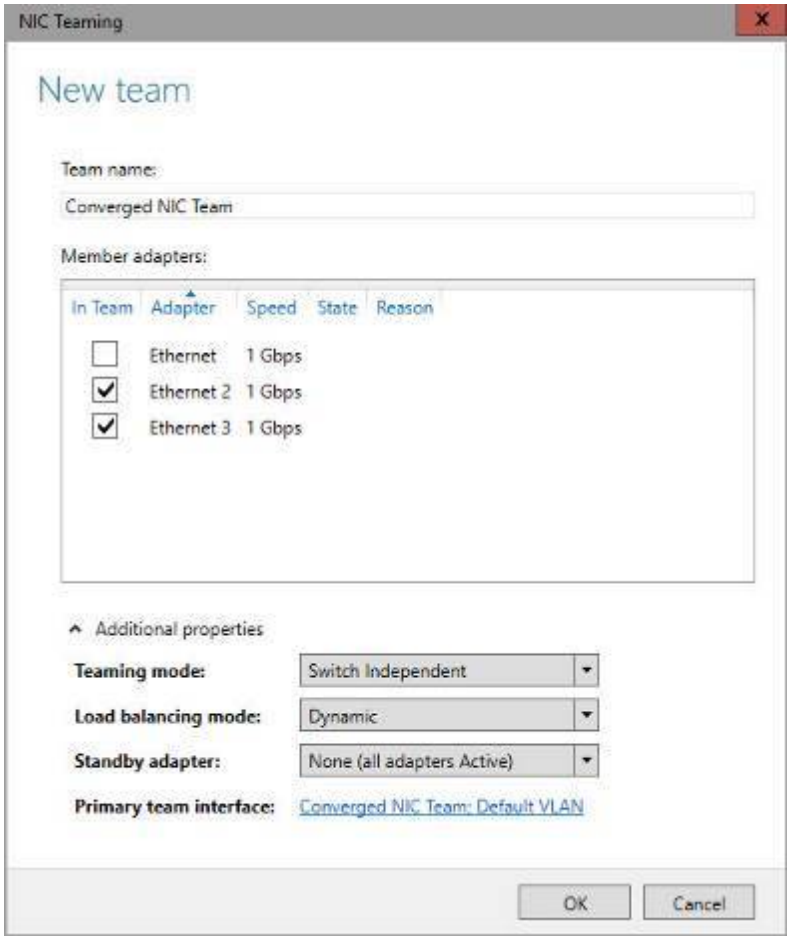

Teaming modes

The options for Teaming mode are **Switch Independent** and **Switch Dependent**. The Switch Dependent mode includes **Static Teaming** and **Link Aggregation Control Protocol (LACP)**.

Tip

For best NIC Team performance, we recommend that you use a Load Balancing mode of Dynamic distribution.

Switch Independent

With Switch Independent mode, the switch or switches to which the NIC Team members are connected are unaware of the presence of the NIC team and do not determine how to distribute network traffic to NIC Team members - instead, the NIC Team distributes inbound network traffic across the NIC Team members.

When you use Switch Independent mode with Dynamic distribution, the network traffic load is distributed based on the TCP Ports address hash as modified by the Dynamic load balancing algorithm. The Dynamic load balancing algorithm redistributes flows to optimize team member bandwidth utilization so that individual flow transmissions can move from one active team member to another. The algorithm takes into account the small possibility that redistributing traffic could cause out-of-order delivery of packets, so it takes steps to minimize that possibility.

Switch Dependent

With Switch Dependent modes, the switch to which the NIC Team members are connected determines how to distribute the inbound network traffic among the NIC Team members. The switch has complete independence to determine how to distribute the network traffic across the NIC Team members.

Important

Switch dependent teaming requires that all team members are connected to the same physical switch or a multi-chassis switch that shares a switch ID among the multiple chassis.

- **Static Teaming.** Static Teaming requires you to manually configure both the switch and the host to identify which links form the team. Because this is a statically configured solution, there is no additional protocol to assist the switch and the host to identify incorrectly plugged cables or other errors that could cause the team to fail to perform. This mode is typically supported by server-class switches.

- **Link Aggregation Control Protocol (LACP).** Unlike Static Teaming, LACP Teaming mode dynamically identifies links that are connected between the host and the switch. This dynamic connection enables the automatic creation of a team and, in theory but rarely in practice, the expansion and reduction of a team simply by the transmission or receipt of LACP packets from the peer entity. All server-class switches support LACP, and all require the network operator to administratively enable LACP on the switch port. When you configure a Teaming mode of LACP, NIC Teaming always operates in LACP's Active mode with a short timer. No option is presently available to modify the timer or change the LACP mode.

When you use Switch Dependent modes with Dynamic distribution, the network traffic load is distributed based on the TransportPorts address hash as modified by the Dynamic load balancing algorithm. The Dynamic load balancing algorithm redistributes flows to optimize team member bandwidth utilization. Individual flow transmissions can move from one active team member to another as part of the dynamic distribution. The algorithm takes into account the small possibility that redistributing traffic could cause out-of-order delivery of packets, so it takes steps to minimize that possibility.

As with all switch dependent configurations, the switch determines how to distribute the inbound traffic among the team members. The switch is expected to do a reasonable job of distributing the traffic across the team members but it has complete independence to determine how it does so.

Load Balancing modes

The options for Load Balancing distribution mode are **Address Hash**, **Hyper-V Port**, and **Dynamic**.

Address Hash

With Address Hash, this mode creates a hash based on address components of the packet, which then get assigned to one of the available adapters. Usually, this mechanism alone is sufficient to create a reasonable balance across the available adapters.

Use Windows PowerShell to specify values for the following hashing function components.

- Source and destination TCP ports and source and destination IP addresses. This is the default when you select **Address Hash** as the Load Balancing mode.
- Source and destination IP addresses only.
- Source and destination MAC addresses only.

The TCP ports hash creates the most granular distribution of traffic streams, resulting in smaller streams that can be independently moved between NIC team members. However, you cannot use the TCP ports hash for traffic that is not TCP or UDP-based, or where the TCP and UDP ports are hidden from the stack, such as with IPsec-protected traffic. In these cases, the hash automatically uses the IP address hash or, if the traffic is not IP traffic, the MAC address hash is used.

Hyper-V Port

With Hyper-V Port, NIC Teams configured on Hyper-V hosts give VMs independent MAC addresses. The VMs MAC address or the VM ported connected to the Hyper-V switch, can be used to divide network traffic between NIC Team members. You cannot configure NIC Teams that you create within VMs with the Hyper-V Port load balancing mode. Instead, use the Address Hash mode.

Because the adjacent switch always sees a particular MAC address on one port, the switch distributes the ingress load (the traffic from the switch to the host) on multiple links based on the destination MAC (VM MAC) address. This is particularly useful when Virtual Machine Queues (VMQs) are used, because a queue can be placed on the specific NIC where the traffic is expected to arrive.

However, if the host has only a few VMs, this mode might not be granular enough to achieve a well-balanced distribution. This mode will also always limit a single VM (i.e., the traffic from a single switch port) to the bandwidth that is available on a single interface. NIC Teaming uses the Hyper-V Virtual Switch Port as the identifier instead of using the source MAC address because, in some instances, a VM might be configured with more than one MAC address on a switch port.

Dynamic

With Dynamic, outbound loads are distributed based on a hash of the TCP ports and IP addresses. Dynamic mode also rebalances loads in real time so that a given outbound flow may move back and forth between team members. Inbound loads, on the other hand, get distributed the same way as Hyper-V Port. In a nutshell, Dynamic mode utilizes the best aspects of both Address Hash and Hyper-V Port and is the highest performing load balancing mode.

The outbound loads in this mode are dynamically balanced based on the concept of flowlets. Just as human speech has natural breaks at the ends of words and sentences, TCP flows (TCP communication streams) also have naturally occurring breaks. The portion of a TCP flow between two such breaks is referred to as a flowlet.

When the dynamic mode algorithm detects that a flowlet boundary has been encountered - such as when a break of sufficient length has occurred in the TCP flow - the algorithm automatically rebalances the flow to another team member if appropriate. In some circumstances the algorithm might also periodically rebalance flows that do not contain any flowlets. Because of this, the affinity between TCP flow and team member can change at any time as the dynamic balancing algorithm works to balance the workload of the team members.

Whether the team is configured with Switch Independent or one of the Switch Dependent modes, it is recommended that you use Dynamic distribution mode for best performance.

There is an exception to this rule when the NIC Team has just two team members, is configured in Switch Independent mode, and has Active/Standby mode enabled, with one NIC active and the other configured for Standby. With this NIC Team configuration, Address Hash distribution provides slightly better performance than Dynamic distribution.

Standby adapter setting

The options for Standby Adapter are **None (all adapters Active)** or your selection of a specific network adapter in the NIC Team that acts as a Standby adapter. When you configure a NIC as a Standby adapter, all other unselected team members are Active, and no network traffic is sent to or processed by the adapter until an Active NIC fails. After an Active NIC fails, the Standby NIC becomes active and processes network traffic. When all team members get restored to service, the standby team member returns to standby status.

If you have a two-NIC team and you choose to configure one NIC as a Standby adapter, you lose the bandwidth aggregation advantages that exist with two active NICs. You do not need to designate a Standby Adapter to achieve fault tolerance; fault tolerance is always present whenever there are at least two network adapters in a NIC Team.

Primary Team interface property

To access the Primary Team Interface dialog box, you must click the link that is highlighted in the illustration below.

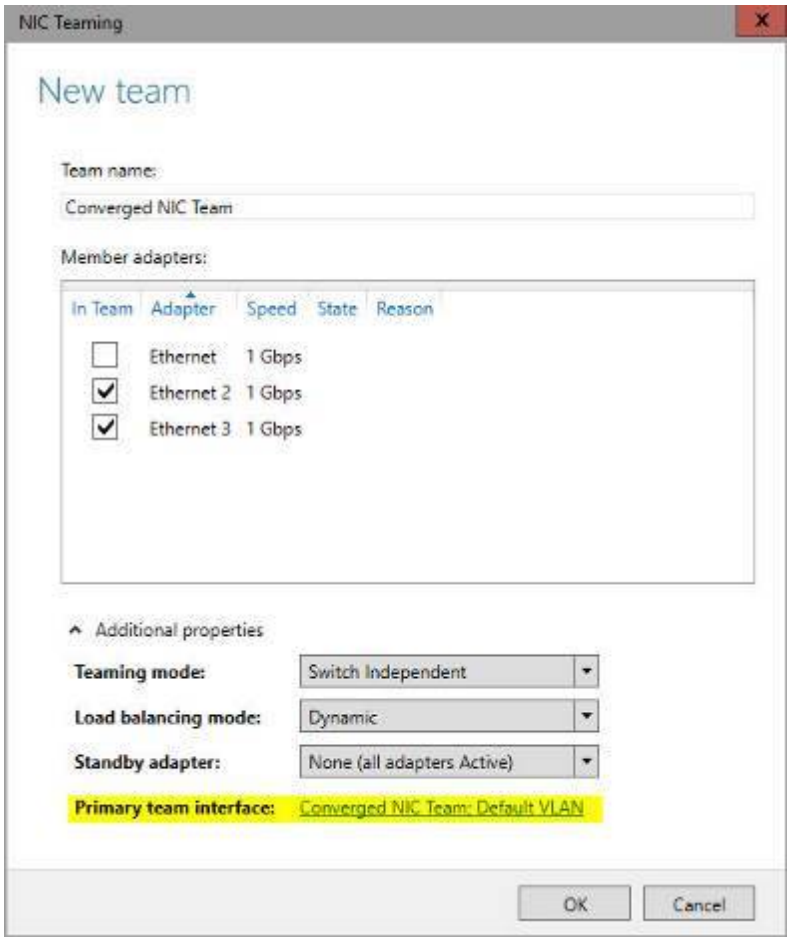

After you click the highlighted link, the following **New Team Interface** dialog box opens.

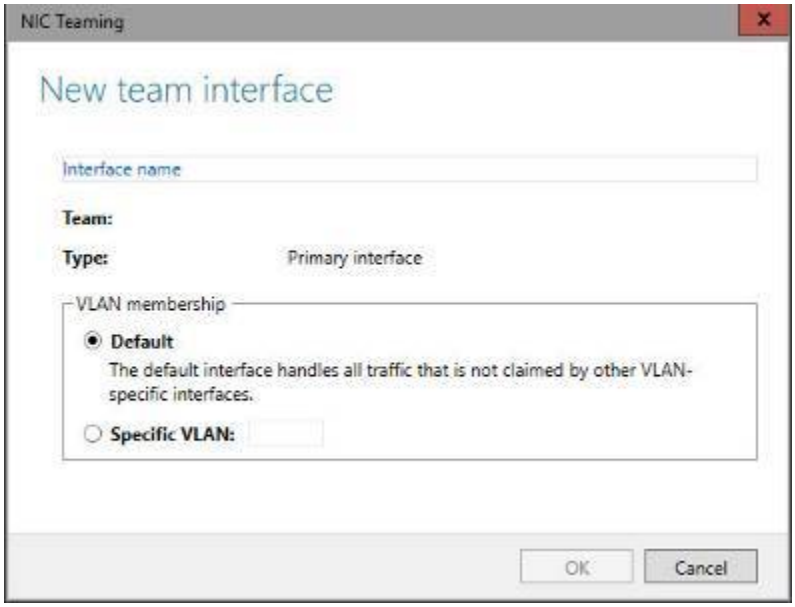

If you are using VLANs, you can use this dialog box to specify a VLAN number.

Whether or not you are using VLANs, you can specify a NIC name for the NIC Team.

NIC Teaming MAC address use and management

https://docs.microsoft.com/en-us/windows-server/networking/technologies/nic-teaming/nic-teaming-mac-address-use-and-management

Applies to: Windows Server 2022, Windows Server 2019, Windows Server 2016, Azure Stack HCI, versions 21H2 and 20H2

When you configure a NIC Team with switch independent mode and either address hash or dynamic load distribution, the team uses the media access control (MAC) address of the primary NIC Team member on outbound traffic. The primary NIC Team member is a network adapter that the operating system selects from the initial set of team members. It is the first team member to bind to the team after you create it or after the host computer is restarted. Because the primary team member might change in a non-deterministic manner at each boot, NIC disable/enable action, or other reconfiguration activities, the primary team member might change, and the MAC address of the team might vary.

In most situations this doesn't cause problems, but there are a few cases where issues might arise.

If the primary team member is removed from the team and then placed into operation, there may be a MAC address conflict. To resolve this conflict, disable and then enable the team interface. The process of disabling and enabling the team interface causes the interface to select a new MAC address from the remaining team members, and eliminates the MAC address conflict.

You can set the MAC address of the NIC team to a specific MAC address by setting it in the primary team interface, just as you can do when configuring the MAC address of any physical NIC.

MAC address use on transmitted packets

When you configure a NIC Team in switch independent mode and either address hash or dynamic load distribution, the packets from a single source (such as a single VM) are simultaneously distributed across multiple team members. To prevent the switches from getting confused and to prevent MAC flapping alarms, the source MAC address is replaced with a different MAC address on the frames transmitted on team members other than the primary team member. Because of this, each team member uses a different MAC address. MAC address conflicts are prevented unless and until failure occurs.

When a failure is detected on the primary NIC, the NIC Teaming software starts using the primary team member's MAC address on the team member that is chosen to serve as the temporary primary team member (that is, the one that will now appear to the switch as the primary team member). This change only applies to traffic that was going to be sent on the primary team member with the primary team member's MAC address as its source MAC address. Other traffic continues to be sent with whatever source MAC address it would have used prior to the failure.

Following are lists that describe NIC Teaming MAC address replacement behavior, based on how the team is configured:

1. **In Switch Independent mode with Address Hash distribution**
    - All ARP and NS packets are sent on the primary team member
    - All traffic sent on NICs other than the primary team member are sent with the source MAC address modified to match the NIC on which they are sent
    - All traffic sent on the primary team member is sent with the original source MAC address (which may be the team's source MAC address)
2. **In Switch Independent mode with Hyper-V Port distribution**
    - Every vmSwitch port is affinitized to a team member
    - Every packet is sent on the team member to which the port is affinitized
    - No source MAC replacement is done
3. **In Switch Independent mode with Dynamic distribution**
    - Every vmSwitch port is affinitized to a team member
    - All ARP/NS packets are sent on the team member to which the port is affinitized
    - Packets sent on the team member that is the affinitized team member have no source MAC address replacement done
    - Packets sent on a team member other than the affinitized team member will have source MAC address replacement done
4. **In Switch Dependent mode (all distributions)**
    - No source MAC address replacement is performed

Create a new NIC Team on a host computer or VM

https://docs.microsoft.com/en-us/windows-server/networking/technologies/nic-teaming/create-a-new-nic-team-on-a-host-computer-or-vm

Applies to: Windows Server 2022, Windows Server 2019, Windows Server 2016, Azure Stack HCI, versions 21H2 and 20H2

In this topic, you create a new NIC Team on a host computer or in a Hyper-V virtual machine (VM) running Windows Server 2016.

Network configuration requirements

Before you can create a new NIC Team, you must deploy a Hyper-V host with two network adapters that connect to different physical switches. You must also configure the network adapters with IP addresses that are from the same IP address range.

The physical switch, Hyper-V Virtual Switch, local area network (LAN), and NIC Teaming requirements for creating a NIC Team in a VM are:

- The computer running Hyper-V must have two or more network adapters.
- If connecting the network adapters to multiple physical switches, the physical switches must be on the same Layer 2 subnet.
- You must use Hyper-V Manager or Windows PowerShell to create two external Hyper-V Virtual Switches, each connected to a different physical network adapter.
- The VM must connect to both external virtual switches you create.
- NIC Teaming, in Windows Server 2016, supports teams with two members in VMs. You can create larger teams, but there is no support.
- If you are configuring a NIC Team in a VM, you must select a **Teaming mode** of *Switch Independent* and a **Load balancing mode** of *Address Hash*.

Step 1. Configure the physical and virtual network

In this procedure, you create two external Hyper-V Virtual Switches, connect a VM to the switches, and then configure the VM connections to the switches.

Prerequisites

You must have membership in **Administrators**, or equivalent.

Procedure

1. On the Hyper-V host, open Hyper-V Manager, and under Actions, click **Virtual Switch Manager**.

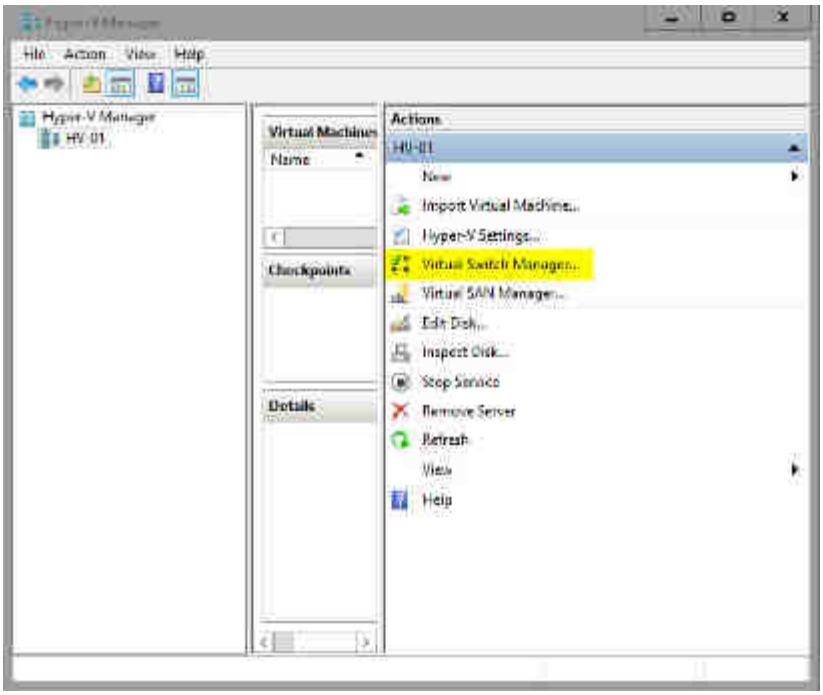

2. In Virtual Switch Manager, make sure **External** is selected, and then click **Create Virtual Switch**.

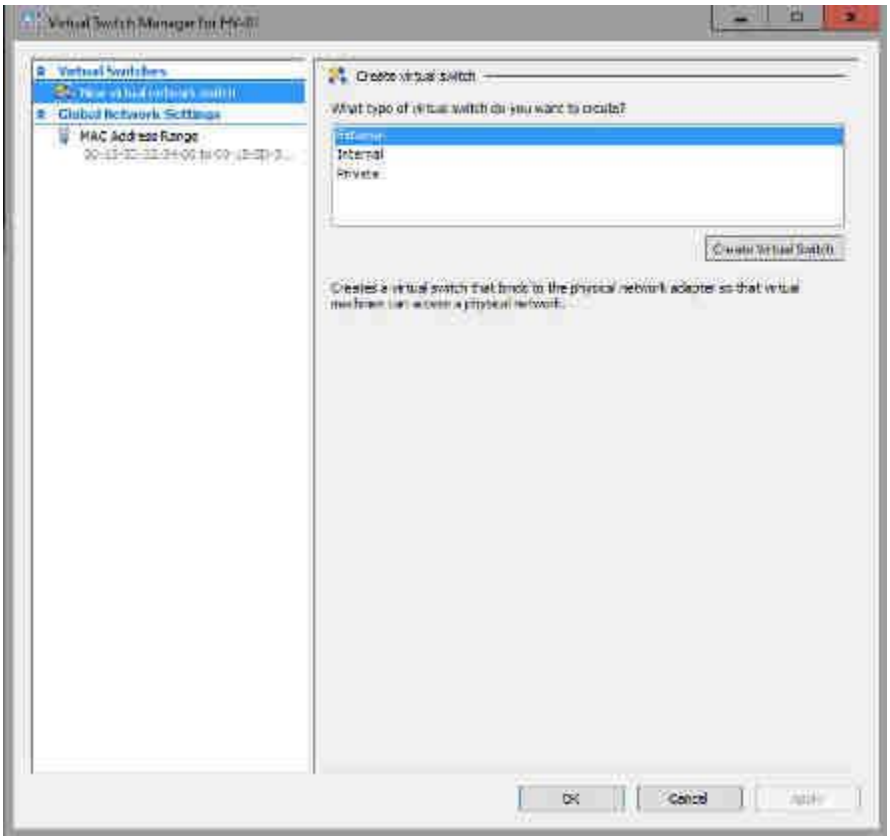

3. In Virtual Switch Properties, type a **Name** for the virtual switch, and add **Notes** as needed.
4. In **Connection type**, in **External network**, select the physical network adapter to which you want to attach the virtual switch.

5. Configure additional switch properties for your deployment, and then click **OK**.
6. Create a second external virtual switch by repeating the previous steps. Connect the second external switch to a different network adapter.
7. In Hyper-V Manager, under **Virtual Machines**, right-click the VM that you want to configure, and then click **Settings**.

   The VM **Settings** dialog box opens.
8. Ensure that the VM is not started. If it is started, perform a shutdown before configuring the VM.
9. In **Hardware**, click **Network Adapter**.

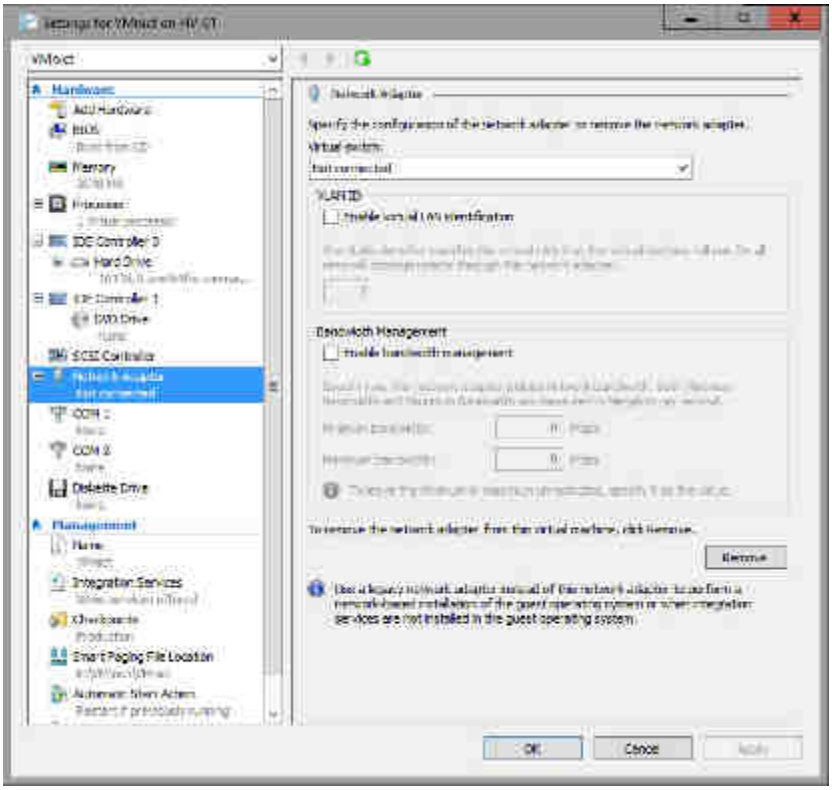

10. In **Network Adapter** properties, select the first virtual switch that you created in previous steps, and then click **Apply**.

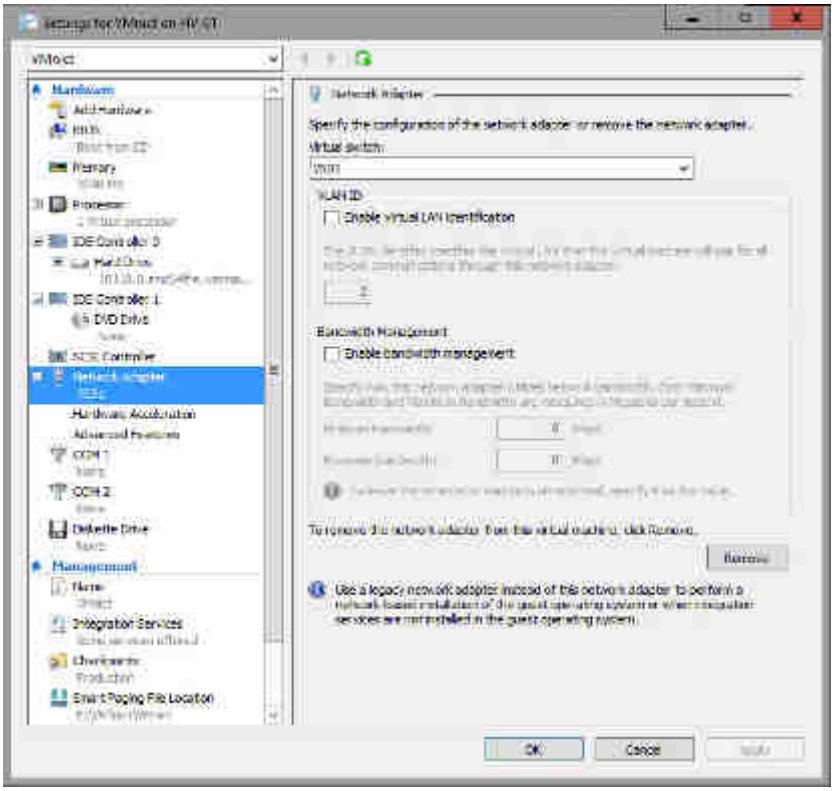

11. In **Hardware**, click to expand the plus sign (+) next to **Network Adapter**.
12. Click **Advanced Features** to enable NIC Teaming by using the graphical user interface.

    Tip
    You can also enable NIC Teaming with a Windows PowerShell command:
    PowerShell
    ```
    Set-VMNetworkAdapter -VMName <VMname> -AllowTeaming On
    ```

12. a. Select **Dynamic** for MAC address.

    b. Click to select **Protected network**.

    c. Click to select **Enable this network adapter to be part of a team in the guest operating system**.

    d. Click **OK**.

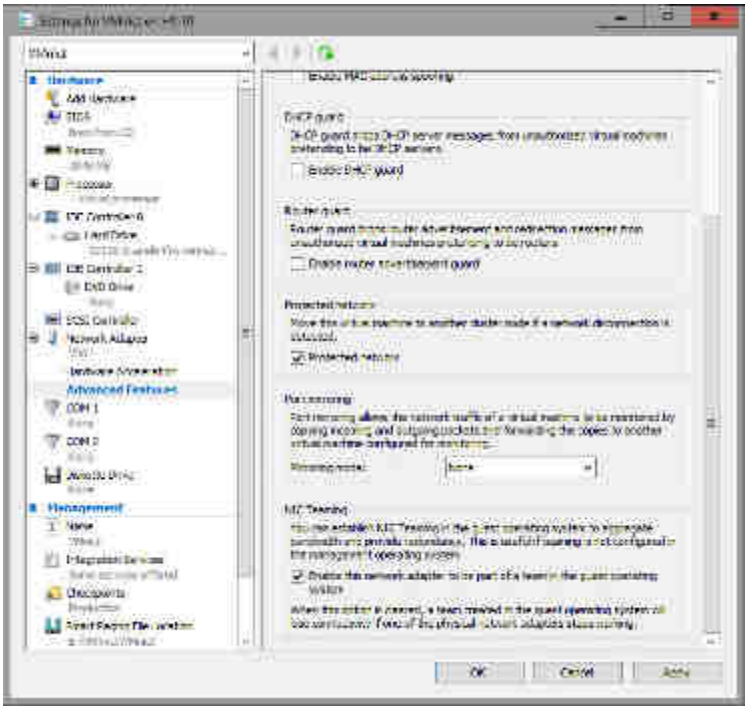

13. To add a second network adapter, in Hyper-V Manager, in **Virtual Machines**, right-click the same VM, and then click **Settings**.

    The VM **Settings** dialog box opens.

14. In **Add Hardware**, click **Network Adapter**, and then click **Add**.

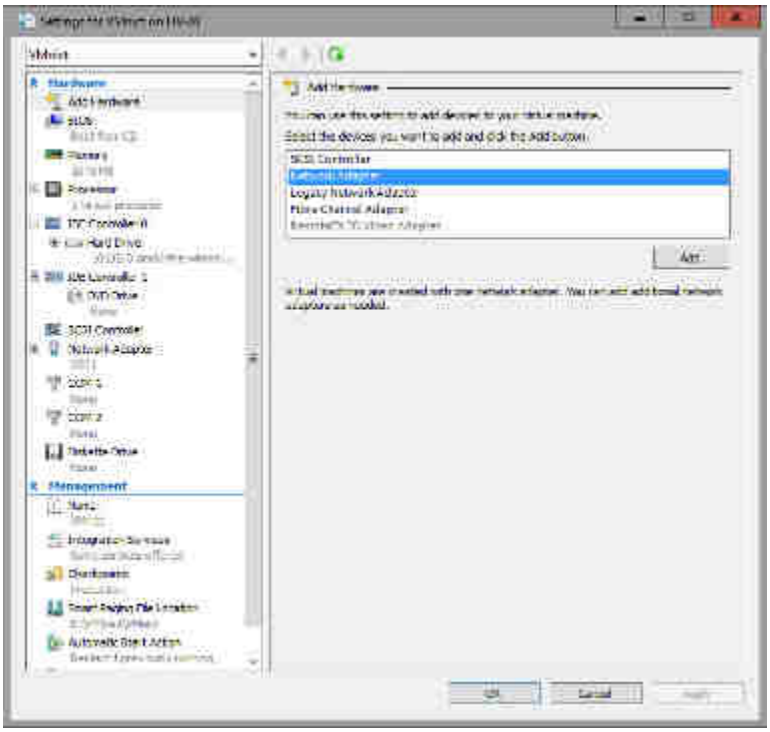

15. In **Network Adapter** properties, select the second virtual switch that you created in previous steps, and then click **Apply**.

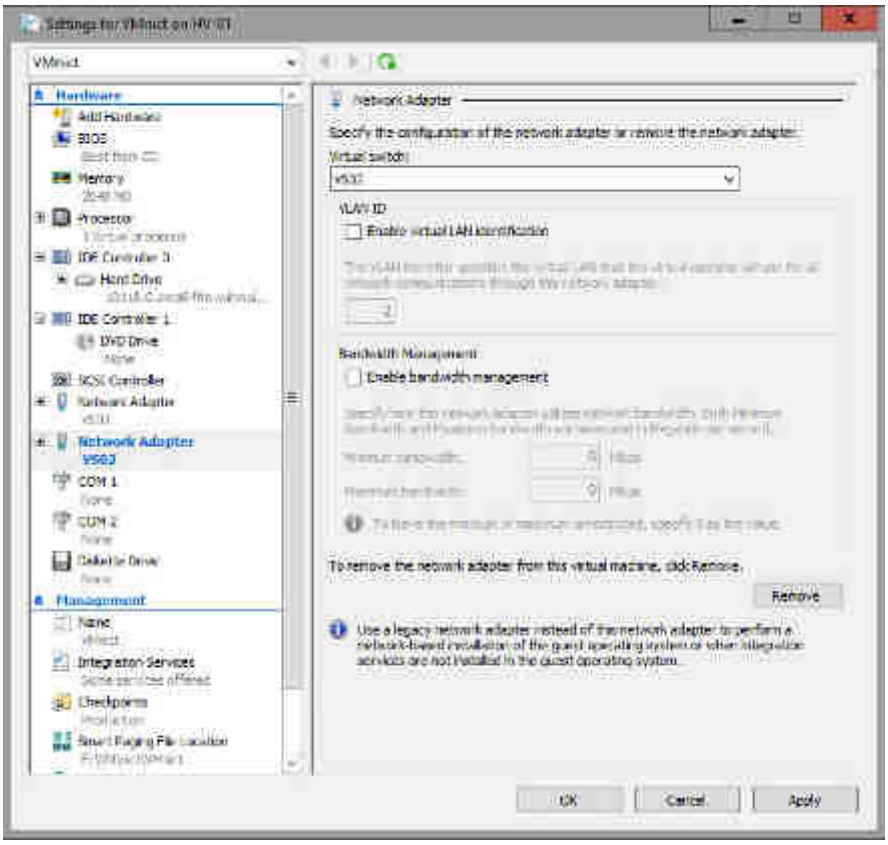

16. In **Hardware**, click to expand the plus sign (+) next to **Network Adapter**.
17. Click **Advanced Features**, scroll down to **NIC Teaming**, and click to select **Enable this network adapter to be part of a team in the guest operating system**.
18. Click **OK**.

Congratulations! You have configured the physical and virtual network. Now you can proceed to creating a new NIC Team.

Step 2. Create a new NIC Team

When you create a new NIC Team, you must configure the NIC Team properties:

- Team name
- Member adapters
- Teaming mode
- Load balancing mode
- Standby adapter

You can also optionally configure the primary team interface and configure a virtual LAN (VLAN) number.

For more details on these settings, see NIC Teaming settings.

Prerequisites

You must have membership in **Administrators**, or equivalent.

Procedure

1. In Server Manager, click **Local Server**.
2. In the **Properties** pane, in the first column, locate **NIC Teaming**, and then click the **Disabled** link.

   The **NIC Teaming** dialog box opens.

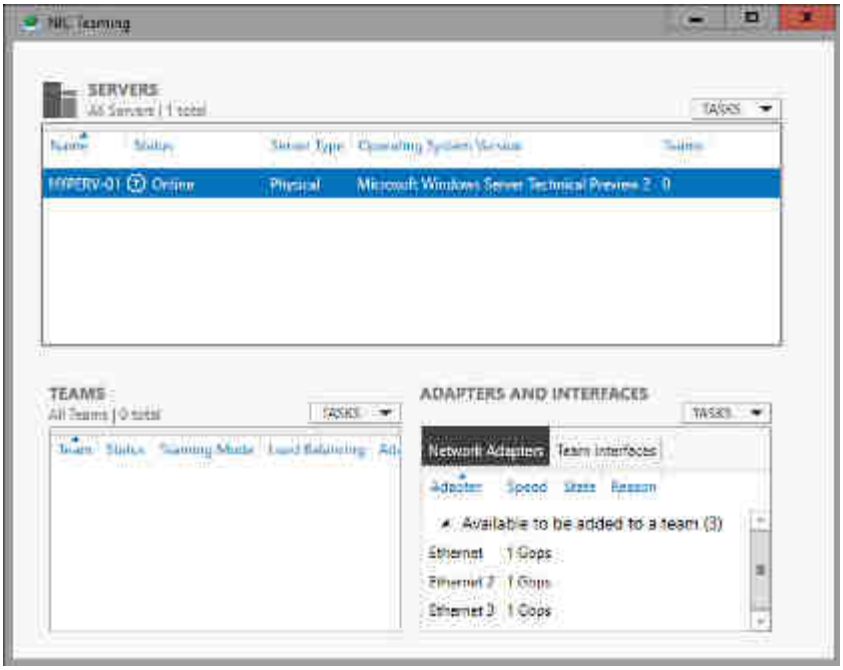

3. In **Adapters and Interfaces**, select the one or more network adapters that you want to add to a NIC Team.
4. Click **TASKS**, and click **Add to New Team**.

   The **New team** dialog box opens and displays network adapters and team members.
5. In **Team name**, type a name for the new NIC Team, and then click **Additional properties**.
6. In **Additional properties**, select values for:
   - **Teaming mode**. The options for Teaming mode are **Switch Independent** and **Switch Dependent**. The Switch Dependent mode includes **Static Teaming** and **Link Aggregation Control Protocol (LACP)**.
     - **Switch Independent.** With Switch Independent mode, the switch or switches to which the NIC Team members are connected are unaware of the presence of the NIC team and do not determine how to distribute network traffic to NIC Team members - instead, the NIC Team distributes inbound network traffic across the NIC Team members.
     - **Switch Dependent.** With Switch Dependent modes, the switch to which the NIC Team members are connected determines how to distribute the inbound network traffic among the NIC Team members. The switch has complete independence to determine how to distribute the network traffic across the NIC Team members.

Table 1

Mode	Description
Static Teaming	Requires you to manually configure both the switch and the host to identify which links form the team. Because this is a statically configured solution, there is no additional protocol to assist the switch and the host to identify incorrectly plugged cables or other errors that could cause the team to fail to perform. This mode is typically supported by server-class switches.
Link Aggregation Control Protocol (LACP)	Unlike Static Teaming, LACP Teaming mode dynamically identifies links that are connected between the host and the switch. This dynamic connection enables the automatic creation of a team and, in theory but rarely in practice, the expansion and reduction of a team simply by the transmission or receipt of LACP packets from the peer entity. All server-class switches support LACP, and all require the network operator to administratively enable LACP on the switch port. When you configure a Teaming mode of LACP, NIC Teaming always operates in LACP's Active mode. By default, NIC Teaming uses a short timer (3 seconds), but you can

Table 1

Mode	Description
	configure a long timer (90 seconds) with `Set-NetLbfoTeam`.

- **Load balancing mode**. The options for Load Balancing distribution mode are **Address Hash**, **Hyper-V Port**, and **Dynamic**.
  - **Address Hash.** With Address Hash, this mode creates a hash based on address components of the packet, which then get assigned to one of the available adapters. Usually, this mechanism alone is sufficient to create a reasonable balance across the available adapters.
  - **Hyper-V Port.** With Hyper-V Port, NIC Teams configured on Hyper-V hosts give VMs independent MAC addresses. The VMs MAC address or the VM ported connected to the Hyper-V switch, can be used to divide network traffic between NIC Team members. You cannot configure NIC Teams that you create within VMs with the Hyper-V Port load balancing mode. Instead, use the Address Hash mode.
  - **Dynamic.** With Dynamic, outbound loads are distributed based on a hash of the TCP ports and IP addresses. Dynamic mode also rebalances loads in real time so that a given outbound flow may move back and forth between team members. Inbound loads, on the other hand, get distributed the same way as Hyper-V Port. In a nutshell, Dynamic mode utilizes the best aspects of both Address Hash and Hyper-V Port and is the highest performing load balancing mode.
- **Standby adapter**. The options for Standby Adapter are **None (all adapters Active)** or your selection of a specific network adapter in the NIC Team that acts as a Standby adapter.

Tip

If you are configuring a NIC Team in a virtual machine (VM), you must select a **Teaming mode** of *Switch Independent* and a **Load balancing mode** of *Address Hash*.

7. If you want to configure the primary team interface name or assign a VLAN number to the NIC Team, click the link to the right of **Primary team interface**.

   The **New team interface** dialog box opens.

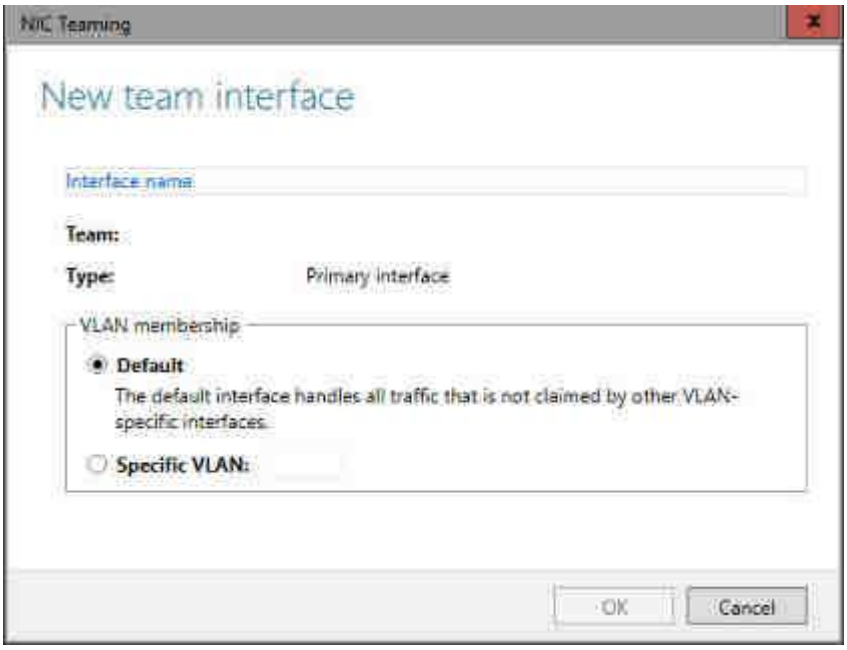

8. Depending on your requirements, do one of the following:
   - Provide a tNIC interface name.
   - Configure VLAN membership: click **Specific VLAN** and type the VLAN information. For example, if you want to add this NIC Team to the accounting VLAN number 44, Type Accounting 44 - VLAN.

9. Click **OK**.

Congratulations! You've created a new NIC Team on a host computer or VM.

Troubleshooting NIC Teaming

https://docs.microsoft.com/en-us/windows-server/networking/technologies/nic-teaming/troubleshooting-nic-teaming

Applies to: Windows Server 2022, Windows Server 2019, Windows Server 2016, Azure Stack HCI, versions 21H2 and 20H2

In this topic, we discuss ways to troubleshoot NIC Teaming, such as hardware and physical switch securities. When hardware implementations of standard protocols don't conform to specifications, NIC Teaming performance might be affected. Also, depending on the configuration, NIC Teaming may send packets from the same IP address with multiple MAC addresses that can trip security features on the physical switch.

Hardware that doesn't conform to specification

During normal operation, NIC Teaming may send packets from the same IP address, yet with multiple MAC addresses. According to protocol standards, the receivers of these packets must resolve the IP address of the host or VM to a specific MAC address rather than responding to the MAC address from the receiving packet. Clients that correctly implement the address resolution protocols (ARP and NDP) send packets with the correct destination MAC addresses—that is, the MAC address of the VM or host that owns that IP address.

Some embedded hardware does not correctly implement the address resolution protocols, and also might not explicitly resolve an IP address to a MAC address using ARP or NDP. For example, a storage area network (SAN) controller might perform in this manner. Non-conforming devices copy the source MAC address from a received packet and then use it as the destination MAC address in the corresponding outgoing packets, which results in packets sent to the wrong destination MAC address. Because of this, the packets are dropped by the Hyper-V Virtual Switch because they don't match any known destination.

If you are having trouble connecting to SAN controllers or other embedded hardware, you should take packet captures to determine if your hardware is correctly implementing ARP or NDP, and contact your hardware vendor for support.

Physical switch security features

Depending on the configuration, NIC Teaming may send packets from the same IP address with multiple source MAC addresses tripping up security features on the physical switch. For example, Dynamic ARP inspection or IP source guard, especially if the physical switch is not aware that the ports are part of a team, which occurs when you configure NIC Teaming in Switch Independent mode. Inspect the switch logs to determine if switch security features are causing connectivity problems.

Disabling and enabling network adapters by using Windows PowerShell

A common reason for a NIC Team to fail is that the team interface is disabled, and in many cases, by accident when running a sequence of commands. This particular sequence of commands does not enable all of the NetAdapters disabled because disabling all of the underlying physical members of NICs removes the NIC team interface.

In this case, the NIC team interface no longer shows in Get-NetAdapter, and because of this, **Enable-NetAdapter *** does not enable the NIC Team. The **Enable-NetAdapter *** command does, however, enable the member NICs, which then (after a short time) recreates the team interface. The team interface remains in the "disabled" state until re-enabled, allowing network traffic to begin flowing.

The following Windows PowerShell sequence of commands may disable the team interface by accident:

PowerShell

```
Disable-NetAdapter *
Enable-NetAdapter *
```