

SonicOS 7.0

High Availability

Administration Guide



Contents

High Availability	4
About High Availability	4
High Availability Terminology	5
High Availability Modes	6
High Availability Encryption	6
Crash Detection	7
Virtual MAC Address	7
Dynamic WAN Interfaces with PPPoE HA	8
Stateful Synchronization with DHCP	8
Stateful Synchronization with DNS Proxy	8
About HA Monitoring	8
About Active/Standby HA	9
Benefits of Active/Standby HA	10
Working of Active/Standby HA	10
About Stateful Synchronization	11
Benefits of Stateful Synchronization	11
How Does Stateful Synchronization Work?	11
Example of Stateful Synchronization	12
Active/Standby Prerequisites	13
Supported Platforms and Licensing for HA	13
Physically Connecting Your Security Appliances	14
Maintenance	15
Removing an HA Association	15
Replacing a SonicWall Security Appliance	16
High Availability Status	18
Active/Standby High Availability Status	18
High Availability Status	19
High Availability Config	20
High Availability Licenses	21
Configuring High Availability	22
Configuring Active/Standby High Availability Settings	22
Configuring HA with Dynamic WAN Interfaces	24
Configuring High Availability in the Cloud Platform	26
Set up an Active/Standby High Availability Configuration Using Azure	26
Install the Custom Template	28

Enable Identity of Both Virtual Machines (HA1 and HA2)	30
Role Assignment	31
Check the Networking Tab	33
Configuring Active NSv Firewall Using the Associated Public IP	34
Configuring Standby NSv Firewall Using the Associated Public IP	36
Enable the L3 Mode	38
Configuring Active NSv Firewall Using the Floating Public IP	39
Configuring HA to Active/Standby with L3 HA link	39
Adding Additional Floating Public IP	40
Fine Tuning High Availability	43
Advanced Settings	43
Configuring Advanced High Availability Settings	43
Monitoring High Availability	46
Configuring Active/Standby High Availability Monitoring	46
IPv6 High Availability Monitoring	47
IPv6 HA Monitoring Considerations	48
Azure Use Cases	49
Use Case 1: Manage Azure HA Firewall	49
Use Case 2: Forward LAN traffic to the External Network through the Gateway after HA Failover	49
Use Case 3: Configure DNAT on Azure HA Firewall	49
Use Case 4: Configure DNAT on Azure HA Firewall (Need to move multiple floating IPs support)	51
SonicWall Support	53
About This Document	54

High Availability

This section provides conceptual information about SonicOS (HA) in SonicOS and describes how to connect the Security Appliances for HA.

Topics:

- [About High Availability](#)
- [About Active/Standby HA](#)
- [About Stateful Synchronization](#)
- [Active/Standby Prerequisites](#)
- [Maintenance](#)

About High Availability

High Availability is designed to alleviate or eliminate:

- System downtime
- Single points of failure
- Increased system load

High Availability (HA) is a redundancy design that allows two identical firewalls running SonicOS to be configured to provide a reliable, continuous connection. One firewall is configured as the Primary unit, and an identical firewall is configured as the Secondary unit. In the event of a failure on the Primary firewall, the Secondary firewall takes over to secure a reliable connection between the connected networks. Two firewalls configured in this way are known as a High Availability Pair (HA Pair).

High Availability provides a way to share **SonicWall licenses** between two firewalls when one is acting as a high-availability system for the other. To use this feature, you must register the firewalls on MySonicWall.com as Associated Products.

① | **NOTE:** Both firewalls must be the same SonicWall model and firmware.

Topics:

- [High Availability Terminology](#)
- [High Availability Modes](#)
- [High Availability Encryption](#)
- [Crash Detection](#)
- [Virtual MAC Address](#)
- [Dynamic WAN Interfaces with PPPoE HA](#)
- [Stateful Synchronization with DHCP](#)
- [Stateful Synchronization with DNS Proxy](#)
- [About HA Monitoring](#)

High Availability Terminology

HIGH AVAILABILITY TERMINOLOGY

Active	The operative condition of a hardware unit. The Active identifier is a logical role that can be assumed by either a Primary or Secondary hardware unit.
Failover	The actual process in which the Standby unit assumes the Active role following a qualified failure of the Active unit. Qualification of failure is achieved by various configurable physical and logical monitoring facilities described in Configuring High Availability .
HA	High Availability: non-stateful, hardware failover capability.
PPP	Point-to-point protocol that provides a standard method for transporting multi-protocol diagrams over point-to-point links.
PPPoE	A method for transmitting PPP over ethernet.
PPPoE HA	HA PPPoE support function without State.
Preempt	Applies to a post-failover condition in which the Primary unit has failed, and the Secondary unit has assumed the Active role. Enabling Preempt causes the Primary unit to seize the Active role from the Secondary after the Primary has been restored to a verified operational state.
Primary	The principal hardware unit itself. The Primary identifier is a manual designation and is not subject to conditional changes. Under normal operating conditions, the Primary hardware unit operates in an Active role.
Secondary (Backup)	The subordinate hardware unit itself. The Secondary identifier is a relational designation and is assumed by a unit when paired with a Primary unit. Under normal operating conditions, the Secondary unit operates in a Standby mode. Upon failure of the Primary unit, the Secondary unit assumes the Active role.
Standby (Idle)	The passive condition of a hardware unit. The Standby identifier is a logical role that can be assumed by either a Primary or Secondary hardware unit. The Standby unit assumes the Active role upon a determinable failure of the Active unit.

High Availability Modes

High Availability has several operation modes, which can be selected on **Device > High Availability > Settings**.

Choosing the right High Availability Operation mode depends on understanding the network in question, its purpose and operational needs. In planning, the administrator should understand:

- Operational requirements for up time
- Repercussions of failure
- Calculated risk to operations

Each operation mode satisfies a different scenario and without knowing the goals of High Availability, administrators risk building an unsatisfactory solution. Understanding the operational mode and how they map to requirements is fundamental. This Active/Standby mode may be further defined as to whether they are stateless or stateful to a secondary device.

Operational modes are:

- **None**—Selecting None activates a standard high availability configuration and hardware failover functionality, with the option of enabling Stateful HA.
- **Active/Standby Stateless**—Active/Standby mode provides basic high availability with the configuration of two identical firewalls as a High Availability Pair. The Active unit handles all traffic, while the Standby unit shares its configuration settings and can take over at any time to provide continuous network connectivity if the Active unit stops working. By default, Active/Standby mode is stateless, meaning that network connections and VPN tunnels must be re-established after a failover.
- **Active/Standby Stateful**—Stateful Synchronization can be licensed and enabled with Active/Standby mode. In this Stateful HA mode, the dynamic state is continuously synchronized between the Active and Standby units.

Network connections and VPN tunnel information are continuously synchronized between the two units so that the Secondary can seamlessly assume all network responsibilities if the Primary firewall fails.

When the Active unit encounters a fault condition, stateful failover occurs as the Standby firewall takes over the Active role with no interruptions to the existing network connections.

① | **NOTE:** Not all information is synchronized in a stateful configuration.

High Availability Encryption

High Availability encryption adds security to the communication between appliances in a HA pair. HA control messages between active and standby firewalls, such as heartbeats, configuration sync and HA state information, are encrypted to ensure security for inter-node communication.

This option is available in Active-Standby HA mode only and does not apply to messages exchanged for stateful synchronization even in Active-Standby mode. Discovery messages (find-peer and found-peer) are transmitted without encryption. After the discovery stage, however, all control messages are encrypted between the firewalls:

- Heartbeats
- Messages used for incremental config updates
- prefSync messages
- Various messages for sending HA commands between the firewall pair
- Firmware sync messages

Crash Detection

The HA feature has a thorough self-diagnostic mechanism for both the Active and Standby Security Appliances. The failover to the standby unit occurs when critical services are affected, physical or logical link failure is detected on monitored interfaces, or when the Security Appliance loses power.

The self-checking mechanism is managed by software diagnostics, which check the complete system integrity of the Security Appliance. The diagnostics check internal system statuses, system process statuses, and network connectivity. There is a weighting mechanism on both sides to decide which side has better connectivity to avoid potential failover looping.

Critical internal system processes such as NAT, VPN, and DHCP (among others) are checked in real time. The failing service is isolated as early as possible, and the failover mechanism repairs it automatically.

Virtual MAC Address

The Virtual MAC address allows the High Availability pair to share the same MAC address, which dramatically reduces convergence time following a failover. Convergence time is the amount of time it takes for the devices in a network to adapt their routing tables to the changes introduced by high availability.

Without Virtual MAC enabled, the Active and Standby Security Appliances each use their own MAC addresses. Because the Security Appliances are using the same IP address, when a failover occurs, it breaks the mapping between the IP address and MAC address in the ARP cache of all clients and network resources. The Standby Security Appliance must issue an ARP request, announcing the new MAC address/IP address pair. Until this ARP request propagates through the network, traffic intended for the Active Security Appliance's MAC address can be lost.

The Virtual MAC address greatly simplifies this process by using the same MAC address for both the Active and Standby Security Appliances. When a failover occurs, all routes to and from the Active Security Appliance are still valid for the Standby Security Appliance. All clients and remote sites continue to use the same Virtual MAC address and IP address without interruption.

By default, this Virtual MAC address is provided by the SonicWall firmware and is different from the physical MAC address of either the Primary or Secondary Security Appliances. This eliminates the possibility of configuration errors and ensures the uniqueness of the Virtual MAC address, which prevents possible conflicts. Optionally, you can manually configure the Virtual MAC address on `DEVICE | High Availability > Monitoring`.

The Virtual MAC setting is available even if Stateful High Availability is not licensed. When Virtual MAC is enabled, it is always used even if Stateful Synchronization is not enabled.

Dynamic WAN Interfaces with PPPoE HA

① **NOTE:** Dynamic WAN interfaces with PPPoE HA is not supported on the NSsp 15700. Only the DHCP Server dynamic WAN mode is supported.

PPPoE can be enabled on interfaces in non-stateful mode, HA Active/Standby mode. PPPoE HA provides HA where a Standby Security Appliance assumes connection to the PPPoE server when the Active Security Appliance fails.

① **NOTE:** One WAN interface must be configured as PPPoE; see Configuring a WAN Interface section in the *SonicOS 7.0 Firewall Network* document available at <https://www.sonicwall.com/support/technical-documentation/>.

After the Active unit connects to the PPPoE server, the Security Appliance synchronizes the PPPoE session ID and server name to the Standby unit.

When the Active Security Appliance fails, it terminates the PPPoE HA connection on the client side by timing out. The Secondary Security Appliance connects to the PPPoE server, terminates the original connection on the server side, and starts a new PPPoE connection. All pre-existing network connections are rebuilt, the PPPoE sessions are re-established, and the PPP process is renegotiated.

Stateful Synchronization with DHCP

DHCP can be enabled on interfaces in both Active/Standby non-stateful and Stateful Synchronization modes.

Only the Active Security Appliance can get a DHCP lease. The Active Security Appliance synchronizes the DHCP IP address along with the DNS and gateway addresses to the Standby Security Appliance. The DHCP client ID is also synchronized, allowing this feature to work even without enabling Virtual MAC.

During a failover, the Active Security Appliance releases the DHCP lease and, as it becomes the Active unit, the Standby Security Appliance renews the DHCP lease using the existing DHCP IP address and client ID. The IP address does not change, and network traffic, including VPN tunnel traffic, continues to pass.

If the Active Security Appliance does not have an IP address when failover occurs, the Standby Security Appliance starts a new DHCP discovery.

Stateful Synchronization with DNS Proxy

DNS Proxy supports stateful synchronization of the DNS cache. When the DNS cache is added, deleted, or updated dynamically, it synchronizes to the idle Security Appliance.

About HA Monitoring

On **DEVICE | High Availability > Monitoring**, you can configure both physical and logical interface monitoring:

- By enabling physical interface monitoring, you enable link detection for the designated HA interfaces. The link is sensed at the physical layer to determine link viability.
- Logical monitoring involves configuring the SonicWall to monitor a reliable device on one or more of the connected networks.

Failure to periodically communicate with the device by the Active unit in the HA Pair triggers a failover to the Standby unit. If neither unit in the HA Pair can connect to the device, no action is taken.

The Primary and Secondary IP addresses configured on **DEVICE | High Availability > Monitoring** can be configured on interfaces, and are used for multiple purposes:

- As independent management addresses for each unit (supported on all physical and virtual interfaces)
- To allow synchronization of licenses between the Standby unit and the SonicWall licensing server
- As the source IP addresses for the probe pings sent out during logical monitoring

Configuring unique management IP addresses for both units in the HA Pair allows you to log in to each unit independently for management purposes. Note that non-management traffic is ignored if it is sent to one of these IP addresses. The Primary and Secondary Security Appliances' unique IP addresses cannot act as an active gateway; all systems connected to the interface need to use IP address configured on the network interfaces as their gateway.

If WAN monitoring IP addresses are configured, then X0 monitoring IP addresses are not required. If WAN monitoring IP addresses are not configured, then X0 monitoring IP addresses are required, because in such a scenario the Standby unit uses the X0 monitoring IP address to connect to the licensing server with all traffic routed through the Active unit.

① | **NOTE:** X0 needs to have a routeable IP.

The management IP address of the Standby unit is used to allow license synchronization with the SonicWall licensing server, which handles licensing on a per-Security Appliance basis (not per-HA Pair). Even if the Secondary unit was already registered on MySonicWall before creating the HA association, you must use the link on **Device | Settings > Licenses** to connect to the SonicWall server while accessing the Secondary Security Appliance through its management IP address (for more information, see *SonicOS 7.0 Settings* document).

When using logical monitoring, the HA Pair pings the specified Logical Probe IP address target from the Primary as well as from the Secondary unit. The IP address set in the Primary IP Address or Secondary IP Address field is used as the source IP address for the ping. If both units can successfully ping the target, no failover occurs. If both cannot successfully ping the target, no failover occurs, as SonicOS assumes that the problem is with the target, and not the Security Appliances. If one Security Appliance can ping the target but the other cannot, however, the HA Pair failovers to the unit that can ping the target.

The configuration tasks on **DEVICE | High Availability > Monitoring** are performed on the Active unit and then are automatically synchronized to the Standby.

About Active/Standby HA

HA allows two identical Security Appliances running SonicOS to be configured to provide a reliable, continuous connection to their connected networks. One Security Appliance is configured as the Active unit, and an identical

Security Appliance is configured as the Standby unit. In the event of the failure of the Active Security Appliance, the Standby Security Appliance takes over to secure a reliable connection between the connected networks. Two Security Appliances configured in this way are also known as a High Availability Pair (HA Pair).

Topics:

- [Benefits of Active/Standby HA](#)
- [Working of Active/Standby HA](#)

Benefits of Active/Standby HA

- **Increased network reliability** - In a High Availability configuration, the Secondary Security Appliance assumes all network responsibilities when the Primary unit fails, ensuring a reliable connection between the connected networks.
- **Cost-effectiveness** - is a cost-effective option for deployments that provide high availability by using redundant Security Appliances. You do not need to purchase a second set of licenses for the Secondary unit in a High Availability Pair.
- **Virtual MAC for reduced convergence time after failover** - The Virtual MAC address setting allows the HA Pair to share the same MAC address, which dramatically reduces convergence time following a failover. Convergence time is the amount of time it takes for the devices in a network to adapt their routing tables to the changes introduced by high availability. By default, the Virtual MAC address is provided by the SonicWall firmware and is different from the physical MAC address of either the Primary or Secondary Security Appliances.

Working of Active/Standby HA

HA requires one SonicWall Security Appliance configured as the Primary SonicWall, and an identical Security Appliance configured as the Secondary SonicWall. During normal operation, the Primary SonicWall is in an Active state and the Secondary SonicWall in an Standby state. If the Primary device loses connectivity, the Secondary SonicWall transitions to Active mode and assumes the configuration and role of Primary, including the interface IP addresses of the configured interfaces.

Basic Active/Standby HA provides stateless high availability. After a failover to the Secondary Security Appliance, all the pre-existing network connections must be re-established, including the VPN tunnels that must be re-negotiated. Stateful Synchronization can be licensed and enabled separately. For more information, see [About Stateful Synchronization](#).

The failover applies to loss of functionality or network-layer connectivity on the Active SonicWall. The failover to the Standby SonicWall occurs when critical services are affected, physical or logical link failure is detected on monitored interfaces, if, not when the Active SonicWall loses power.

There are two types of synchronization for all configuration settings:

- **Incremental** - If the timestamps are in sync and a change is made on the Active unit, an incremental synchronization is pushed to the Standby unit.

- **Complete** - If the timestamps are out of sync and the Standby unit is available, a complete synchronization is pushed to the Standby unit. When incremental synchronization fails, a complete synchronization is automatically attempted
- ① | **NOTE:** The complete synchronization reboots the Standby unit.

About Stateful Synchronization

Stateful Synchronization provides dramatically improved failover performance. When enabled, the network connections and VPN tunnel information is continuously synchronized between the two units so that the Standby can seamlessly assume all network responsibilities if the Active Security Appliance fails, with no interruptions to existing network connections.

Topics:

- [Benefits of Stateful Synchronization](#)
- [How Does Stateful Synchronization Work?](#)
- [Example of Stateful Synchronization](#)

Benefits of Stateful Synchronization

- **Improved reliability** - By synchronizing most critical network connection information, Stateful Synchronization prevents down time and dropped connections in case of Security Appliance failure.
- **Faster failover performance** - By maintaining continuous synchronization between the Active and Standby Security Appliances, Stateful Synchronization enables the Standby Security Appliance to take over in case of a failure with virtually no down time or loss of network connections.
- **Minimal impact on CPU performance** - Typically less than 1% usage.
- **Minimal impact on bandwidth** - Transmission of synchronization data is throttled so as not interfere with other data.

How Does Stateful Synchronization Work?

Stateful Synchronization is not load-balancing. It is an active-standby configuration where the Primary Security Appliance handles all traffic. When Stateful Synchronization is enabled, the Primary Security Appliance actively communicates with the Secondary to update most network connection information. As the Primary Security Appliance creates and updates network connection information (such as VPN tunnels, active users, connection cache entries), it immediately informs the Secondary Security Appliance. This ensures that the Secondary Security Appliance is always ready to transition to the Active state without dropping any connections.


The synchronization traffic is throttled to ensure that it does not interfere with regular network traffic. All configuration changes are performed on the Active Security Appliance and automatically propagated to the Standby Security Appliance. The High Availability pair uses the same LAN and WAN IP addresses—regardless of which Security Appliance is currently Active.

When using SonicWall Network Security Manager (NSM) to manage the Security Appliances, NSM logs into the shared WAN IP address. In case of a failover, NSM administration continues seamlessly, and NSM administrators currently logged into the Security Appliance are not logged out; however, **Get** and **Post** commands may result in a time out with no reply returned.

Synchronized and non-synchronized information table lists the information that is synchronized and information that is not currently synchronized by Stateful Synchronization.

SYNCHRONIZED AND NON-SYNCHRONIZED INFORMATION

Information that is Synchronized	Information that is not Synchronized
VPN information	Dynamic WAN clients (L2TP, PPPoE, and PPTP)
Basic connection cache	Deep Packet Inspection (GAV, IPS, and Anti Spyware)
FTP	IPHelper bindings (such as NetBIOS and DHCP)
Oracle SQL*NET	SYNFlood protection information
Real Audio	Content Filtering Service information
RTSP	VoIP protocols
GVC information	Dynamic ARP entries and ARP cache time outs
Dynamic Address Objects	Active wireless client information
DHCP server information	Wireless client packet statistics
Multicast and IGMP	Rogue AP list
Active users	
ARP	
SonicPoint and SonicWave status	
Wireless guest status	
Weighted Load Balancing information	
Dynamic Routing Configuration	

 **WARNING:** The configuration is synchronized, but the routing table has to be rebuilt in a failover.

Example of Stateful Synchronization

In case of a failover, the following sequence of events occurs:

1. A PC user connects to the network, and the Active Security Appliance creates a session for the user.
2. The Active Security Appliance synchronizes with the Standby Security Appliance. The Standby now has all of the user's session information.

3. The administrator restarts the Active unit.
4. The Standby unit detects the restart of the Active unit and switches from Standby to Active.
5. Now Active Security Appliance begins to send gratuitous ARP messages to the connected switches using the same Virtual MAC address and IP address as the Active Security Appliance. No routing updates are necessary for downstream or upstream network devices.
6. When the PC user attempts to access a Web page, now Active Security Appliance has all of the user's session information and is able to continue the user's session without interruption.

Active/Standby Prerequisites

This section lists the supported platforms, provides recommendations and requirements for physically connecting the units, and describes how to register, associate, and license the units.

Topics:

- [Supported Platforms and Licensing for HA](#)
- [Physically Connecting Your Security Appliances](#)

Supported Platforms and Licensing for HA

Licenses included with the purchase of a SonicWall Security Appliance are shown in **HA licenses available with SonicWall Security Appliances** table. Some platforms require additional licensing to use the HA features.

The HA licenses included with the purchase of the SonicWall Security Appliance are shown in **HA licenses available with SonicWall Security Appliances**. Some platforms require additional licensing to use the Stateful Synchronization feature. SonicOS Expanded licenses or High Availability licenses can be purchased on MySonicWall or from a SonicWall reseller.

NOTE: Stateful High Availability licenses must be activated on each Security Appliance, either by registering the unit on MySonicWall from the SonicOS management interface, or by applying the license keyset to each unit if Internet access is not available.

HA Licenses Available With Sonicwall Network Security Firewalls

Platform	Active/Standby HA	Stateful HA
TZ270/TZ270 W	Included	Optional
TZ370/TZ370 W	Included	Optional
TZ470/TZ470 W	Included	Optional
TZ570/TZ570 W/TZ570 P	Included	Optional
TZ670	Included	Optional
NSA 2700	Included	Optional
NSA 3700	Included	Optional

HA Licenses Available With Sonicwall Network Security Firewalls

NSA 4700	Included	Included
NSA 6700	Included	Included
NSsp 10700	Included	Included
NSsp 11700	Included	Included
NSSP 13700	Included	Included
NSSP 15700	Included	Included
NSv 270	Included	Included
NSv 470	Included	Included
NSv 870	Included	Included

- ① **NOTE:** HA Stateful licensing is not standard across all models. Enterprise class models often include HA Stateful licensing.

You can view system licenses on **DEVICE | Settings > Licenses**. This page also provides a way to log into MySonicWall and to apply licenses to a Security Appliance. For further information, see *SonicOS 7.0 Settings* document.

There is also a way to synchronize licenses for an HA pair whose Security Appliances do not have Internet access. When live communication with SonicWall's licensing server is not permitted due to network policy, you can use license keysets to manually apply security services licenses to your Security Appliances. When you register a Security Appliance on MySonicWall, a license keyset is generated for the Security Appliance. If you add a new security service license, the keyset is updated. However, until you apply the licenses to the Security Appliance, it cannot perform the licensed services.

- ① **IMPORTANT:** In a High Availability deployment without Internet connectivity, you must apply the license keyset to both of the Security Appliances in the HA pair.
- ① **IMPORTANT:** Even if you first register your Security Appliances on MySonicWall, you must individually register both the Primary and the Secondary Security Appliances from the SonicOS management interface while logged into the individual management IP address of each Security Appliance. This allows the Secondary unit to synchronize with the SonicWall license server and share licenses with the associated Primary Security Appliance. When Internet access is restricted, you can manually apply the shared licenses to both Security Appliances.

Physically Connecting Your Security Appliances

- ① **NOTE:** For complete procedures for connecting your Security Appliances, see the Quick Start Guide for your Security Appliance.
- ① **NOTE:** If you are connecting the Primary and Secondary Security Appliances to an Ethernet switch that uses the spanning tree protocol, be aware that it may be necessary to adjust the link activation time on the switch port to which the SonicWall interfaces connect.

High Availability requires additional physical connections among the affected SonicWall Security Appliances.

In any High Availability deployment, you must physically connect the interfaces of all units to the appropriate switches.

It is important that the X0 interfaces from all units be connected to the same broadcast domain. Otherwise, traffic failover does not work. Also, X0 is the default redundant HA port; if the normal HA Control link fails, X0 interfaces are used to communicate heartbeats between units. Without X0 in the same broadcast domain, both units would become active if the HA Control link fails.

① | **NOTE:** If X0 interfaces are not used to communicate, the units should be connected directly to each other.

① | **TIP:** SonicOS Security Appliances now allow heartbeats to be exchanged between an HA pair across the MGMT interface in addition to the HA control interface.

A WAN connection to the Internet is useful for registering your Security Appliances on MySonicWall and for synchronizing licensing information. Unless live communication with SonicWall's licensing server is not permitted due to network policy, the WAN interface should be connected before registration and licensing are performed.

Maintenance

Topics:

- [Removing an HA Association](#)
- [Replacing a SonicWall Security Appliance](#)

Removing an HA Association

You can remove the association between two SonicWall Security Appliances on MySonicWall at any time. You might need to remove an existing HA association if you replace a Security Appliance or reconfigure your network. For example, if one of your SonicWall Security Appliances fails and you need to replace it, or you might need to switch the HA Primary Security Appliance with the Secondary, or HA Secondary, unit after a network reconfiguration. In either case, you must first remove the existing HA association, and then create a new association that uses a new Security Appliance or changes the parent-child relationship of the two units (see [Replacing a SonicWall Security Appliance](#)).

To remove the association between two registered SonicWall Security Appliances:

1. Log in to MySonicWall.
2. In the left navigation bar, navigate to **My Workspace > Tenant Products**.
3. Scroll down to find the secondary Security Appliance from which you want to remove associations. Click the **serial number**.
4. On the **Products Details** page, scroll down to the **Parent Products** section, just below the **Associated Products** section.

5. Under **Parent Products**, to remove the association for this Security Appliance:
 - a. Click **Remove** under **ACTIONS**.
 - b. Wait for the page to reload.
 - c. Scroll down.
 - d. Click **Remove** again.

Replacing a SonicWall Security Appliance

If your SonicWall Security Appliance has a hardware failure while still under warranty, SonicWall will replace it. In this case, you need to remove the HA association containing the failed Security Appliance in MySonicWall, and add a new HA association that includes the replacement. If you contact SonicWall Technical Support to arrange the replacement (known as an RMA), Support can help with this process.

After replacing the failed Security Appliance in your equipment rack with the new unit, you can update MySonicWall and your SonicOS configuration.

Replacing a failed HA Primary unit is slightly different than replacing an HA Secondary unit. Both procedures are provided in these sections:

- [Replacing an HA Primary Unit](#)
- [Replacing an HA Secondary Unit](#)

Replacing an HA Primary Unit

To replace an HA Primary unit:

1. In the SonicOS management interface of the remaining SonicWall Security Appliance (the Secondary unit), on the High Availability page, uncheck **Enable High Availability** to disable it.
2. Check **Enable High Availability**.
The old Secondary unit now becomes the Primary unit. Its serial number is automatically displayed in the Primary SonicWall Serial Number field.
3. Type the serial number for the replacement unit into the **Secondary Device** field.
4. Click **Synchronize Settings**.
5. On MySonicWall, remove the old HA association. See [Removing an HA Association](#).
6. On MySonicWall, register the replacement SonicWall Security Appliance and create an HA association with the new Primary (original Secondary) unit as the HA Primary, and the replacement unit as the HA Secondary.
7. Contact SonicWall Technical Support to transfer the security services licenses from the former HA Pair to the new HA Pair.
This step is required when the HA Primary unit has failed because the licenses are linked to the Primary unit in an HA Pair.

Replacing an HA Secondary Unit

To replace an HA Secondary unit:

1. On MySonicWall, remove the old HA association as described in [Removing an HA Association](#).
2. On MySonicWall, register the replacement SonicWall Security Appliance.
3. Create an HA association with the original HA Primary, using the replacement unit as the HA Secondary as described in [Replacing an HA Primary Unit](#).


High Availability Status

Topics:

- [Active/Standby High Availability Status](#)

The **DEVICE | High Availability > Status** page displays the current status of the High Availability pair, including state of primary and secondary units, mode and link configuration, and licenses.

At the top of the page, you can see which unit you are logged into, **Primary** or **Secondary**, and whether the unit is in the **Active** or **Standby** state.



In the event that the Primary unit has a failure, you can view the status by accessing the management interface of the Secondary unit at the Primary unit virtual IP address or the Secondary unit unique IP address. When the Active unit restarts after a failure, it is accessible using the unique IP address created on the **DEVICE | High Availability > Monitoring** page. If preempt mode is enabled, the Primary unit immediately takes over as the Active firewall and the Secondary unit returns to Standby status.

Active/Standby High Availability Status

Active/Standby High Availability provides basic high availability with the configuration of two identical firewalls as a High Availability pair. On a firewall that belongs to an Active/Standby HA pair, the **DEVICE | High Availability > Status** page displays information about the state, configuration, and licenses on the HA pair.

HIGH AVAILABILITY STATUS	
Status	Primary ACTIVE
Primary State	ACTIVE
Secondary State	STANDBY
Active Up Time	11 Days 18:42:43
Found Peer	Yes
Settings Synchronized	Yes
Stateful HA Synchronized	Yes
HIGH AVAILABILITY CONFIG	
HA Mode	Active / Standby
HA Control Link	X6 1000 Mbps full-duplex
HA Data Link	X7 1000 Mbps full-duplex
HIGH AVAILABILITY LICENSES	
Primary Stateful HA Licensed	Yes
Secondary Stateful HA Licensed	Yes

Topics:

- [High Availability Status](#)
- [High Availability Config](#)
- [High Availability Licenses](#)

High Availability Status

The **High Availability Status** section on the **DEVICE | High Availability > Status** page displays the following information:

- **Status** - Indicates the High Availability status of the current firewall. The possible values are:
 - **Primary Active** - Indicates that the current appliance is the Primary unit in the ACTIVE state.
 - **Primary Standby** - Indicates that the current appliance is the Primary unit in the STANDBY state.
 - **Primary Disabled** - Indicates that the current appliance is the Primary unit, but High Availability has not been enabled.
 - **Primary not in a steady state** - Indicates that the current appliance is the Primary unit, HA is enabled, and the appliance is neither in the ACTIVE nor the STANDBY state.
- **Primary State** - Indicates the current state of the Primary appliance as a member of an HA Pair. The Primary State field is displayed on both the Primary and the Secondary appliances. The possible values are:
 - **ACTIVE** - Indicates that the Primary unit is handling all the network traffic except management/monitoring/licensing traffic destined to the standby unit.

- **STANDBY** - Indicates that the Primary unit is passive and is ready to take over on a failover.
 - **ELECTION** - Indicates that the Primary and Secondary units are negotiating which should be the ACTIVE unit.
 - **SYNC** - Indicates that the Primary unit is synchronizing settings or firmware to the Secondary.
 - **ERROR** - Indicates that the Primary unit has reached an error condition.
 - **REBOOT** - Indicates that the Primary unit is rebooting.
 - **NONE** - When viewed on the Primary unit, **NONE** indicates that HA is not enabled on the Primary. When viewed on the Secondary unit, **NONE** indicates that the Secondary unit is not receiving heartbeats from the Primary unit.
- **Secondary State** - Indicates the current state of the Secondary appliance as a member of an HA Pair. The Secondary State field is displayed on both the Primary and the Secondary appliances. The possible values are:
 - **ACTIVE** - Indicates that the Secondary unit is handling all the network traffic except management/monitoring/licensing traffic destined to the standby unit.
 - **STANDBY** - Indicates that the Secondary unit is passive and is ready to take over on a failover.
 - **ELECTION** - Indicates that the Secondary and Primary units are negotiating which should be the ACTIVE unit.
 - **SYNC** - Indicates that the Secondary unit is synchronizing settings or firmware with the Primary.
 - **ERROR** - Indicates that the Secondary unit has reached an error condition.
 - **REBOOT** - Indicates that the Secondary unit is rebooting.
 - **NONE** - When viewed on the Secondary unit, **NONE** indicates that HA is not enabled on the Secondary. When viewed on the Primary unit, **NONE** indicates that the Primary unit is not receiving heartbeats from the Secondary unit.
 - **Active Up Time** - Indicates how long the current Active firewall has been Active, since it last became Active. If the unit is not part of an HA pair, this line displays High Availability Disabled.
 - **Found Peer** - Indicates if the Primary unit has discovered the Secondary unit. Possible values are **Yes** and **No**.
 - **Settings Synchronized** - Indicates if HA settings are synchronized between the Primary and Secondary units. Possible values are **Yes** and **No**.
 - **Stateful HA Synchronized** - Indicates if stateful synchronization settings are synchronized between the Primary and Secondary units. Possible values are **Yes** and **No**.

High Availability Config

The **High Availability Config** section on the **Device > High Availability > Settings** page provides the following information:

- **HA Mode** - Indicates one of:
 - **None** - High Availability is not enabled on the unit.
 - **Active/Standby** - Active/Standby mode provides basic high availability with the configuration of two identical firewalls as a High Availability Pair. By default, Active/Standby mode is stateless, meaning that network connections and VPN tunnels must be re-established after a failover. To avoid this, Stateful Synchronization can be licensed and enabled with Active/Standby mode.
- **HA Control Link** - Indicates the port, speed, and duplex settings of the HA control link, such as **X6 1 Gbps Full Duplex**. When High Availability is not enabled, the field displays not configured. The HA control link is used to communicate heartbeats and other control traffic between the units. If the HA control link fails, X0 is used to communicate heartbeats between units; therefore heartbeats on both units should be in the same broadcast domain.
- **HA Data Link** - Indicates the port, speed, and duplex settings of the HA data link, such as **X7 1 Gbps Full Duplex**. When High Availability is not enabled, the field displays not configured. The HA data link is used to transfer stateful data to keep session data synchronized between the units. The HA Data Link is not required when running in non-stateful HA.

High Availability Licenses

The **High Availability Licenses** section on the **DEVICE | High Availability > Status** page provides the following information:

- **Primary Stateful HA Licensed** - Indicates if the Primary appliance is licensed for Stateful HA. Possible values are **Yes** or **No**. With Stateful HA licensed and enabled, the dynamic state is continuously synchronized between the Active and Standby units. When the Active unit encounters a fault condition, stateful failover occurs as the Standby firewall takes over the Active role with no interruptions to the existing network connections.
- **Secondary Stateful HA Licensed** - Indicates if the Secondary appliance has a Stateful HA license. Possible values are Yes or No. Note that the Stateful HA license is shared with the Primary, but you must access MySonicWall at <https://www.mysonicwall.com> while logged into the unique management IP address of the Secondary unit in order to synchronize with the SonicWall licensing server.

Configuring High Availability

Topics:

- [Configuring Active/Standby High Availability Settings](#)
- [Configuring HA with Dynamic WAN Interfaces](#)


IMPORTANT: High Availability cannot be used along with PortShield except with the SonicWall Switches. Before configuring HA, remove any existing PortShield configuration from **NETWORK | System > PortShield Groups**. For more information, go to <https://www.sonicwall.com/support/technical-documentation/> and search for the SonicWall TZ Series in the Select A Product field.

Configuring Active/Standby High Availability Settings

The configuration tasks on **DEVICE | High Availability > Settings** are performed on the Primary firewall and then are automatically synchronized to the Secondary firewall.

To configure Active/Standby:

1. Navigate to **DEVICE | High Availability > Settings**.
2. In **GENERAL SETTINGS** section, do the following:
 - a. Select **Active / Standby** from the **Mode** drop-down field.
 - b. (Optional) If Licensed, select **Enable Stateful Synchronization**. This option is not selected by default.
When Stateful High Availability is not enabled, session state is not synchronized between the Primary and Secondary firewalls. If a failover occurs, any session that had been active at the time of failover needs to be renegotiated.
 - c. (Optional) Click **OK** in the information dialog displayed.



- d. (Optional) To configure the High Availability Pair so that the Primary firewall takes back the Primary role when it restarts after a failure, select **Enable Preempt Mode**. This option is not selected by default.

① | **TIP:** It is recommended that preempt mode be disabled when enabling Stateful High Availability because preempt mode can be over-aggressive about failing over to the Standby firewall.

- e. (Optional) Click **OK**.


- f. (Optional) Select **Enable Virtual MAC** to allow the Primary and Secondary firewalls to share a single MAC address. This greatly simplifies the process of updating network ARP tables and caches when a failover occurs. This option is not selected by default.

① | **IMPORTANT:** If PPPoE Unnumbered is configured, you must select Enable Virtual MAC. Only the switch to which the two firewalls are connected needs to be notified. All outside devices continue to route to the single shared MAC address.

- g. (Optional) To encrypt HA control communication between the active and standby firewalls, select **Enable Encryption for Control Communication**. This option is not selected by default.

① | **IMPORTANT:** Firewall performance may be affected if you choose encryption.


A confirmation message displays:



- h. (Optional) Click **OK**.

- 3. In the **HA DEVICES** section, enter the Serial Number of the **SECONDARY DEVICE**.

The serial number for the Primary Device is displayed, but the field is dimmed and cannot be edited.



- 4. In the **HA INTERFACES** section:

- a. Select the interface for the **HA Control Interface**.

This option is dimmed and the interface displayed if the firewall detects that the interface is already configured.

- b. (Optional) Select the interface for the **HA Data Interface**.

- c. When finished with all High Availability configuration, click **Accept**. All settings are synchronized to the Secondary firewall, and the Secondary firewall reboots.

Configuring HA with Dynamic WAN Interfaces

The configuration tasks on **DEVICE | High Availability > Settings** are performed on the Active firewall and then are automatically synchronized to the Standby firewall.

To configure HA with a dynamic WAN interface:

1. Navigate to **NETWORK | System > Interfaces**.
2. Configure a WAN interface as PPPoE, as described in *Configuring a WAN Interface* in the *SonicOS 7.0 Firewall Network* document available at <https://www.sonicwall.com/support/technical-documentation/>.
3. Navigate to **DEVICE | High Availability > Settings**.
4. In **GENERAL SETTINGS** section, do the following:
 - a. Select **Active/Passive** from the **Mode** drop-down field.
 - b. Click **OK**.
 - c. Ensure **Enable Stateful Synchronization** is not selected. This option is not selected by default and may require additional licensing.
 - d. Ensure **Enable Preempt Mode** is not selected. This option is not selected by default.
 - e. Select **Enable Virtual MAC** to allow the Primary and Secondary firewalls to share a single MAC address. This greatly simplifies the process of updating network ARP tables and caches when a failover occurs. This option is not selected by default.
 - f. If PPPoE Unnumbered is configured, you must select Enable Virtual MAC.
 - g. Only the switch to which the two firewalls are connected needs to be notified. All outside devices continue to route to the single shared MAC address.
5. Configure HA Devices and HA Interfaces options as described in [Configuring Active/Standby High Availability Settings](#).
6. Click **Accept**.
7. Navigate to **DEVICE | High Availability > Monitoring**.

NAME	PRIMARY IP ADDRESS	SECONDARY IP ADDRESS	PROBE IP ADDRESS	PHYSICAL/LINK MONITORING	LOGICAL/PROBE MONITORING	MANAGEMENT
X0	0.0.0.0	0.0.0.0	0.0.0.0	✓		
X1	0.0.0.0	0.0.0.0	0.0.0.0	✓		
X2	0.0.0.0	0.0.0.0	0.0.0.0			
X3	0.0.0.0	0.0.0.0	0.0.0.0			
X4	0.0.0.0	0.0.0.0	0.0.0.0			

8. Hover over the PPPoE interface and click **Edit** icon. **Interface Monitoring Settings** dialog is displayed.

Interface X3 Monitoring Settings

Physical/Link Monitoring

Primary IPv4 Address

Secondary IPv4 Address

Allow Management on Primary/Secondary IPv4 Address

Logical/Probe IPv4 Address

Override Virtual MAC

Cancel OK

9. Enable **Physical/Link Monitoring**. This option is not selected by default.
10. Ensure the **Primary IPv4 Address** and **Secondary IPv4 Address** fields are set to 0.0.0.0.
11. Ensure none of the other options are selected.
12. Click **OK**.

Configuring High Availability in the Cloud Platform

Topics:

- [Set up an Active/Standby High Availability Configuration Using Azure](#)

Set up an Active/Standby High Availability Configuration Using Azure

SonicWall NSv series brings industry-leading NGFW capabilities, such as application intelligence and control, real-time monitoring, IPS, TLS/SSL decryption and inspection, advanced threat protection, VPN, and Network segmentation capabilities, to protect your Azure environment. The following scenario will show how to deploy a high-availability environment using two Sonicwall NSv in Microsoft Azure's cloud platform.

Azure lets you add cloud capabilities to your existing network through its platform as a service (PaaS) model or entrust Microsoft with all your computing and network needs with Infrastructure as a Service (IaaS).

Product Matrix Table:

Product Models	NSv 270	NSv 470	NSv 870
Maximum Cores	2	4	8
Minimum Total Cores	2	2	2
Management Cores	1	1	1
Maximum Data Plane Cores	1	3	7
Minimum Data Plane Cores	1	1	1

① **NOTE:** HA requires a minimum of three interfaces for High-Availability Exchange Messages. Hence, the VM size should be selected as Standard D3_V2 for GEN 7 template deployment. By default, the SonicWall custom template already sets the value Standard_D3_v2.

For example, the following IP addresses are used in this guide.

GEN7NSvHA-01

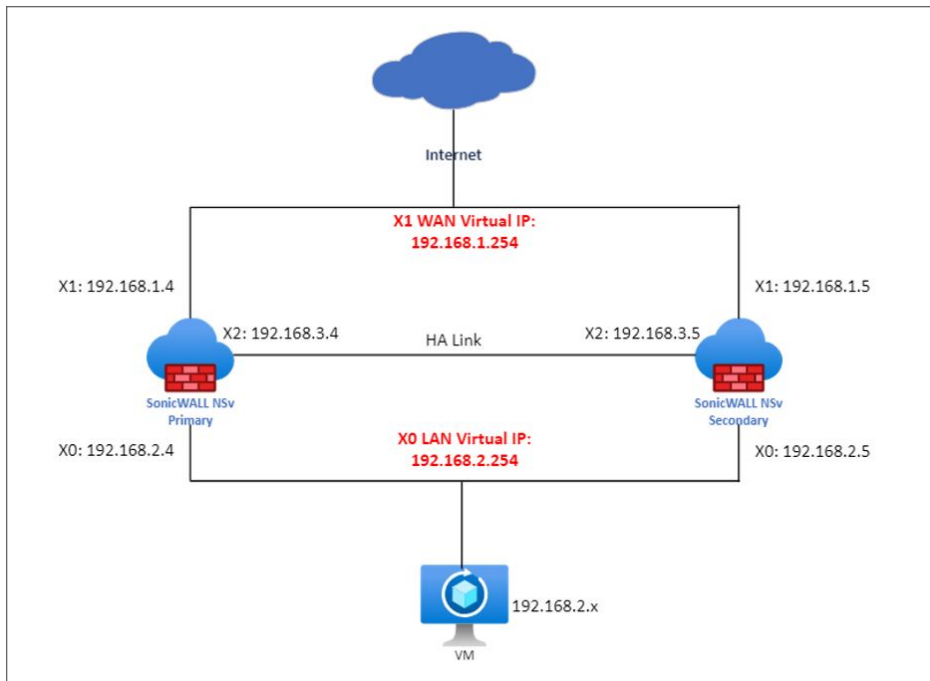
Vnet	192.168.0.0/16
Resource Group	High Availability Standby
WAN IP X1	192.168.1.4/24
LAN IP X0	192.168.2.4/24
HA IP X2	192.168.3.4/24

GEN7NSvHA-02

Vnet	192.168.0.0/16
Resource Group	High Availability Standby
WAN IP X1	192.168.1.5/24
LAN IP X0	192.168.2.5/24
HA IP X2	192.168.3.5/24

① **NOTE:** For the HA interface, use only /24 subnet. There is no such limitation for other interfaces like X0 or X1.

Topology:



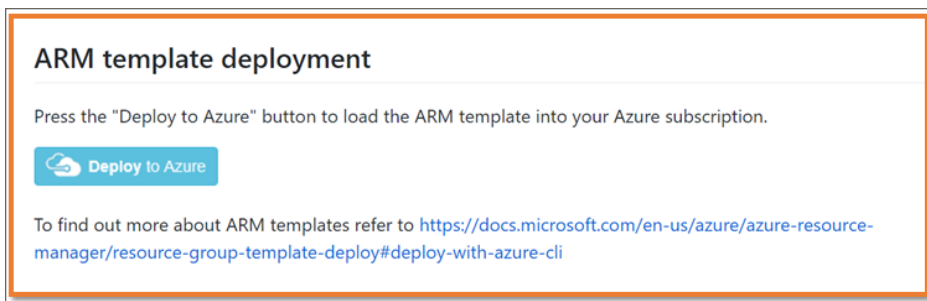
Topics:

- [Install the Custom Template](#)
- [Enable Identity of Both Virtual Machines \(HA1 and HA2\)](#)
- [Role Assignment](#)
- [Check the Networking Tab](#)
- [Configuring Active NSv Firewall Using the Associated Public IP](#)
- [Configuring Standby NSv Firewall Using the Associated Public IP](#)
- [Enable the L3 Mode](#)
- [Configuring Active NSv Firewall Using the Floating Public IP](#)
- [Configuring HA to Active/Standby with L3 HA link](#)
- [Adding Additional Floating Public IP](#)

Install the Custom Template

SonicWall provides a custom template with default values that have already set up according to best practices. To install the custom Template, please access the SonicWall GitHub Repository through the link below, following the steps:

1. In your browser, navigate to <https://github.com/sonicwall-NSv/azure-template/tree/feature/HA>, scroll down, and click **Deploy to Azure**.



2. Log in to the **Azure Portal** using valid credentials.
3. The following custom **Template** will appear. Proceed through the tabs and fill in the blank spaces, such as Resource Group (SonicWall recommends creating a NEW one), Storage Account, etc.

Custom deployment

Deploy from a custom template

New! Deployment Stacks let you manage the lifecycle of your deployments. Try it now →

Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription * ⓘ

Resource group * ⓘ [Create new](#)

Instance details

Region * ⓘ

Location ⓘ

Storage Account * ⓘ ✓

Storage Account Type ⓘ ✓

Storage Account New Or Existing ⓘ

Vm Name Prefix * ⓘ ✓

SSH User Name ⓘ ✓

Authentication Type

SSH Password ⓘ

Ssh Key ⓘ

Image Sku ⓘ ✓

Image Version ⓘ ✓

Management Access IP Source ⓘ ✓

Nsv Model ⓘ

- ① **NOTE:** Please store the "SSH Password" and "Key pair" safely. They will be required when Console or SSH access to the firewall is needed.
- ① **NOTE:** Please leave the Image Version tab as the default value "latest" to install the Gen 7 NSv Firewall.


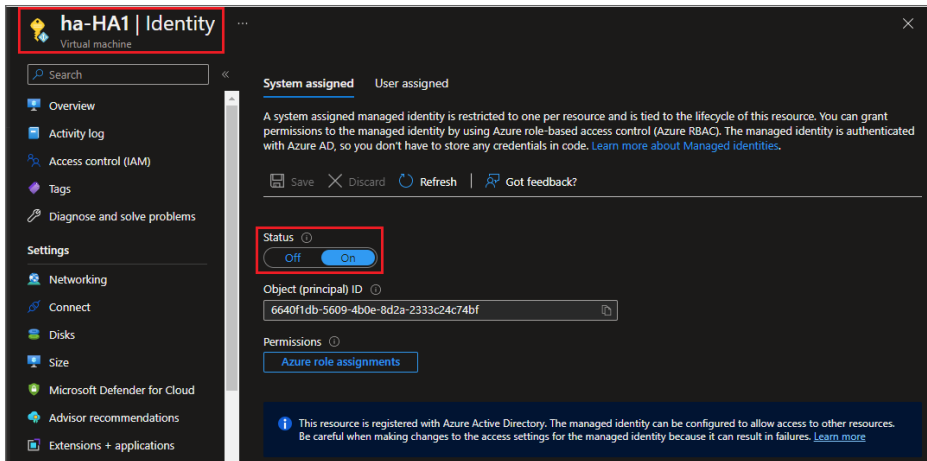
Virtual Network New Or Existing	new	✓
Virtual Network Name	firewallVnet	✓
Virtual Network Address Prefix	192.168.0.0/16	✓
Virtual Network Resource Group Name	[resourceGroup().name]	
Subnet WAN Name	WAN-X1	✓
Subnet LAN Name	LAN-X0	✓
Subnet HA Name	HA-X2	✓
Subnet WAN Prefix	192.168.1.0/24	✓
Subnet LAN Prefix	192.168.0.0/24	✓
Subnet HA Prefix	192.168.2.0/24	✓

The custom template brings up two Virtual Machines (HA1 and HA2) with LAN, WAN, and HA Interfaces, which are all necessary to complete the deployment successfully.

Enable Identity of Both Virtual Machines (HA1 and HA2)

To enable Identity:

1. Access the **All Services > Virtual machine** page on the left panel.
2. Search for the Primary VM you created during deployment. On the left panel, select **Identity** and change the status to **ON** (if it is already configured through the template, please leave it as the default).
3. Repeat the steps 1 and 2 to the **Secondary VM**.

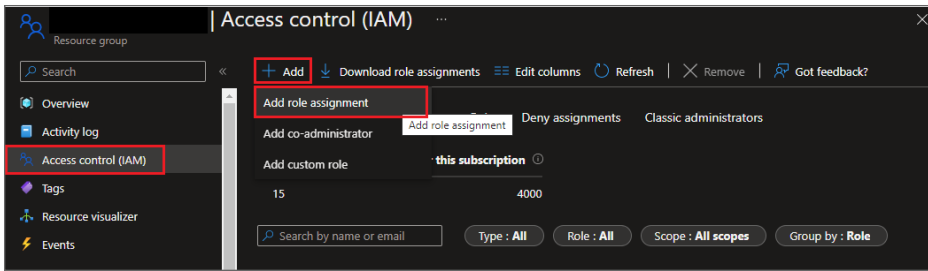


Role Assignment

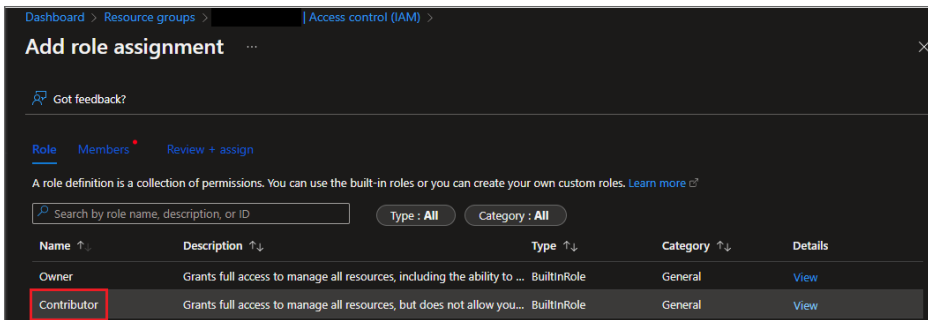
As the next step, you will add Role Assignment to the Resource Group.

The role assignment “Contributor” should be set to the Resource Group since it will allow the firewalls to exchange High Availability information. To do that, follow the steps below:


1. Navigate to the **Dashboard**. On the search bar, search for the Resource Group you created during deployment. On the left panel, select **Access Control (IAM)**.
2. Click **Add > Add role assignment**.



3. Click Contributor.



4. In Contributor, select **Managed Identity**, check if you are in the right subscription, and drop Managed Identity. Select Virtual Machine and select **HA1 and H2 VMs to provide permissions** (if it is already configured through the template, please leave it as the default).




NOTE: It is recommended that VNet and NSv Virtual Machines be in the same Resource Group. If they are in separate Resource Groups, NSv Virtual Machines must also be added as Contributors to the VNet Resource Group.


Check the Networking Tab

For the next step, we are going to check the **Networking** tab.

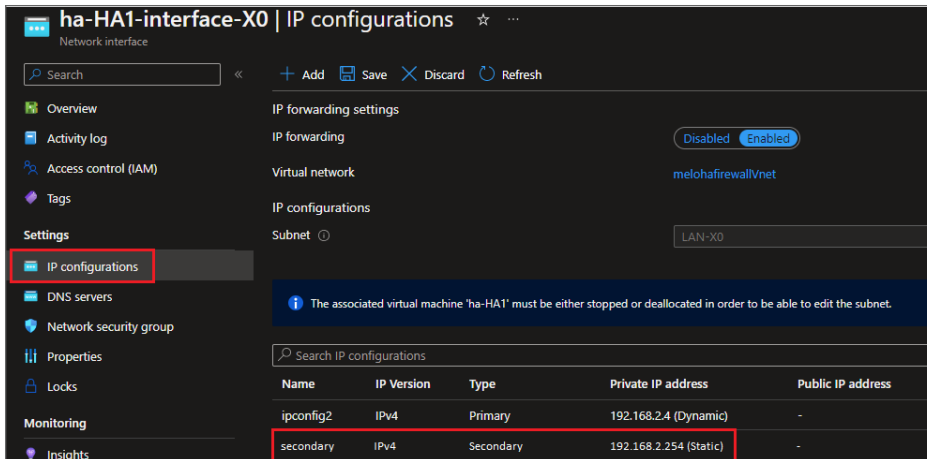
1. Access **All Services** > **Virtual machine** page is on the left panel.
2. Search for the **Primary VM (HA1)** you created during deployment.
3. On the left menu, access the **Networking** tab.
4. Select **ha-HA1-Interface-X1** > **Network Interface ha-HA1-Interface-X1** > **IP Configurations**.



You will note that the template automatically configures an additional Secondary Interface, which acts as a virtual IP address. Virtual IP address should be for both the WAN and LAN Interfaces. Therefore, these IPs will be necessary for the next part of the configuration.



Repeat the previous steps to access **HA1-interface-X0** and check the **Secondary Interface**.



Now you've check that the **Secondary Interface** IP address is provided on both X1 and X0 interfaces, you will set those IPs on the Active NSv Firewall as shown in the next steps.

ⓘ | **NOTE:** For the first login, it's going to be necessary to Register the appliance.

Configuring Active NSv Firewall Using the Associated Public IP

Once you are logged in, **register the appliance first**. After the registration, follow the steps below.

1. Navigate to **Network > System > Interfaces**. Change the X0 configuration first, and then X1, as shown below. You will lose access after you change X1.

Edit Interface - X0

General | Advanced

INTERFACE 'X0' SETTINGS

Zone: LAN

Mode / IP Assignment: Static IP Mode

IP Address: 192.168.2.254

Subnet Mask: 255.255.255.0

Default Gateway (Optional): 0.0.0.0

Comment: HA1-interface-X0-ipconfig-float

Domain Name: ⓘ

Add rule to enable redirect from HTTP to HTTPS:

MANAGEMENT

- HTTPS:
- Ping:
- SNMP:
- SSH:

USER LOGIN

- HTTP:
- HTTPS:

Cancel OK

Edit Interface - X1

General | Advanced

INTERFACE 'X1' SETTINGS

Zone: WAN

Mode / IP Assignment: Static

IP Address: 192.168.1.254

Subnet Mask: 255.255.255.0

Default Gateway: 192.168.1.1

DNS Server 1: 168.63.129.16

DNS Server 2: 0.0.0.0

DNS Server 3: 0.0.0.0

Comment: HA1-interface-X1-ipconfig-float

Domain Name:

Add rule to enable redirect from HTTP to HTTPS:

MANAGEMENT

HTTPS:

Ping:

SNMP:

SSH:

USER LOGIN

HTTP:

HTTPS:

Cancel OK

2. Before changing the X1 Interface IP, take a **screenshot** of the current configuration. The Default Gateway and DNS Server 1 should remain the same and ensure you won't lose access to the active firewall during the following steps.
3. After the changes, the Interfaces should look similar to the screenshot below:

System | Interface Settings | Traffic Statistics

IPv4 | IPv6

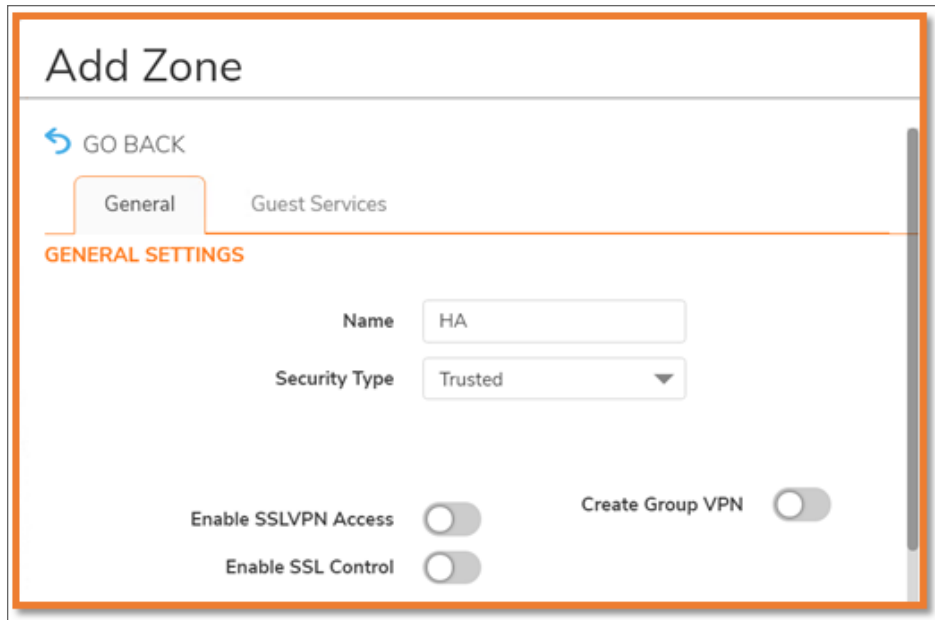
+ Add Interface

NAME	ZONE	GROUP	IP ADDRESS	SUBNET MASK	IP ASSIGNMENT	STATUS	ENABLED	COMMENT
X0	LAN	N/A	192.168.2.254	255.255.255.0	Static IP	10 Gbps Full Duplex	<input checked="" type="checkbox"/>	HA1-interface-X0-ip-float
X1	WAN	Default LB Group	192.168.1.254	255.255.255.0	Static IP	10 Gbps Full Duplex	<input checked="" type="checkbox"/>	HA1-interface-X1-ip-float

Configuring Standby NSv Firewall Using the Associated Public IP

You will access the standby firewall for the first time, so you should also register the appliance before proceeding with the rest of the configuration.

1. Navigate to **Network > System > Interfaces**, and create the HA interface with the custom zone "HA" using the below IP address schema details (SonicWall uses X2 Interface in the customer template).
 - a. Security Type: **Trusted**



Add Zone

[GO BACK](#)

General Guest Services

GENERAL SETTINGS

Name: HA

Security Type: Trusted

Enable SSLVPN Access:

Enable SSL Control:

Create Group VPN:

2. The IP Address is the **ha-HA2-Interface-X2** found on the **Networking** tab.

Edit Interface - X2

General | Advanced

INTERFACE 'X3' SETTINGS

Zone: HA

Mode / IP Assignment: Static IP Mode

IP Address: 192.168.3.5

Subnet Mask: 255.255.255.0

Default Gateway (Optional): 0.0.0.0

Comment: ha-HA2-interface-X2

Domain Name:

Add rule to enable redirect from HTTP to HTTPS:

MANAGEMENT

HTTPS:

Ping:

SNMP:

SSH:

USER LOGIN

HTTP:

HTTPS:

Cancel OK

System | Interface Settings | Traffic Statistics

IPv4 | IPv6

+ Add Interface Refresh

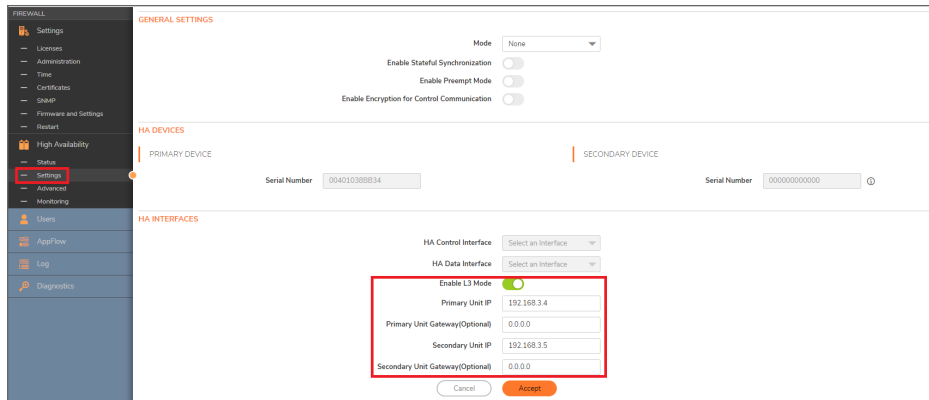
NAME	ZONE	GROUP	IP ADDRESS	SUBNET MASK	IP ASSIGNMENT	STATUS	ENABLED	COMMENT
X0	LAN	N/A	192.168.2.254	255.255.255.0	Static IP	10 Gbps Full Duplex	<input checked="" type="checkbox"/>	HA1-interface-X0-standby- Role
X1	WAN	Default LB Group	192.168.1.254	255.255.255.0	Static IP	10 Gbps Full Duplex	<input checked="" type="checkbox"/>	HA1-interface-X1-standby- Role
X2	HA-Link	N/A	192.168.3.5	255.255.255.0	Static IP	10 Gbps Full Duplex	<input checked="" type="checkbox"/>	HA-Data-Link
X3	Unassigned	N/A	0.0.0.0	0.0.0.0		Unplugged	<input checked="" type="checkbox"/>	N/A
X4	Unassigned	N/A	0.0.0.0	0.0.0.0		Unplugged	<input checked="" type="checkbox"/>	N/A
X5	Unassigned	N/A	0.0.0.0	0.0.0.0		Unplugged	<input checked="" type="checkbox"/>	N/A
X6	Unassigned	N/A	0.0.0.0	0.0.0.0		Unplugged	<input checked="" type="checkbox"/>	N/A
X7	Unassigned	N/A	0.0.0.0	0.0.0.0		Unplugged	<input checked="" type="checkbox"/>	N/A

Enable the L3 Mode

Enable L3 mode option on the Standby appliance.

1. Navigate to **Device > High Availability > Settings > HA Interfaces**, and select the **Enable L3 Mode** option on the standby firewall.
2. Under HA Interfaces, fill out the primary Unit IP and Secondary Unit IP, providing the **HA IP X2**

addresses of both firewalls HA1 and HA2.



Configuring Active NSv Firewall Using the Floating Public IP

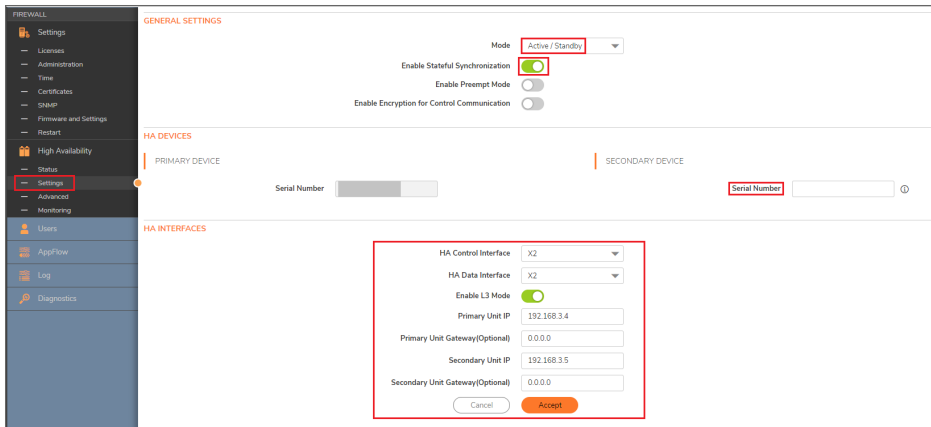
On the Active firewall, Log in using the Floating Public IP address. To get the IP:

1. Navigate to **Resource Groups**.
2. Select the **VM ha-HA1**.
3. On the left panel, press **Networking**.
4. Network Interface: ha-HA1-Interface-X1.
5. On the left panel, press **IP configurations**.
6. Copy and paste the **Floating Public IP** address into the browser.

Configuring HA to Active/Standby with L3 HA link

The following steps configure HA to Active/Standby with the L3 HA link.

1. Configure HA to Active/Standby with the L3 HA link. To do so, browse **Manage > High Availability** and select **Enable Stateful Synchronization**.
2. Navigate to the HA interfaces section and switch the HA Control link to L3 mode. A gateway address is unnecessary if two HA Interfaces are in the same subnet. However, if two HA interfaces are in different subnets, a proper gateway address is needed, and the default is X.X.X.1 on Azure.



3. Add monitoring IPs for X0 and X1 is mandatory.

NAME	PRIMARY IP ADDRESS	SECONDARY IP ADDRESS	PROBE IP ADDRESS	PHYSICAL/LINK MONITORING	LOGICAL/PROBE MONITORING	MANAGEMENT
X0	192.168.2.4	192.168.2.5	0.0.0.0	✓	✓	✓
X1	192.168.1.4	192.168.1.5	0.0.0.0	✓	✓	✓
X3	0.0.0.0	0.0.0.0	0.0.0.0			
X4	0.0.0.0	0.0.0.0	0.0.0.0			
X5	0.0.0.0	0.0.0.0	0.0.0.0			
X6	0.0.0.0	0.0.0.0	0.0.0.0			
X7	0.0.0.0	0.0.0.0	0.0.0.0			


4. Navigate to **Device > High Availability > Status** page to check whether the pair is coming together. The secondary will reboot, and seeing the device pair up may take a while.

NOTE: It is recommended that you use the latest 7.1.1-7051 or a newer firmware version to use multiple floating IP addresses in Azure.

Adding Additional Floating Public IP

Steps to add Additional Floating public IP:

1. Navigate to **Resource Groups**.
2. Select the **VM ha-HA1**.
3. On the left panel, press **Networking**.
4. Network Interface: ha-HA1-Interface-X1.
5. On the left panel, press **IP configurations**.
6. Click **Add** to Add IP configuration and specify the name and IP allocation method.



7. **Associate the public IP** by using an existing public IP or create a new public IP address.

Add IP configuration ✕

nsw470-HA1-interface-X1

i A primary IP configuration already exists. Any additional IP configurations will be secondary. The virtual network this network interface is attached to only supports IPv4. [Learn more](#) 🔗

Name

IP version

IPv4

IPv6

Type

Primary

Secondary

Private IP address settings

Allocation

Dynamic

Static

Private IP address

Public IP address settings

Associate public IP address

Public IP address

[Create a public IP address](#)

Add
Cancel

Fine Tuning High Availability

Topics:

- [Advanced Settings](#)
- [Configuring Advanced High Availability Settings](#)

Advanced Settings

DEVICE | High Availability > Advanced provides the ability to fine-tune the High Availability configuration as well as synchronize settings and firmware between the High Availability Security Appliances

The **Heartbeat Interval** and **Failover Trigger Level (missed heartbeats)** settings apply to the HA heartbeats.

For more information on High Availability, see [About High Availability](#) and [Active/Standby Prerequisites](#).

Configuring Advanced High Availability Settings

To configure advanced settings:

1. Log in as an administrator to the SonicOS Management Interface on the Active Node.
2. Navigate to **DEVICE | High Availability > Settings**.

ADVANCED SETTINGS

Heartbeat Interval (milliseconds)	<input type="text" value="1000"/>	?
Failover Trigger Level (missed heartbeats)	<input type="text" value="5"/>	?
Probe Interval (seconds)	<input type="text" value="20"/>	?
Probe Count	<input type="text" value="3"/>	?
Election Delay Time (seconds)	<input type="text" value="3"/>	?
Dynamic Route Hold-Down Time (seconds)	<input type="text" value="45"/>	?
SD-WAN Probes Hold-Down Time (seconds)	<input type="text" value="10"/>	
Active/Standby Failover only when ALL aggregate links are down	<input type="checkbox"/>	
Include Certificates/Keys	<input checked="" type="checkbox"/>	?

- Set the **Probe Interval** to the interval, in seconds, between probes sent to specified IP addresses to monitor that the network critical path is still reachable. This interval is used in logical monitoring for the local HA pair. The default is **20** seconds, and the allowed range is 5 to 255 seconds.
You can set the Probe IP Address(es) on **DEVICE | High Availability > Advanced**. See [Monitoring High Availability](#).
- Set the **Probe Count** to the number of consecutive probes before SonicOS concludes that the network critical path is unavailable or the probe target is unreachable. This count is used in logical monitoring for the HA pair. The default is **3**, and the allowed range is 3 to 10.
- Set the **Election Delay Time** to the number of seconds the Active Security Appliance waits to consider an interface up and stable. The default is 3 seconds, the minimum is **3** seconds, and the maximum is 255 seconds.
This timer is useful with switch ports that have a spanning-tree delay set.
- Set the **Dynamic Route Hold-Down Time** to the number of seconds the newly-active Security Appliance keeps the dynamic routes it had previously learned in its route table. The default value is **45** seconds, the minimum is 0 seconds, and the maximum is 1200 seconds (20 minutes).
 - NOTE:** The **Dynamic Route Hold-Down Time** setting is displayed only when the **Advanced Routing Mode** option is selected on **NETWORK | System > Dynamic Routing > Settings**.
 - TIP:** In large or complex networks, a larger value may improve network stability during a failover. This setting is used when a failover occurs on a High Availability pair that is using either a dynamic routing protocol. During this time, the newly-active appliance relearns the dynamic routes in the network. When the **Dynamic Route Hold-Down Time** duration expires, SonicOS deletes the old routes and implements the new routes it has learned from routing protocols.
- If you want Failover to occur only when ALL aggregate links are down, select **Active/Standby Failover only when ALL aggregate links are down**. This option is not selected by default.
- To have the appliances synchronize all certificates and keys within the HA pair. select **Include Certificates/Keys**. This option is selected by default.
- (Optional) To force synchronize the SonicOS preference settings between your primary and secondary HA firewalls, click **Synchronize Settings**.

① | **NOTE:** This will cause a restart of the standby device.

10. (Optional) To synchronize the firmware version between your primary and secondary HA firewalls, click **Synchronize Firmware**.
11. (Optional) To test the HA failover functionality is working properly by attempting an Active/Standby HA failover to the secondary Security Appliance, click **Force Active/Standby Failover**.
12. When finished with all High Availability configuration, click **Accept**. All settings are synchronized to the Secondary Security Appliance.

Monitoring High Availability

On **DEVICE | High Availability > Monitoring**, you can configure independent management IP addresses for each unit in the HA Pair. You can also configure physical/link monitoring and logical/probe monitoring.

Topics:

- [Configuring Active/Standby High Availability Monitoring](#)
- [IPv6 High Availability Monitoring](#)
- [IPv6 HA Monitoring Considerations](#)

Configuring Active/Standby High Availability Monitoring

To set the independent LAN management IP addresses and configure physical and/or logical interface monitoring:

1. Log in as an administrator to the SonicOS Management Interface on the Active SonicWall Security Appliance.
2. Navigate to **DEVICE | High Availability > Monitoring**.

Monitoring Ipv4 Settings		Monitoring IPv6 Settings				
NAME	PRIMARY IP ADDRESS	SECONDARY IP ADDRESS	PROBE IP ADDRESS	PHYSICAL/LINK MONITORING	LOGICAL/PROBE MONITORING	MANAGEMENT
X0	0.0.0.0	0.0.0.0	0.0.0.0	✓		
X1	0.0.0.0	0.0.0.0	0.0.0.0	✓		
X2	0.0.0.0	0.0.0.0	0.0.0.0			
X3	0.0.0.0	0.0.0.0	0.0.0.0			
X4	0.0.0.0	0.0.0.0	0.0.0.0			
X5	0.0.0.0	0.0.0.0	0.0.0.0			
X6	0.0.0.0	0.0.0.0	0.0.0.0			
X7	0.0.0.0	0.0.0.0	0.0.0.0			
X8	0.0.0.0	0.0.0.0	0.0.0.0			
X9	0.0.0.0	0.0.0.0	0.0.0.0			
U0	0.0.0.0	0.0.0.0	0.0.0.0			

3. Click the **Edit** icon for an interface on the LAN, such as X0. The **Interface Monitoring Settings** dialog is displayed.
4. To enable link detection between the designated HA interfaces on the Primary and Secondary units, leave **Physical/Link Monitoring** selected. This option is selected by default.
5. In the Primary IPv4/v6 Address field, enter the unique LAN management IP address of the Primary unit and it should be same subnet as interface IP address. The default is **0.0.0.0**.
6. In the Secondary IPv4/v6 Address field, enter the unique LAN management IP address of the Secondary unit it should be same subnet as interface IP address. The default is **0.0.0.0**.
7. Select **Allow Management on Primary/Secondary IP Address**. When this option is enabled for an interface, a green icon appears in the interface's **Management** column in the **Monitoring Settings** table. Management is only allowed on an interface when this option is enabled. This option is not selected by default.
8. In the **Logical/ Probe IPv4/v6 Address** field, enter the IP address of a downstream device on the network that should be monitored for connectivity. Typically, this should be a downstream router or server. (If probing is desired on the WAN side, an upstream device should be used.) This option is not selected by default.

The Primary and Secondary Security Appliances regularly ping this probe IP address. If both successfully ping the target, no failover occurs. If neither successfully ping the target, no failover occurs, because it is assumed that the problem is with the target, and not the Security Appliances. But, if one Security Appliance can ping the target but the other cannot, failover occurs to the Security Appliance that can ping the target.

The Primary IPv4/v6 Address and Secondary IPv4/v6 Address fields must be configured with independent IP addresses on a the interface, such as X0, (or a WAN interface, such as X1, for probing on the WAN) to allow logical probing to function correctly.

9. (Optional) To manually specify the virtual MAC address for the interface, select **Override Virtual MAC** and enter the MAC address in the field. The format for the MAC address is six pairs of hexadecimal numbers separated by colons, such as A1:B2:C3:d4:e5:f6. This option is not selected by default.

ⓘ | IMPORTANT: Care must be taken when choosing the Virtual MAC address to prevent configuration errors.

When **Enable Virtual MAC** is selected on **DEVICE | High Availability > Settings**, the SonicOS firmware automatically generates a Virtual MAC address for all interfaces. Allowing the SonicOS firmware to generate the Virtual MAC address eliminates the possibility of configuration errors and ensures the uniqueness of the Virtual MAC address, which prevents possible conflicts.

10. Click **OK**.
11. Click **Close**.

IPv6 High Availability Monitoring

For complete information on the SonicOS implementation of IPv6, see IPv6.

IPv6 High Availability (HA) Monitoring is implemented as an extension of HA Monitoring in IPv4. After configuring HA Monitoring for IPv6, both the primary and backup Security Appliances can be managed from the IPv6 monitoring address, and IPv6 Probing is capable of detecting the network status of HA pairs.

For easy configuration of both IP versions, toggle between IPv6 and IPv4 displays in **DEVICE | High Availability > Monitoring**.

The IPv6 HA Monitoring configuration page is inherited from IPv4, so the configuration procedures are almost identical. Just select IPv6 and refer to [About High Availability](#) and [IPv6 HA Monitoring Considerations](#) for configuration details.

IPv6 HA Monitoring Considerations

Consider the following when configuring IPv6 HA Monitoring:

- In the **Interface Settings** dialog, enable **Physical/Link Monitoring** and **Override Virtual MAC** are dimmed because they are layer 2 properties. That is, the properties are used by both IPv4 and IPv6, so you configure them in the IPv4 monitoring page.
- The Primary/Secondary IPv6 address must be in the same subnet of the interface, and it can not be same as the global IP and Link-Local-IP of the Primary/Secondary Security Appliance.
- If the Primary/Secondary monitoring IP is set to (not ::), then they cannot be the same.
- If **Allow Management on Primary/Secondary IPv6 Address** is enabled, then Primary/Secondary monitoring IPv6 addresses cannot be unspecified (that is, ::).
- If **Logical/Probe IPv6 Address** is enabled, then the probe IP cannot be unspecified.

Azure Use Cases

Topics:

- [Use Case 1: Manage Azure HA Firewall](#)
- [Use Case 2: Forward LAN traffic to the External Network through the Gateway after HA Failover](#)
- [Use Case 3: Configure DNAT on Azure HA Firewall](#)
- [Use Case 4: Configure DNAT on Azure HA Firewall \(Need to move multiple floating IPs support\)](#)

Use Case 1: Manage Azure HA Firewall

Use floating IP (aka interface IP on SonicOS) to manage Azure firewall, which state is active.

Use two HA monitor IPs (aka Primary IP on Azure, aka monitor IP on SonicOS) to manage the Azure firewall. The primary/secondary IP manages firewalls in active and standby states.

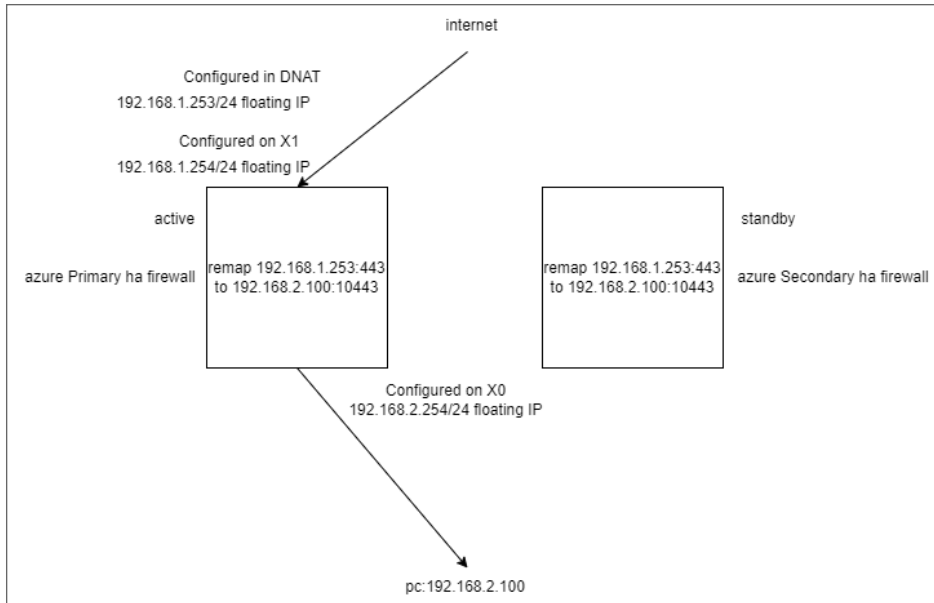
Use Case 2: Forward LAN traffic to the External Network through the Gateway after HA Failover

The interface IP (aka floating IP) in the Azure HA environment is set as the default gateway of the LAN subnet. When LAN traffic forwards to the external network through the gateway, traffic will automatically flow into the active HA Azure firewall without flowing into the standby HA Azure firewall. This is because the interface IP always floats to the active HA Azure firewall. After HA failover, the active HA Azure firewall will change to the standby firewall, so traffic will be forwarded to the newly active firewall.

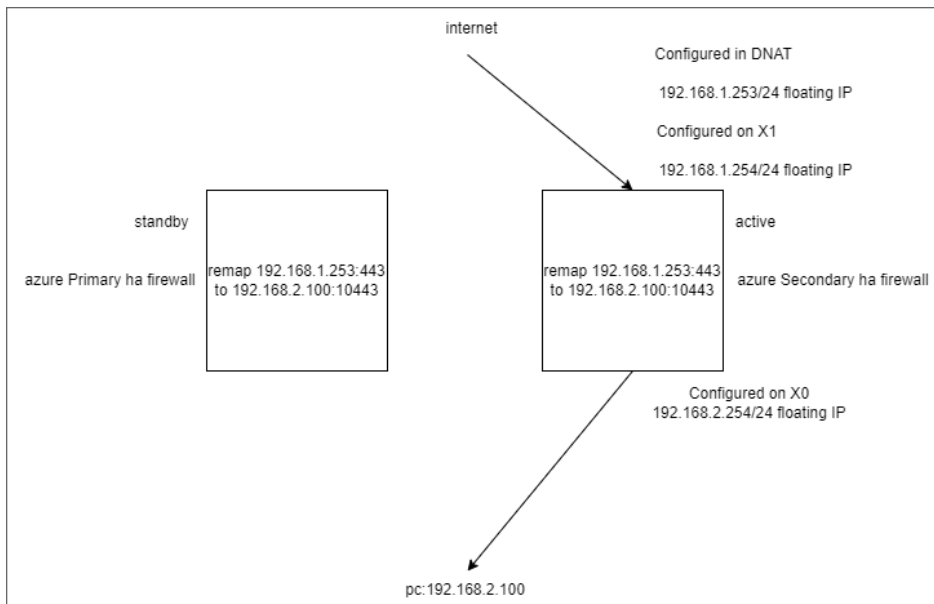
Use Case 3: Configure DNAT on Azure HA Firewall

Configure a DNAT to remap the traffic of an interface IP (192.168.1.254, aka floating IP) accessing the firewall from the Internet to the traffic of a particular machine (192.168.2.200) in the LAN.

FLOATING IP ALWAYS ON ACTIVE FIREWALL




AFTER HA FAILOVER



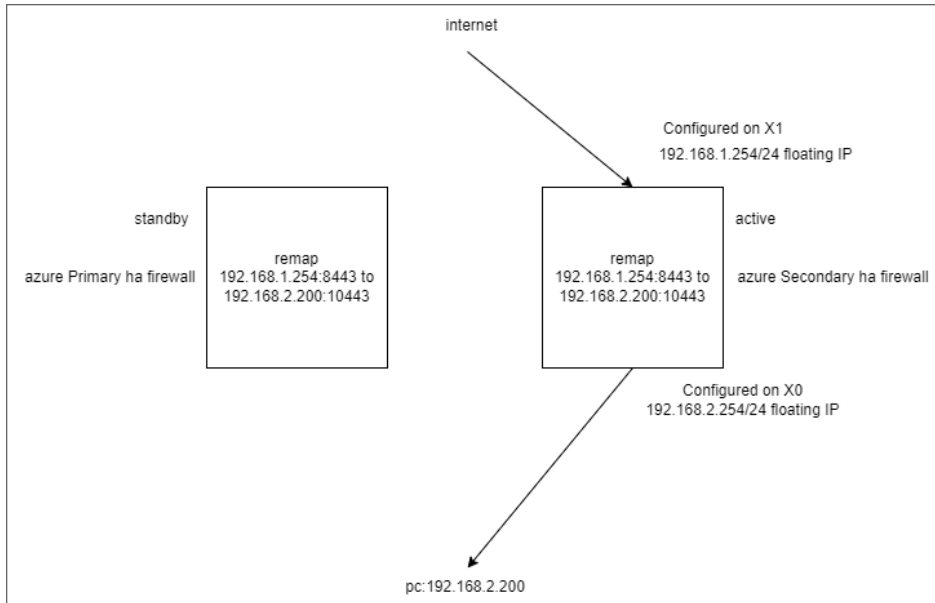
Use Case 4: Configure DNAT on Azure HA Firewall (Need to move multiple floating IPs support)

Configure a DNAT to remap the traffic of floating IP (192.168.1.253, not interface IP) accessing the firewall from the Internet to the traffic of a particular machine (192.168.2.100) in the LAN.

FLOATING IP ALWAYS ON ACTIVE FIREWALL



AFTER HA FAILOVER



SonicWall Support

Technical support is available to customers who have purchased SonicWall products with a valid maintenance contract.

The [Support Portal](#) provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year.

The [Support Portal](#) enables you to:

- View [Knowledge Base articles](#) and [Technical Documentation](#)
- View and participate in the [Community Forum](#) discussions
- View [Video Tutorials](#)
- Access [MySonicWall](#)
- Learn about [SonicWall Professional Services](#)
- Review [SonicWall Support services and warranty information](#)
- Register at [SonicWall University](#) for training and certification

About This Document

SonicOS High Availability Administration Guide
Updated - August 2024
Software Version - 7.0
232-005335-10 Rev B

Copyright © 2024 SonicWall Inc. All rights reserved.

The information in this document is provided in connection with SonicWall and/or its affiliates' products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, SONICWALL AND/OR ITS AFFILIATES ASSUME NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL SONICWALL AND/OR ITS AFFILIATES BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF SONICWALL AND/OR ITS AFFILIATES HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. SonicWall and/or its affiliates make no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. and/or its affiliates do not make any commitment to update the information contained in this document.

For more information, visit <https://www.sonicwall.com/legal>.

End User Product Agreement

To view the SonicWall End User Product Agreement, go to: <https://www.sonicwall.com/legal/end-user-product-agreements/>.

Open Source Code

SonicWall Inc. is able to provide a machine-readable copy of open source code with restrictive licenses such as GPL, LGPL, AGPL when applicable per license requirements. To obtain a complete machine-readable copy, send your written requests, along with certified check or money order in the amount of USD 25.00 payable to "SonicWall Inc.", to:

General Public License Source Code Request
Attn: Jennifer Anderson
1033 McCarthy Blvd
Milpitas, CA 95035