SONICWALL®

SonicOS 8

High Availability

Administration Guide

Contents

About SonicOS **4**
Working with SonicOS 4
SonicOS Workflow 5
How to Use the SonicOS Administration Guides 6
Guide Conventions 8

High Availability **9**
About High Availability 9
- High Availability Terminology 10
- High Availability Modes 11
- High Availability Encryption 12
- Crash Detection 12
- Virtual MAC Address 13
- Dynamic WAN Interfaces with PPPoE HA 13
- Stateful Synchronization with DHCP 13
- Stateful Synchronization with DNS Proxy 14
- About HA Monitoring 14

About Active/Standby HA 15
- Benefits of Active/Standby HA 15
- Working of Active/Standby HA 16

About Stateful Synchronization 16
- Benefits of Stateful Synchronization 16
- How Does Stateful Synchronization Work? 17
- Example of Stateful Synchronization 18

About Active/Active DPI 18
- Benefits of Active/Active DPI 19

Active/Standby and Active/Active DPI Prerequisites 20
- Supported Platforms and Licensing for HA 20
- Physically Connecting Your Firewalls 21
- Connecting the Active/Active DPI Interface 22

Maintenance 22
- Removing an HA Association 22
- Replacing a SonicWall firewall 23

High Availability Status **25**
Active/Standby High Availability Status 25
- High Availability Status 26
- High Availability Config 28

High Availability Licenses ... 28

Configuring High Availability ... 29
Configuring Active/Standby High Availability Settings ... 29
Configuring HA with Dynamic WAN Interfaces ... 31
Configuring Active/Active DPI High Availability Settings ... 32

Configuring High Availability in the Cloud Platform ... 34
Set up an Active/Standby High Availability Configuration Using Azure ... 34
Install the Custom Template ... 36
Enable Managed Identity on Both Virtual Machines (HA1 and HA2) ... 38
Role Assignment ... 39
Check the Networking Tab ... 40
Configuring Active NSv Firewall Using the Associated Public IP ... 41
Configuring Standby NSv Firewall Using the Associated Public IP ... 43
Enable the L3 Mode ... 45
Deploying NSv XS in Azure Using L3 ... 46
Configuring Active NSv Firewall Using the Floating Public IP ... 50
Configuring HA to Active/Standby with L3 HA link ... 50
Adding Additional Floating Public IP ... 51

Fine Tuning High Availability ... 54
Advanced Settings ... 54
Configuring Advanced High Availability Settings ... 54

Monitoring High Availability ... 58
Configuring Active/Standby High Availability Monitoring ... 58
IPv6 High Availability Monitoring ... 60
IPv6 HA Monitoring Considerations ... 60
Verifying Active/Active DPI Traffic Offloading ... 60

Azure Use Cases ... 63
Use Case 1: Manage Azure HA Firewall ... 63
Use Case 2: Forward LAN traffic to the External Network through the Gateway after HA Failover ... 63
Use Case 3: Configure DNAT on Azure HA Firewall ... 64
Use Case 4: Configure DNAT on Azure HA Firewall ... 65

SonicWall Support ... 67
About This Document ... 68

1

About SonicOS

This guide is a part of the SonicOS collection of administrative guides that describe how to administer and monitor the SonicWall family of firewalls. SonicOS provides network administrators with the management interface, API (Application Program Interface), and Command Line Interface (CLI) for firewall configuration. You can configure and manage your firewall by setting objects to secure and protect the network services, manage traffic, and provide the desired level of network service.

Topics:

- Working with SonicOS
- SonicOS Workflow
- How to Use the SonicOS Administration Guides
- Guide Conventions

Working with SonicOS

SonicOS provides a web management interface for configuring, managing, and monitoring the features, policies, security services, connected devices, and outside threats to your network. SonicOS functions in conjunction with SonicCore, SonicWall's secure underlying operating system.

The SonicOS management interface facilitates:

- Setting up and configuring your firewall
- Configuring external devices such as access points or switches
- Configuring networks and external system options that connect to your firewall
- Defining objects and policies for protection
- Monitoring the health and status of the security appliance, network, users, and connections
- Monitoring traffic, users, and threats
- Investigating events

The following table details the SonicWall firewall series and the specific performance, security, and feature sets they provide within the enterprise.

Firewall Type	Classic Mode	Comments
TZ Series	yes	The entry level TZ Series, also known as desktop firewalls, delivers revamped features such as 5G readiness, better connectivity options, improved threat protection, SSL and decryption performance that addresses HTTPS bandwidth issues; built-in SD-WAN, and lawful TLS 1.3 decryption support. It provides advanced networking and security features, like the multi-engine Capture Advanced Threat Protection (ATP) cloud-based sandbox service with patent-pending Real-Time Deep Memory Inspection (RTDMI™).
NSa Series	yes	NSa firewalls provide your mid-sized network with enhanced security. These firewalls are designed specifically for businesses with more than 250 workstations. The NSa provides cloud-based and on-box capabilities such as TLS/SSL decryption and inspection, application intelligence and control, SD-WAN, real-time visualization, and WLAN management.
NSv Series	yes	The NSv series firewalls offers all the security advantages of a physical firewall with the operational and economic benefits of virtualization. The NSv firewalls can operate in either Policy Mode or Classic Mode. You can switch between modes, but some configuration information from extra interfaces is removed.

In addition to the management interface, SonicOS also has a full-featured API and a command-line interface (CLI) to manage the firewalls.

For more information, refer to:

- SonicOS API Reference Guide
- SonicOS CLI Reference Guide

SonicOS Workflow

When working with SonicWall products, you can use the following workflow as a guide for setting up your security solution.

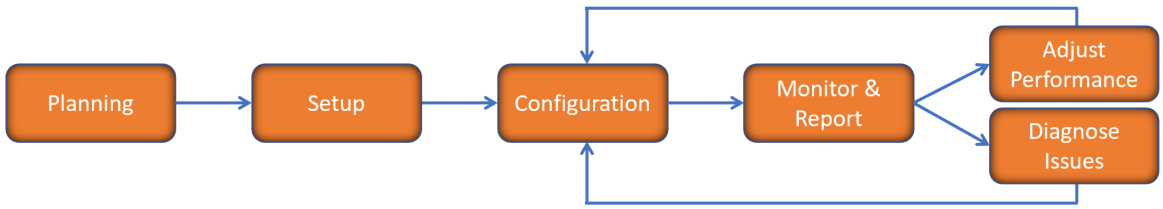

You begin your planning as you start making your purchasing decisions. Your sales partners can help you assess your network and make recommendations based on the kinds of security services you need. You can learn more about SonicWall products by reviewing product information and solutions. After selecting the solution, you can schedule your implementation.

After planning and scheduling your solution, you begin setting up the firewalls. The Getting Started Guides for your products can help you begin setting up the pieces to your solution. The getting started guides are designed to help you install the firewall to a minimal level of operation. Before performing any detailed configuration tasks described in the SonicOS Administration Guides, you should have your firewall set up and basic operation validated.

The configuration block of the workflow refers to the many tasks that combine to define how your firewall is integrated into your security solution and how it behaves when protecting your environment. Depending on the features of your security solution, this task can be quite complex. The System Administration Guides are broken into the key command sets and features. Some documents may be used for all solutions, but others may be used only if you integrated that feature into your solution. For example, High Availability or Wireless Access Points are not necessarily used by all customers. More information about a feature's workflow is presented in the feature administration guide. For more information, refer to the specific Administration Guide for a SonicOS feature.

Configuration tends to be a one-time activity, although you might make minor adjustments after monitoring performance or after diagnosing an issue. The configuration activity can be broken down into the more detailed flow as the following figure shows. This also mirrors the key functions that are listed across the top of the management interface.

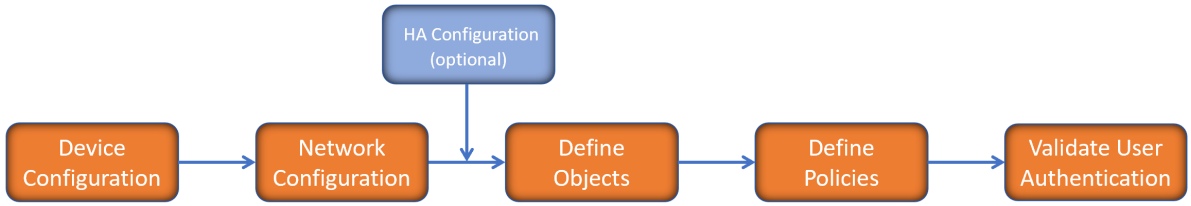

There is some flexibility in the order in which you do things, but this is the general work-flow you would follow when configuring your firewall. Start by defining the settings on the firewall. Next you set up the system and other devices that your firewall is connected to, and you can choose to implement High Availability when done. After your device, network, and system is configured, you should define the objects that you want to monitor. Then you use those objects to define the policies that protect your network. The final step to preparing your setup is to validate the user authentication.

How to Use the SonicOS Administration Guides

The *SonicOS Administration Guide* is a collection of guides that detail the features represented by each of the main menu items in the management interface. Within each guide, you can find topics covering commands in that menu group, along with procedures and in-depth information. The exceptions are the SonicOS 8 Monitor Guide and the SonicOS 8 Objects Guide which combine the topics for each of those functions into a single book.

To help you understand how the books align with the features and commands, the following figure shows the books organized like the SonicOS management interface.

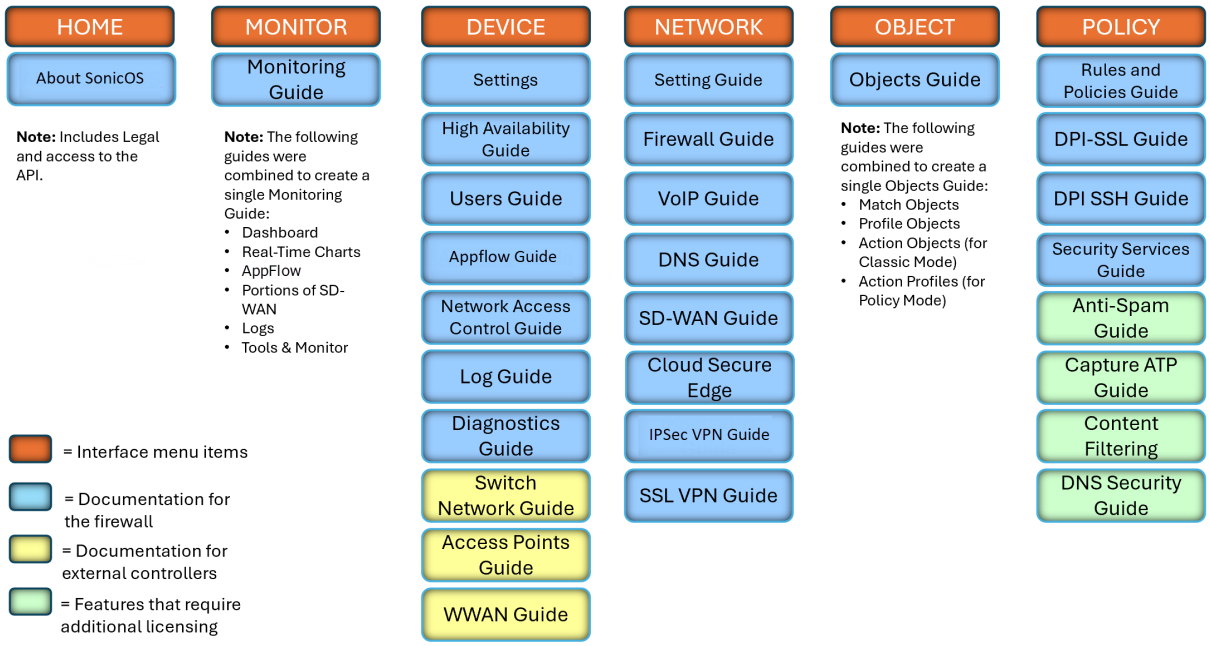

The SonicOS Administration Guides, along with related documentation, such as the getting started guides, are available on the Technical Documentation Portal.

Guide Conventions

These text conventions are used in this guide:

NOTE: A NOTE icon indicates supporting information.

IMPORTANT: An IMPORTANT icon indicates information that is essential for successful completion of a task, proper operation of the product, or avoidance of configuration errors.

TIP: A TIP icon indicates helpful information.

CAUTION: A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

WARNING: A WARNING icon indicates a potential for property damage, personal injury, or death.

Convention	Description
Bold text	Used in procedures to identify elements in the management interface like dialog boxes, windows, screen names, messages, and buttons. Also used for file names and text or values you are being instructed to select or type into the interface.
Function | Menu group > Menu item	Indicates a multiple step menu choice on the user interface. For example, **NETWORK | System > Interfaces** means to select the **NETWORK** functions at the top of the window, then click **System** in the left navigation menu to open the menu group (if needed) and select **Interfaces** to display the page.
`Code`	Indicates sample computer programming code. If bold, it represents text to be typed in the command line interface.
<Variable>	Represents a variable name. The variable name and angle brackets need to be replaced with an actual value. For example in the segment **serialnumber=***<your serial number>*, replace the variable and brackets with the serial number from your device, such as serialnumber=2CB8ED000004.
Italics	Indicates the name of a technical manual. Also indicates emphasis on certain words in a sentence, such as the first instance of a significant term or concept.

2

High Availability

This chapter explains High Availability (HA) in SonicOS and how to connect the firewalls.

Topics:

- About High Availability
- About Active/Standby HA
- About Stateful Synchronization
- About Active/Active DPI
- Active/Standby and Active/Active DPI 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.

ⓘ | **NOTE:** Both firewalls must be the same SonicWall model and run the same firmware version.

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.
HA Control Interface	The interface that carries HA control communication, such as heartbeats, between the two firewalls.
HA Data Interface	The interface that carries stateful synchronization data between the two firewalls.
Active/Active DPI	An HA mode in which the Standby firewall inspects offloaded traffic while the Active firewall forwards all traffic.
Active/Active DPI Interface	The dedicated interface that carries offloaded DPI traffic between the two firewalls.
PPP	Point-to-point protocol that provides a standard method for transporting multi-protocol datagrams over point-to-point links.
PPPoE	A method for transmitting PPP over ethernet.
PPPoE HA	HA PPPoE support function without state synchronization.
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	The passive condition of a hardware unit. The Standby identifier is a logical role that either a Primary or a Secondary hardware unit can assume. The Standby unit assumes the Active role when the Active unit fails.

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.

Operational modes are:

- **None**—When None is selected, HA is disabled.
- **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 is enabled with the Active/Standby configuration to maintain session continuity during failover. In this mode, the dynamic state information 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.

  ⓘ | **NOTE:** Not all information is synchronized in a stateful configuration.
- **Active/Active DPI**—Active/Active Deep Packet Inspection (DPI) is a distinct High Availability operational mode in which both appliances participate concurrently in traffic processing. In this mode, the Active unit handles firewall, NAT, and session processing, while processor-intensive DPI services—such as Intrusion Prevention (IPS), Gateway Anti-Virus (GAV), and Anti-Spyware—are offloaded to the peer appliance for parallel inspection.

This DPI offload architecture improves overall performance by distributing inspection workloads across both security appliances while maintaining a single Active traffic-forwarding role. If a failure occurs, the system transitions to a standard Active/Standby failover behavior until the affected appliance recovers.

High Availability Encryption

High Availability encryption adds security to the communication between the firewalls in an HA pair. HA control messages between the 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 and Active/Active DPI HA modes and does not apply to messages exchanged for stateful synchronization. To enable it, select **Enable Encryption for Control Communication** in the **GENERAL SETTINGS** section on **DEVICE | High Availability > Settings**. Discovery messages (find-peer and found-peer) are transmitted without encryption. After the discovery stage, the firewalls encrypt all control messages:

- 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 firewalls. 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 firewall loses power.

The self-checking mechanism is managed by software diagnostics, which check the complete system integrity of the firewall. The diagnostics check internal system statuses, system process statuses, and network connectivity. Each firewall calculates a connectivity score, and the firewall with the better score becomes Active. This prevents repeated failovers.

Critical internal system processes such as NAT, VPN, and DHCP (among others) are checked in real time. The firewall isolates the failing service as early as possible and attempts to restart it.

In Active/Active DPI mode, crash detection monitors the health of both firewalls. In addition to standard system, process, and link monitoring, the HA subsystem verifies the DPI offload communication path between the two firewalls.

If the Standby firewall that processes the offloaded DPI traffic has a critical failure, such as a system crash, loss of monitored interfaces, or loss of communication between the firewalls, the HA pair stops DPI offloading and continues as a standard Active/Standby pair.

Virtual MAC Address

The Virtual MAC address allows the High Availability pair to share the same MAC address, which 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 firewalls each use their own MAC addresses. Because the firewalls 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 firewall must send a gratuitous ARP message that announces the new MAC address and IP address pair. Until this ARP request propagates through the network, traffic intended for the Active firewall'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 firewalls. When a failover occurs, all routes to and from the Active firewall are still valid for the Standby firewall. 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 firewalls. This reduces 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

PPPoE can be enabled on interfaces in Active/Standby mode without stateful synchronization. PPPoE HA provides HA where a Standby firewall assumes connection to the PPPoE server when the Active firewall fails.

NOTE: One WAN interface must be configured as PPPoE; see Configuring a WAN Interface section in the document on the SonicWall technical documentation site.

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

When the Active firewall fails, its PPPoE connection on the client side times out. The Secondary firewall 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 Active/Standby mode with or without Stateful Synchronization.

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

During a failover, the former Active firewall releases the DHCP lease, and the Standby firewall, which becomes Active, renews the 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 firewall does not have an IP address when failover occurs, the Standby firewall 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 firewall.

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 firewalls' 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.

ⓘ | **NOTE:** X0 needs a routable IP address.

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-firewall 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 firewall through its management IP address (for more information, see *SonicOS 8 Device 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 firewalls. If one firewall can ping the target but the other cannot, however, the HA Pair fails over 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 firewalls running SonicOS to be configured to provide a reliable, continuous connection to their connected networks. One firewall is configured as the Primary unit and acts as the Active firewall. An identical firewall is configured as the Secondary unit and acts as the Standby firewall. In the event of the failure of the Active firewall, the Standby firewall takes over to secure a reliable connection between the connected networks. Two firewalls 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 firewall assumes all network responsibilities when the Primary unit fails, ensuring a reliable connection between the connected networks.
- **Cost-effectiveness** - Active/Standby HA is a cost-effective option for deployments that provide high availability with redundant firewalls. 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 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 firewalls.

Working of Active/Standby HA

HA requires one SonicWall firewall configured as the Primary SonicWall, and an identical firewall configured as the Secondary SonicWall. During normal operation, the Primary SonicWall is in an Active state and the Secondary SonicWall in a 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 firewall, 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, or when the firewall 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.

  ⓘ | **NOTE:** The complete synchronization reboots the Standby unit.

About Stateful Synchronization

Stateful Synchronization provides 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 firewall fails. Most existing network connections continue after the failover.

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 firewall failure.

- **Faster failover performance** - By maintaining continuous synchronization between the Active and Standby firewalls, Stateful Synchronization enables the Standby firewall to take over in case of a failure with virtually no down time or loss of network connections.
- **Minimal impact on CPU performance** - Stateful Synchronization uses little CPU capacity.
- **Minimal impact on bandwidth** - Transmission of synchronization data is throttled so that it does 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 firewall handles all traffic. When Stateful Synchronization is enabled, the Primary firewall actively communicates with the Secondary to update most network connection information. As the Primary firewall creates and updates network connection information (such as VPN tunnels, active users, connection cache entries), it immediately informs the Secondary firewall. This ensures that the Secondary firewall 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 firewall and automatically propagated to the Standby firewall. The High Availability pair uses the same LAN and WAN IP addresses—regardless of which firewall is currently Active.

When using SonicWall Network Security Manager (NSM) to manage the firewalls, NSM logs into the shared WAN IP address. In case of a failover, NSM administration continues, and NSM administrators who are logged in to the firewall stay logged in. However, **GET** and **POST** commands may time out without a reply.

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

Information that is Synchronized	Information that is not Synchronized
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 firewall creates a session for the user.
2. The Active firewall synchronizes with the Standby firewall. 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. The new Active firewall begins to send gratuitous ARP messages to the connected switches. It uses the same Virtual MAC address and IP address as the former Active firewall. No routing updates are necessary for downstream or upstream network devices.
6. When the PC user attempts to access a web page, the new Active firewall has all of the user's session information and is able to continue the user's session without interruption.

About Active/Active DPI

Active/Active Deep Packet Inspection (DPI) is a High Availability (HA) mode that requires Stateful Synchronization. It distributes DPI processing between the two firewalls of the HA pair. The Active firewall handles routing, firewall enforcement, Network Address Translation (NAT), session management, and other core functions. The Standby firewall processes the offloaded DPI traffic.

Active/Active DPI does not provide two independent traffic-forwarding firewalls. Network traffic enters and exits through the Active firewall only. Both firewalls inspect the traffic.

Active/Active DPI offloads eligible TCP traffic. Traffic is eligible when a DPI security service inspects it and it uses a supported application protocol. The following security services inspect offloaded traffic:

- Gateway Anti-Virus (GAV)
- Gateway Anti-Spyware
- Intrusion Prevention Service (IPS)
- Application Control
- Capture Advanced Threat Protection (Capture ATP)

Active/Active DPI supports the following TCP application protocols: HTTP, HTTPS, FTP, IMAP, SMTP, POP3, CIFS/SMB, and SSH. HTTPS inspection requires DPI-SSL. SSH inspection requires DPI-SSH. The Active firewall processes UDP traffic and all other traffic locally.

Active/Active DPI requires a dedicated physical interface, the **Active/Active DPI Interface**. For the requirements, see Connecting the Active/Active DPI Interface.

The Active firewall forwards eligible traffic to the Standby firewall through this interface. The Standby firewall performs the DPI processing and returns the inspection results over the same interface. The Active firewall handles all remaining packet processing.

The amount of traffic offloaded depends on the Active/Active DPI Traffic Offload % and Active/Active DPI CPU Threshold % settings. For more information, see Configuring Advanced High Availability Settings.

ⓘ **NOTE:** Active/Active DPI is supported only on the NSa 4800, NSa 5800, and NSa 6800 firewalls, which use x86 processors. ARM-based firewalls and NSv virtual firewalls do not support it. For licensing information, see Supported Platforms and Licensing for HA.

Benefits of Active/Active DPI

Active/Active DPI uses the unused CPU resources of the Standby firewall in a High Availability (HA) pair. All network traffic enters and exits through the Active firewall. The Active firewall offloads DPI workloads to the Standby firewall, which does not forward traffic.

In Active/Active DPI mode:

- The **Standby** firewall performs heavy DPI operations such as IPS, GAV, Gateway Anti-Spyware, Application Control, and Capture ATP.
- The **Active** firewall handles all non-DPI components, including routing, NAT, session management, and firewall enforcement.
- This division of processing distributes DPI workloads across both firewalls and reduces the CPU load on the Active firewall.

Active/Active DPI does not change HA failover behavior. The offloaded traffic uses the dedicated Active/Active DPI interface, which is separate from the HA control and data links.

Active/Standby and Active/Active DPI 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 Firewalls
- Connecting the Active/Active DPI Interface

Supported Platforms and Licensing for HA

Licenses included with the purchase of a SonicWall security appliance are shown in **HA licenses available with SonicWall firewalls** table.

HA Licenses Available With Sonicwall Network Security Firewalls

Platform	Active/Standby HA	Stateful HA	Active/Active DPI
TZ80	Included	Included	Not supported
TZ280 / TZ280W / TZ280P	Included	Included	Not supported
TZ380 / TZ380W	Included	Included	Not supported
TZ480	Included	Included	Not supported
TZ580	Included	Included	Not supported
TZ680	Included	Included	Not supported
NSa 2800	Included	Included	Not supported
NSa 3800	Included	Included	Not supported
NSa 4800	Included	Included	Included
NSa 5800	Included	Included	Included
NSa 6800	Included	Included	Included
NSv XS	Included	Included	Not supported
NSv S	Included	Included	Not supported
NSv M	Included	Included	Not supported
NSv L	Included	Included	Not supported

ⓘ | **NOTE:** HA Stateful licensing is standard across all models.

NOTE: Active/Active DPI is supported only on the SonicOS 8 x86-based **NSa 4800**, **NSa 5800**, and **NSa 6800** appliances. ARM-based appliances and NSv virtual firewalls do not support Active/Active DPI. Support for Active/Standby HA or Stateful HA does not imply support for Active/Active DPI.

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 firewall. For further information, see *SonicOS 8 Device 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.

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.

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 Firewalls

NOTE: For complete procedures for connecting your firewalls, see the Quick Start Guide for your firewall. For procedures for connecting firewalls in an Active/Active Clustering deployment, see the Active/Active Clustering documentation.

NOTE: If you are connecting the Primary and Secondary firewalls 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 firewalls.

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.

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

TIP: SonicOS firewalls 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 firewalls on MySonicWall and for synchronizing licensing information. If network policy permits live communication with the SonicWall licensing server, connect the WAN interface before you register and license the firewalls.

Connecting the Active/Active DPI Interface

Active/Active DPI requires an additional dedicated physical connection between the two firewalls of the High Availability (HA) pair. This connection is the **Active/Active DPI Interface**.

The Active/Active DPI interface must be separate from the **HA Control Interface** and the **HA Data Interface**. The HA Control Interface carries HA control communication, and the HA Data Interface carries stateful synchronization data. Neither interface carries Active/Active DPI offloaded traffic.

The Active/Active DPI interfaces must meet the following requirements:

- Use the same interface number on both firewalls, for example, X5 on both.
- Appear initially as unassigned, unused interfaces on **NETWORK | System > Interfaces**.
- Connect directly with a standard Ethernet cable.

For example, if X5 is unassigned on both firewalls, connect X5 on the Primary firewall directly to X5 on the Secondary firewall. Select X5 as the **Active/Active DPI Interface** when you configure HA. After you enable Active/Active DPI, SonicOS assigns the interface to the **HA Data-Link** zone, which isolates HA traffic from other zones. The Active/Active DPI interface remains a separate connection from the HA Data Interface.

The Active firewall forwards eligible traffic to the Standby firewall through this connection. The Standby firewall performs DPI processing and returns the inspection results over the same connection.

Maintenance

Topics:

- Removing an HA Association
- Replacing a SonicWall firewall

Removing an HA Association

You can remove the association between two SonicWall firewalls on MySonicWall at any time. You might need to remove an existing HA association if you replace a firewall or reconfigure your network. For example, if one of your SonicWall firewalls fails and you need to replace it, or you might want to swap the Primary and Secondary roles 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 firewall or changes which unit is the Primary firewall and which is the Secondary firewall (see Replacing a SonicWall firewall).

To remove the association between two registered SonicWall firewalls:

1. Log in to MySonicWall.
2. In the left navigation bar, navigate to **My Workspace > Tenant Products**.
3. Scroll down to find the secondary firewall 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 firewall:
    a. Click **Remove** under **ACTIONS**.
    b. Wait for the page to reload.
    c. Scroll down.
    d. Click **Remove** again.

Replacing a SonicWall firewall

If your SonicWall firewall 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 firewall 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 firewall 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 firewall (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 firewall 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 firewall.
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.

3

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.

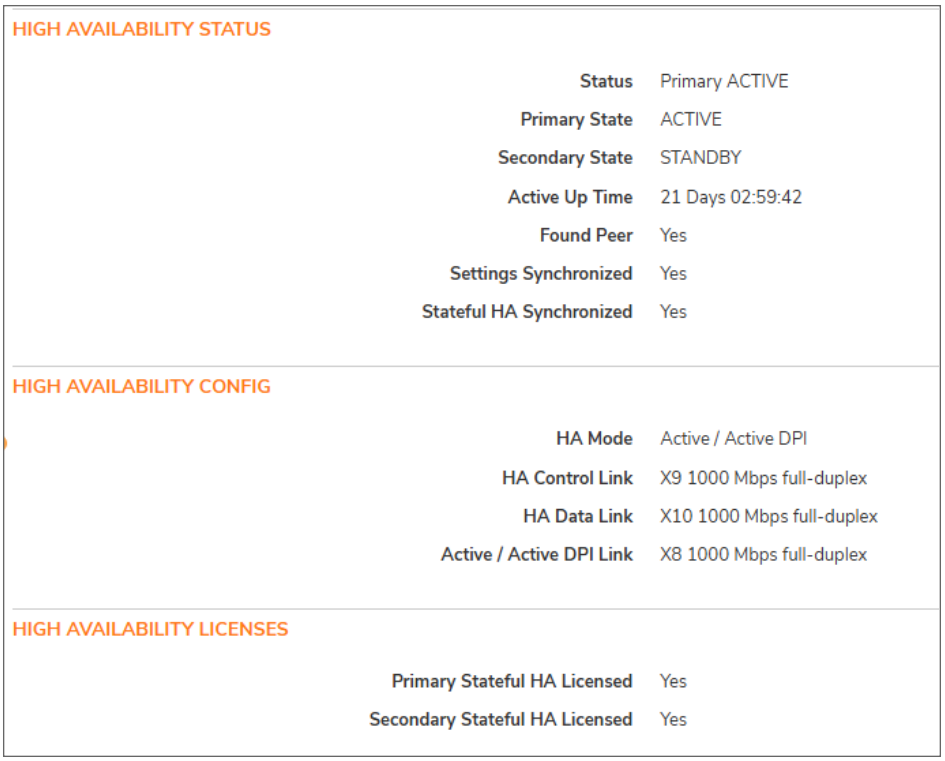

Topics:

- 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**.

NOTE: When Active/Active DPI is enabled, the system continues to report the firewalls with the standard Primary and Secondary roles and the Active and Standby states. The system indicates Active/Active DPI only in the **HA Mode** field. For more information, see High Availability Config.

High Availability Config

The **High Availability Config** section on the **DEVICE | High Availability > Status** 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.
  - **Active/Active DPI** - Distributes eligible DPI processing between the two firewalls of a Stateful HA pair. The Active firewall forwards all traffic and handles routing, firewall enforcement, NAT, and session management. The Standby firewall processes offloaded DPI traffic through the dedicated Active/Active DPI interface. Offloading is controlled by the **Active/Active DPI Traffic Offload %** and **Active/Active DPI CPU Threshold %** settings. The pair has one traffic-forwarding firewall.
- **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 transfers stateful data to keep session data synchronized between the units. The HA Data Link is not required in non-stateful HA. It is required in Stateful HA and Active/Active DPI.
- **Active/Active DPI Link** - Indicates the port, speed, and duplex settings of the dedicated Active/Active DPI link when the mode is Active/Active DPI. This field is not displayed in other modes.

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.

4

Configuring High Availability

Topics:

- Configuring Active/Standby High Availability Settings
- Configuring HA with Dynamic WAN Interfaces
- Configuring Active/Active DPI High Availability Settings

IMPORTANT: You cannot use High Availability with PortShield, except with 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. SonicOS synchronizes them to the Secondary firewall.

To configure Active/Standby:

1. Navigate to **DEVICE | High Availability > Settings**.
2. In the **GENERAL SETTINGS** section, do the following:
    a. From the **Mode** drop-down menu, select **Active/Standby**.
    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. Click **OK** in the information dialog that displays.

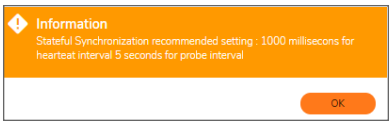

d. (Optional) To configure the High Availability Pair so that the Primary firewall takes back the Primary role when it restarts after a failure, select **Enable Preempt Mode**. This option is not selected by default.

ⓘ **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.

ⓘ **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.

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

A confirmation message displays:

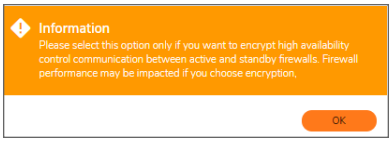

h. (Optional) Click **OK**.

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 8 System* document.
3. Navigate to **DEVICE | High Availability > Settings**.
4. In the **GENERAL SETTINGS** section, do the following:
    a. From the **Mode** drop-down menu, select **Active/Standby**.
    b. If a dialog displays, read the message and click **OK**.
    c. Verify that **Enable Stateful Synchronization** is not selected. This option is not selected by default. Stateful HA requires a license on some platforms.
    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**.
    ⓘ **NOTE:** 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**.

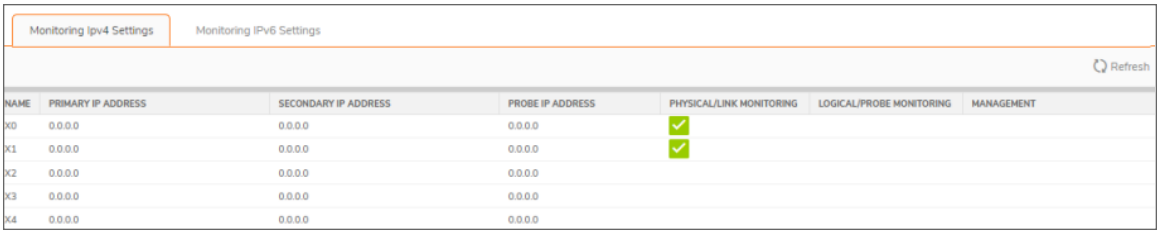

8. Select the PPPoE interface and click the **Edit** icon.
   Interface Monitoring Settings dialog is displayed.

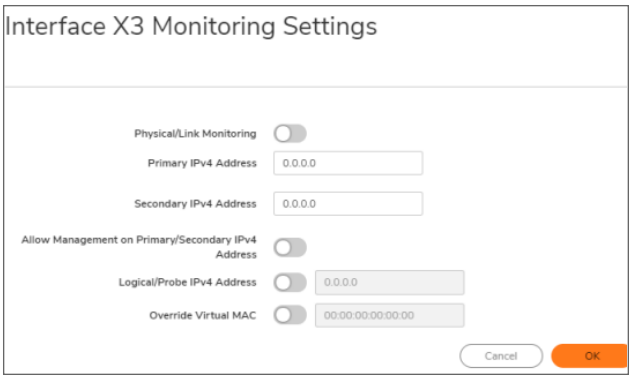

9. Verify that **Physical/Link Monitoring** is selected. This option is selected by default.
10. Ensure the **Primary IPv4 Address** and **Secondary IPv4 Address** fields are set to 0.0.0.0.
11. Ensure none of the other options are selected.
12. Click **OK**.

Configuring Active/Active DPI High Availability Settings

Configure Active/Active DPI on the Primary firewall. SonicOS synchronizes the settings to the Secondary firewall.

Before you configure Active/Active DPI, verify that both firewalls are supported models and that each firewall has an unassigned interface for the dedicated connection. For the requirements and cabling, see Connecting the Active/Active DPI Interface.

To configure Active/Active DPI:

1. Navigate to **DEVICE | High Availability > Settings**.
2. In the **GENERAL SETTINGS** section, do the following:
    a. From the **Mode** drop-down menu, select **Active/Active DPI**. A message displays that states that Active/Active DPI requires current licenses and signature updates on both firewalls.
    b. Click **OK**.
    SonicOS turns on **Enable Stateful Synchronization** automatically and dims the option.
    c. Verify that **Enable Preempt Mode** is not selected. This option is not selected by default.
    ⓘ **NOTE:** Do not enable Preempt Mode for Active/Active DPI. Preempt Mode instructs the Primary firewall to take back the Primary role when it restarts after a failure. If you enable it, the Primary firewall takes the Active role again after each restart, and the pair fails over a second time.
    d. (Optional) Select **Enable Virtual MAC** to allow both 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.

ⓘ **NOTE:** 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.

e. (Optional) Select **Enable Encryption for Control Communication** to encrypt HA control communication between the active and standby firewalls. This option is not selected by default.

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

A confirmation message displays:

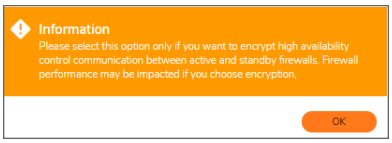

f. Click **OK**.

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 is displayed if the firewall detects that the interface is already configured.

b. Select the interface for the **HA Data Interface**. This interface is required for Active/Active DPI.

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

c. From the **Active/Active DPI Interface** drop-down menu, select the dedicated interface.

If the interface is already configured, the field displays the configured interface and is dimmed.

The Active/Active DPI interface carries offloaded DPI traffic between the appliances. It must be different from the interfaces selected for **HA Control Interface** and **HA Data Interface**.

Only available, unassigned interfaces appear in the drop-down menu. For the interface requirements, see Connecting the Active/Active DPI Interface.

d. Click **Accept** to complete the High Availability configuration.

SonicOS synchronizes all settings to the Secondary firewall, and the Secondary firewall reboots.

To configure when offloading begins and the percentage of eligible DPI traffic to offload, see Configuring Advanced High Availability Settings.

5

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

This section describes how to deploy a High Availability environment with two SonicWall NSv firewalls in the Microsoft Azure cloud platform.

Azure provides cloud capabilities through platform as a service (PaaS) and infrastructure as a service (IaaS) models.

Product Matrix Table:

Product Models	NSv XS	NSv S	NSv M	NSv L
Maximum Cores	1	2	4	8
Minimum Total Cores	1	2	4	8
Management Cores	1	1	1	1
Maximum Data Plane Cores	1	1	3	7
Minimum Data Plane Cores	1	1	1	1

NOTE: HA requires a minimum of three interfaces for HA exchange messages. Select the Standard DS3_v2 virtual machine size for the Gen 8 Active/Standby template deployment. The SonicWall custom template sets this value by default.

For an Active/Active deployment with an Azure load balancer, the minimum number of interfaces is two. This deployment is separate from Active/Active DPI. For more information, see How to Deploy SonicWall NSv High Availability Using Azure Load Balancers.

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

Gen 8 NSvHA-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
Gen 8 NSvHA-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

NOTE: For the HA interface, use only a /24 subnet. Other interfaces, such as X0 and X1, have no such limitation.

Topology:

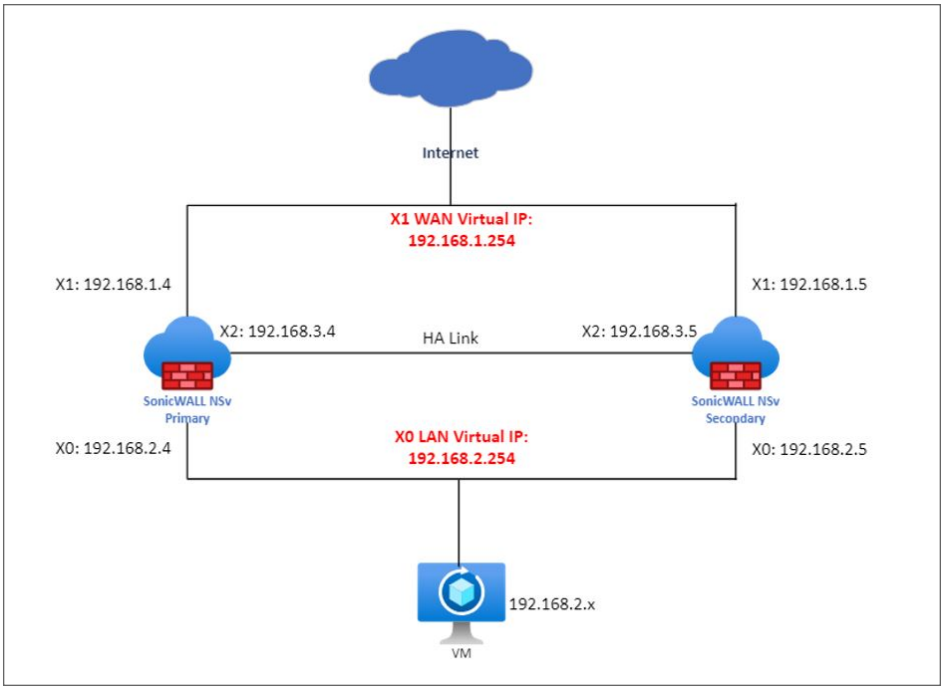

Topics:

- Install the Custom Template
- Enable Managed Identity on 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 are already set up according to best practices. To install the custom template, follow these steps:

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

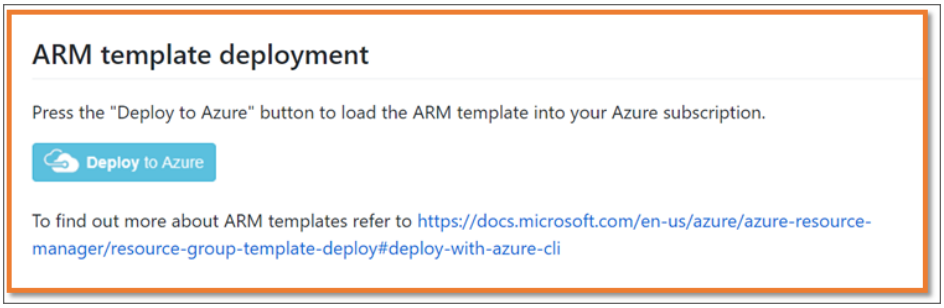

2. Log in to the **Azure Portal** using valid credentials.
3. The custom **Template** displays. Proceed through the tabs and complete the fields, such as Resource Group (SonicWall recommends that you create a new one) and Storage Account.

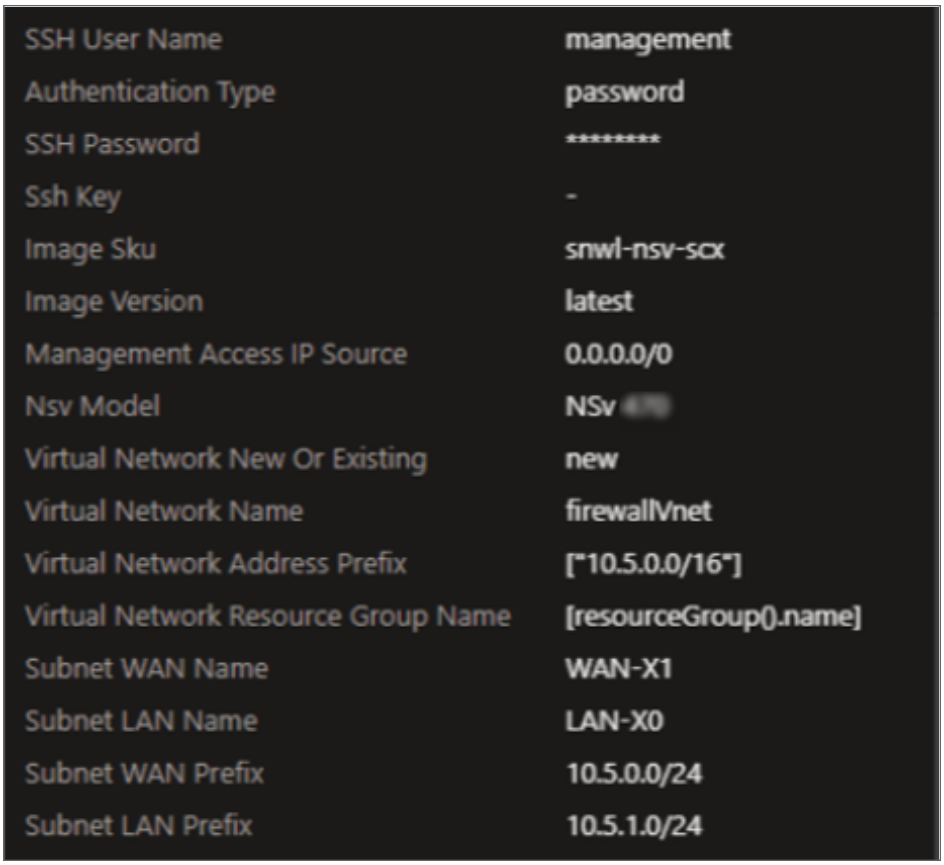

NOTE: Store the SSH password and key pair securely. You need them for console or SSH access to the firewall.

NOTE: Leave the **Image Version** field set to **latest** to install the Gen 8 NSv firewall.

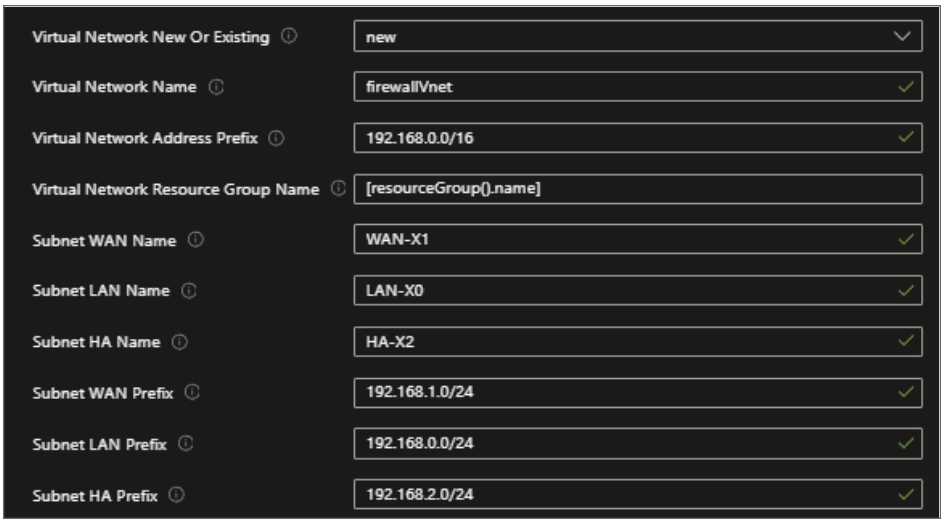

The custom template creates two virtual machines (HA1 and HA2) with the LAN, WAN, and HA interfaces that the deployment requires.

Enable Managed Identity on Both Virtual Machines (HA1 and HA2)

To enable managed identity:

1. In the left panel, select **All Services** > **Virtual machines**.
2. Search for the Primary VM you created during deployment. On the left panel, select **Identity** and change the status to **On**. If the template already configures this setting, leave it unchanged.
3. Repeat steps 1 and 2 for the **Secondary VM**.

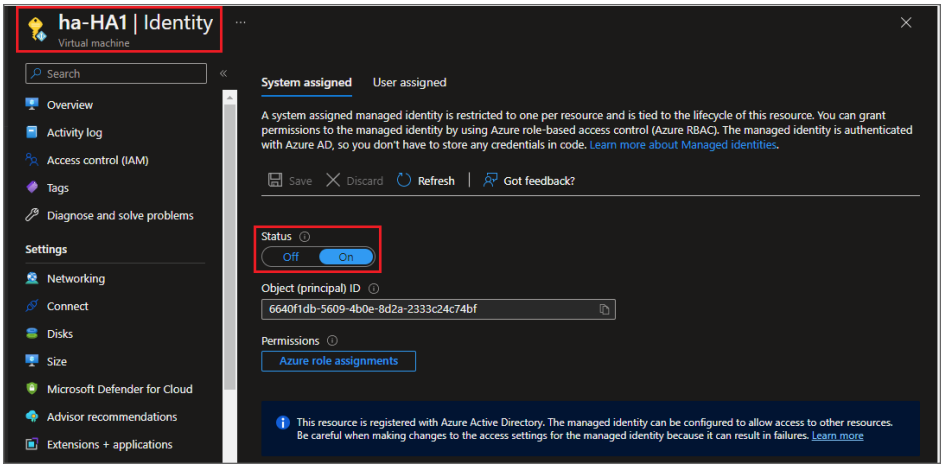

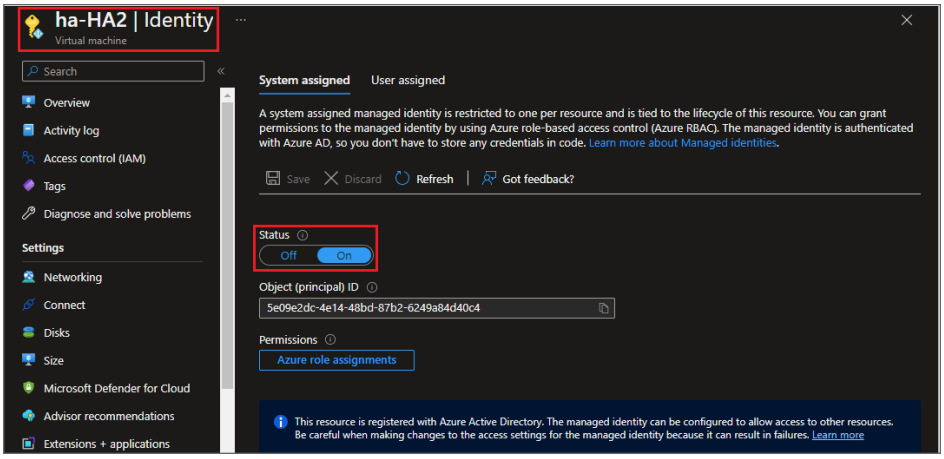

Role Assignment

Add a role assignment to the Resource Group.

Set the **Contributor** role assignment on the Resource Group. The role allows the firewalls to modify network interface settings and to exchange HA information during a failover. To assign the role:

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**.

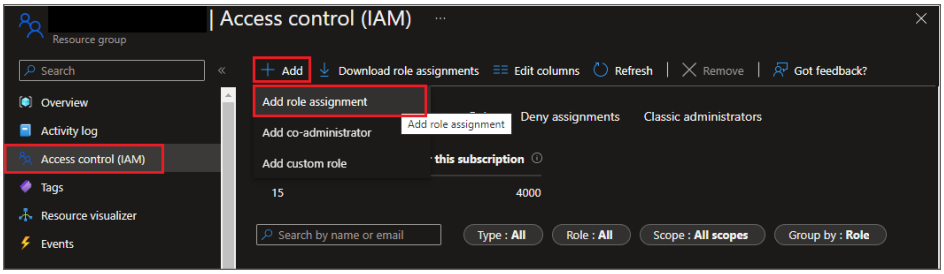

3. Click **Contributor**.

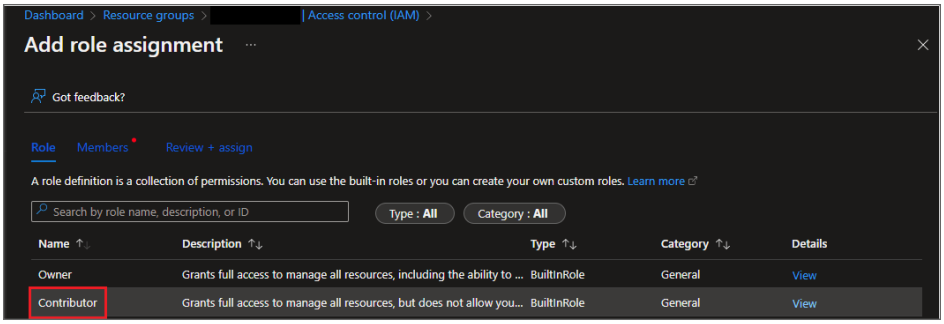

4. On the **Contributor** page, select **Managed identity**, verify that the subscription is correct, and select the managed identity type **Virtual machine**. Select the virtual machines **HA1** and **HA2**. If the template already configures this setting, leave it unchanged.

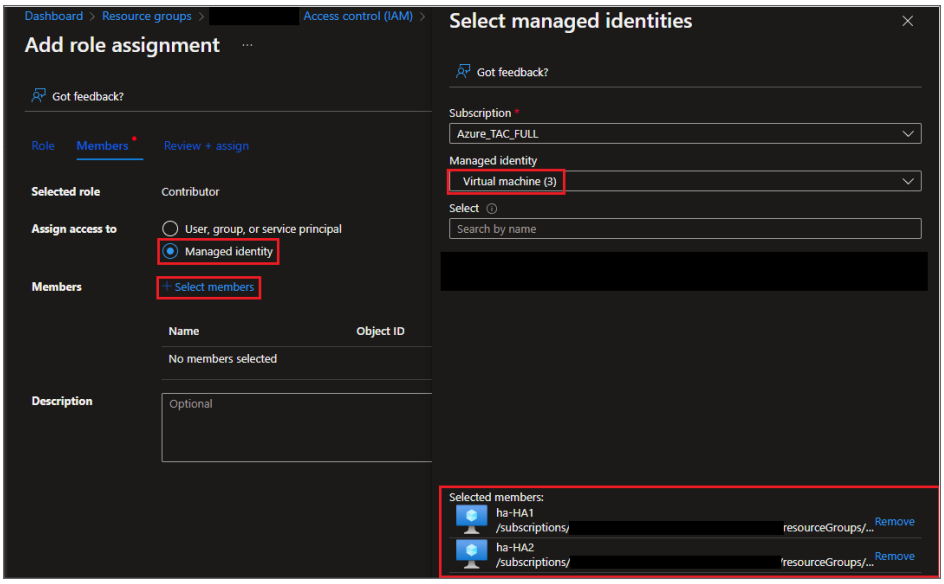

NOTE: It is recommended that VNet and NSv Virtual Machines be in the same Resource Group. If they are in separate Resource Groups, NSv Virtual Machines must also be added as Contributors to the VNet Resource Group.

Check the Networking Tab

Check the **Networking** tab.

1. In the left panel, select **All Services** > **Virtual machines**.
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**.

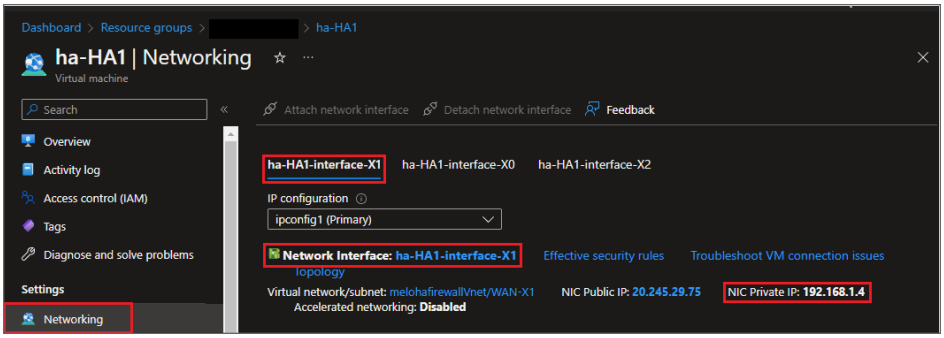

The template automatically configures an additional secondary IP configuration, which acts as the floating IP address. The template creates a floating IP address for both the WAN and LAN interfaces. You need these IP addresses in the next configuration steps.

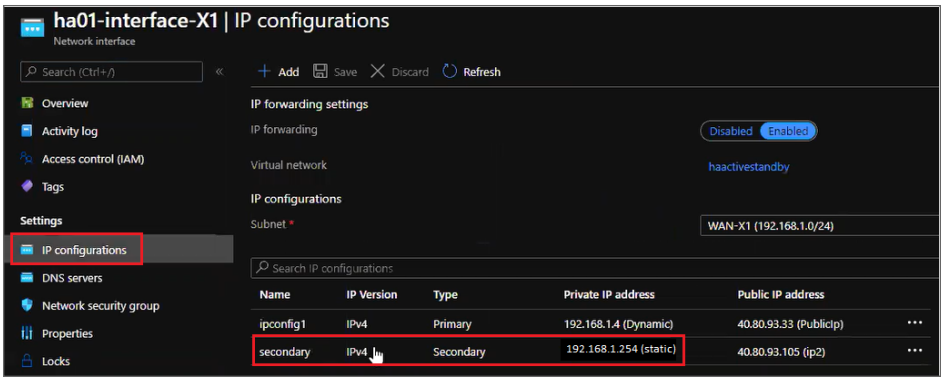

Repeat the previous steps to access **HA1-interface-X0** and check the **Secondary Interface**.

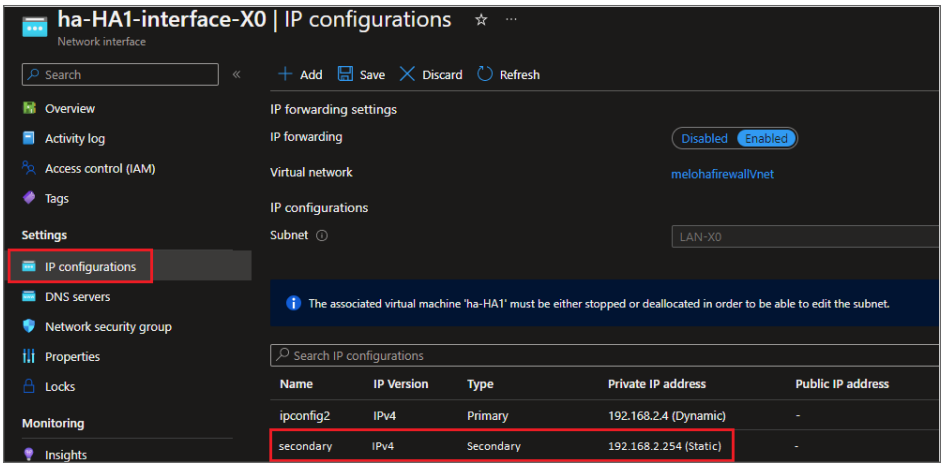

You have now confirmed that the secondary IP address is set on both X1 and X0. Set those IP addresses on the Active NSv firewall in the next steps.

ⓘ | **NOTE:** Register the firewall the first time you log in.

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 lose access to the firewall after you change X1. Before you change X1, take a screenshot of the current configuration.

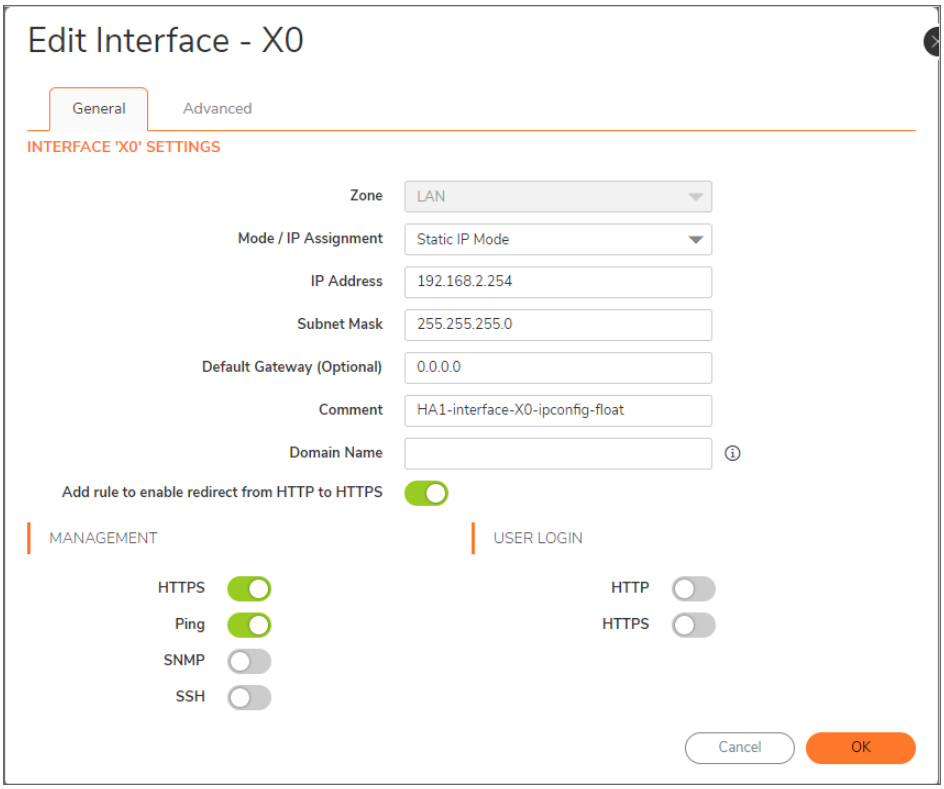
Edit Interface - X0
General
Advanced
INTERFACE 'X0' SETTINGS
Zone
LAN
Mode / IP Assignment
Static IP Mode
IP Address
192.168.2.254
Subnet Mask
255.255.255.0
Default Gateway (Optional)
0.0.0.0
Comment
HA1-interface-X0-ipconfig-float
Domain Name
Add rule to enable redirect from HTTP to HTTPS
MANAGEMENT
HTTPS
Ping
SNMP
SSH
USER LOGIN
HTTP
HTTPS
Cancel
OK

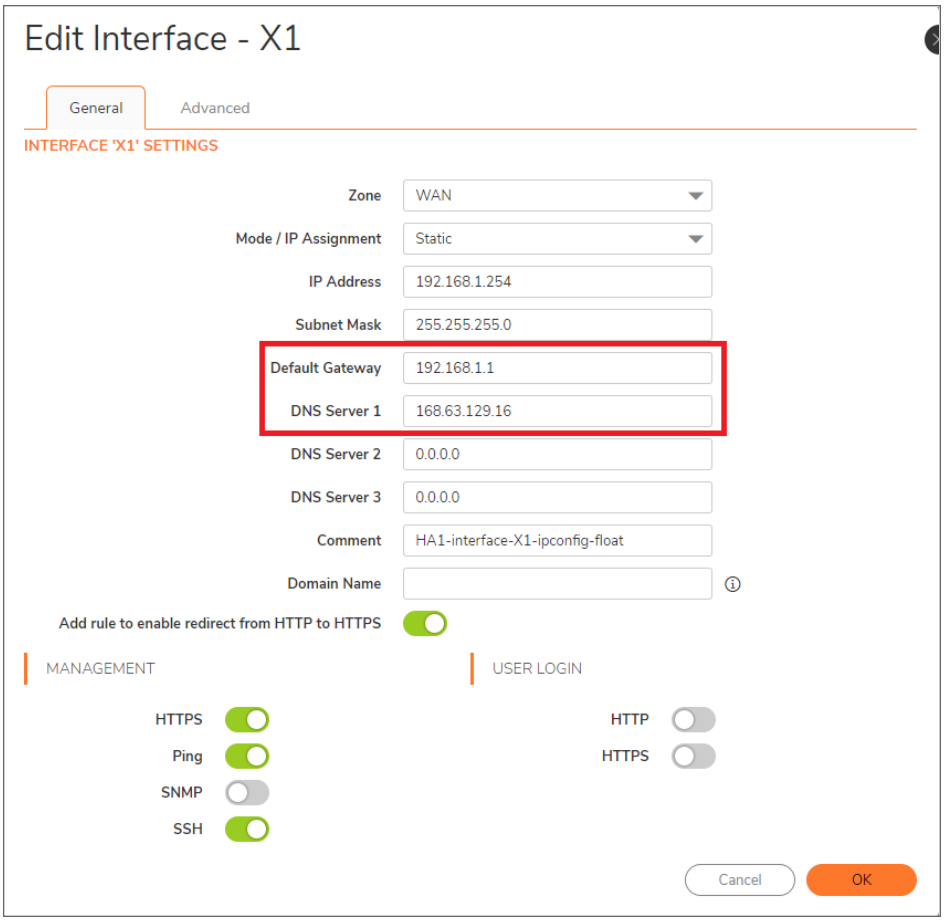

2. Keep the Default Gateway and DNS Server 1 unchanged. To keep access to the Active firewall, use the Floating Public IP address to log in again after you change X1.
3. After the changes, the Interfaces should look similar to the screenshot below:

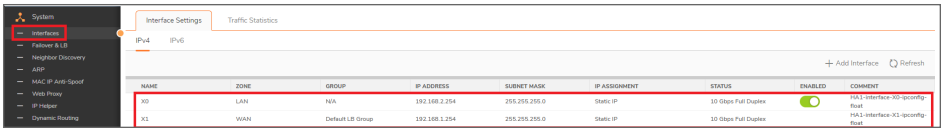

Configuring Standby NSv Firewall Using the Associated Public IP

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**

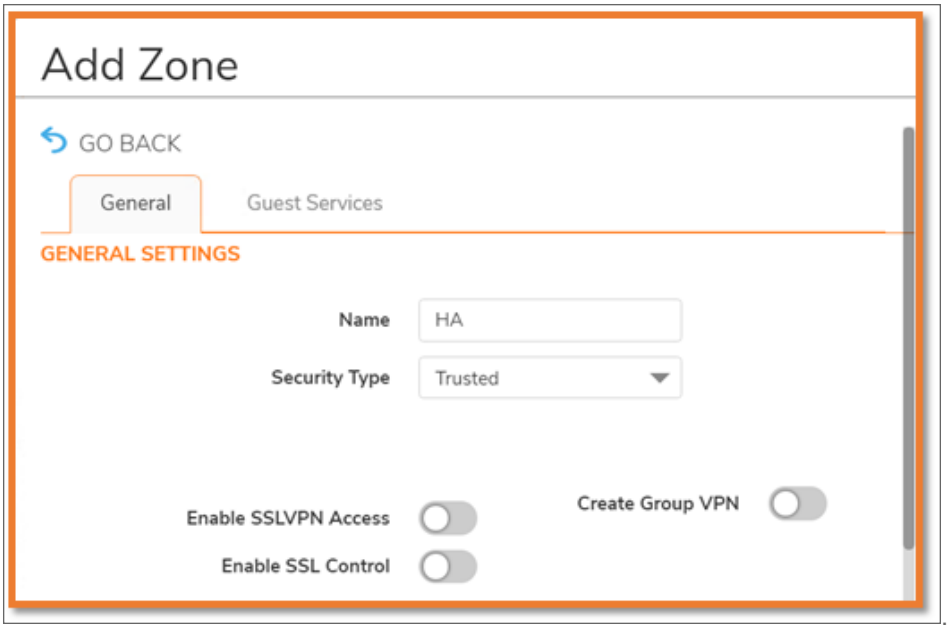

2. The IP Address is the **ha-HA2-Interface-X2** found on the **Networking** tab.

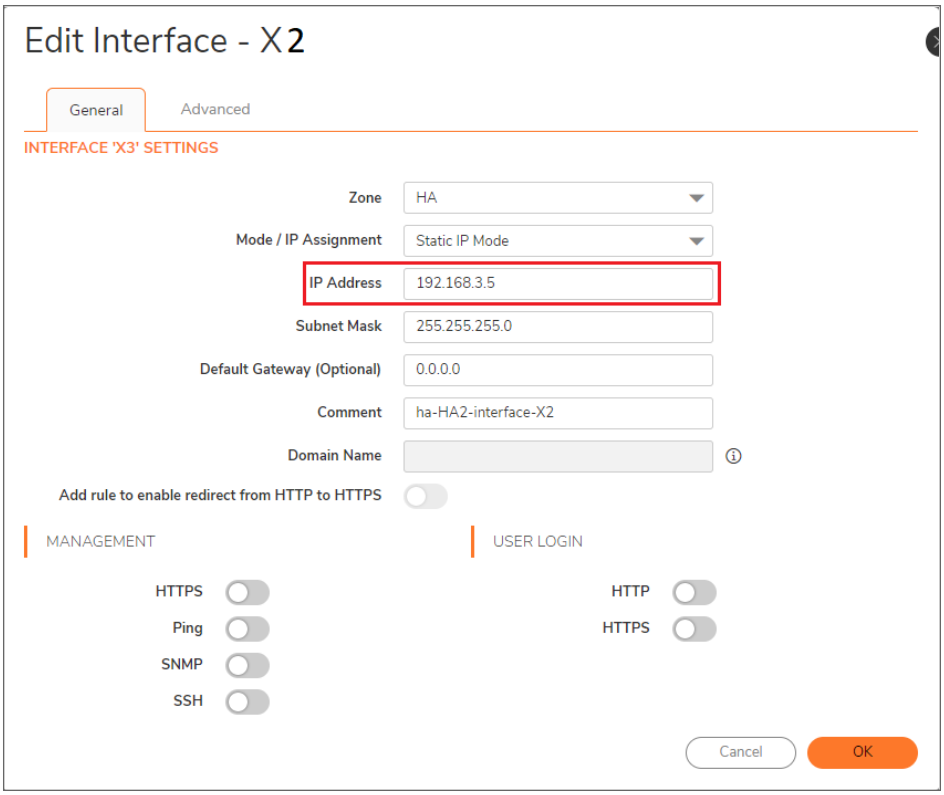

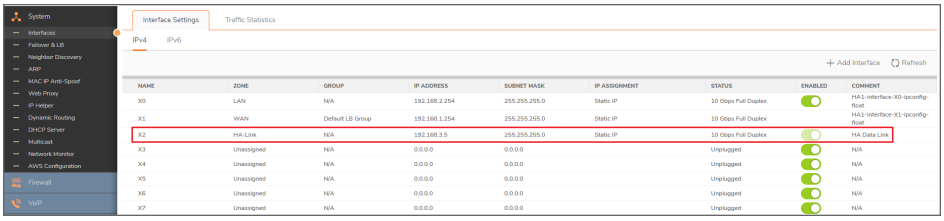

Enable the L3 Mode

Enable Layer 3 (L3) mode on the Standby firewall. L3 mode allows the HA control link to use IP routing between the firewalls.

1. Navigate to **DEVICE | High Availability > Settings**, and go to the **HA INTERFACES** section, and select **Enable L3 Mode** on the Standby firewall.

   NOTE: Configuration of the Primary IP and Secondary IP on the standby unit is blocked. These IPs

ⓘ | must be configured only on the active unit.

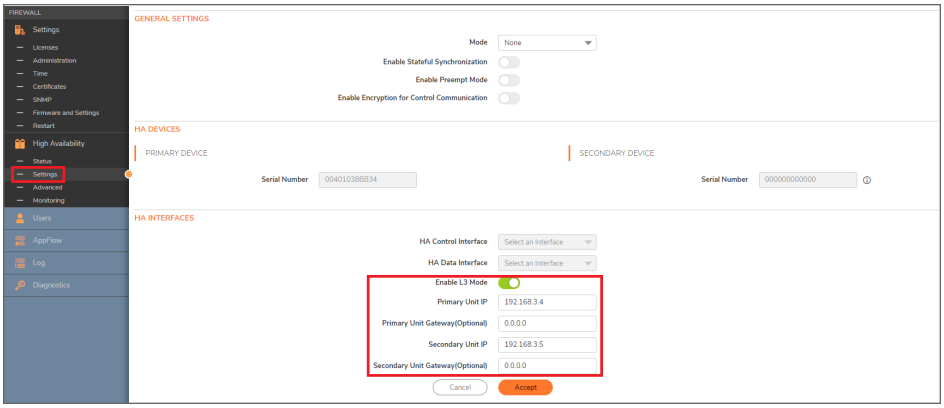

Deploying NSv XS in Azure Using L3

This section describes how to deploy SonicWall NSv XS firewalls in Azure using L3 Active/Standby High Availability.

These procedures apply only to Azure deployments and are not required for on-prem environments.

Topics:

- Prerequisites
- Primary NSv Configuration
- Secondary NSv Configuration
- Azure Permissions for Floating IP Support
- Monitoring IP Configuration

Prerequisites

Before deploying NSv High Availability in Azure, ensure the following requirements are met:

- Virtual Machine Size: Supported Azure VM size meeting minimum core and interface requirements (for example, Standard D4s_v5 or recommended compute-optimized series)
- Firmware Version: SonicOS 8.2.3 or later (both NSv instances must run the same firmware version)
- Resource Group: Both NSv instances must reside in the same Azure Resource Group
- High Availability Mode: Active/Standby
- Layer 3 (L3) HA: Enabled

Interface Mapping

Azure requires a dedicated network interface for each SonicWall interface used in High Availability.

ⓘ | **NOTE:** Ensure all NICs are attached to the appropriate subnets and security groups before continuing.

Azure NIC	SonicWall Interface	Description
NIC 1	X1	WAN
NIC 2	X0	LAN
NIC 3	X2	HA Control/Data Link

IP Address Allocation

The following table shows an example IP addressing scheme for Active/Standby HA in Azure.

Interface	Primary NSv (Static)	Secondary NSv (Static)	Floating IP (HA)
X1 (WAN)	192.168.1.35	192.168.1.53	192.168.1.12
X0 (LAN)	192.168.180.27	192.168.180.38	192.168.180.14
X2 (HA)	192.168.183.7	192.168.183.8	N/A

Floating IP addresses represent the production traffic IPs that migrate between units during an HA failover.

Primary NSv Configuration

On the **Primary NSv**:

1. Log in to the management interface.
2. Navigate to **DEVICE | High Availability > Settings**.
3. Verify the following:
    - **Mode**: Active/Standby
    - **Enable L3 Mode**: Enabled
    - **Secondary Device Serial Number**: Configured
    - **Primary Unit IP** Address: 192.168.183.7 (X2)
    - **Secondary Unit IP** Address: 192.168.183.8 (X2)
4. Click **Accept**.

The NSv should report a status of **Primary Active**.

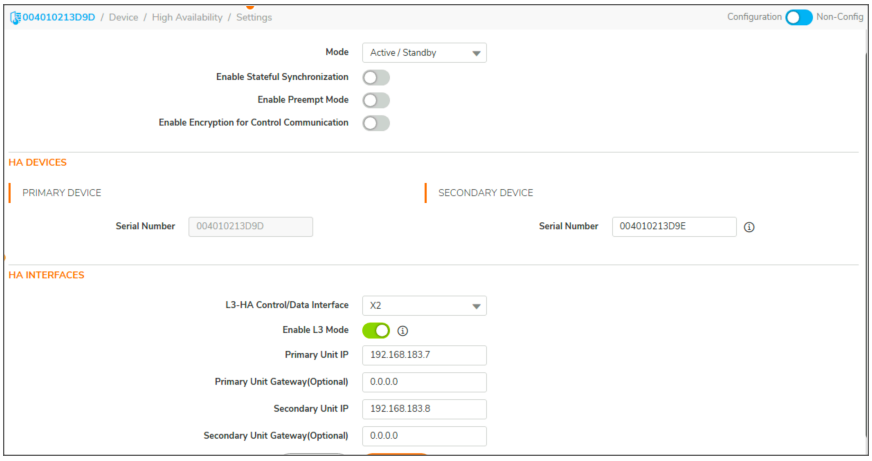

Secondary NSv Configuration

On the Secondary NSv:

1. Log in to the management interface.
2. Navigate to **NETWORK | System > Interfaces** and configure interface **X2** with the static IP address 192.168.183.8.
3. Navigate to **DEVICE | High Availability > Settings**.
4. Select **Enable L3 Mode** if it is not already selected.
5. Click **Accept**. The Secondary firewall should report a status of **Secondary Standby**.
6. Save the configuration.
   The secondary unit should transition to Standby state.

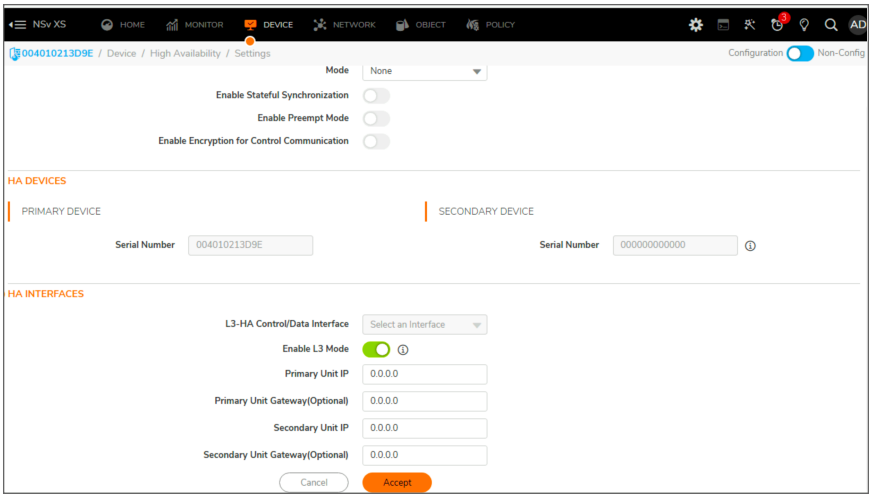

Azure Permissions for Floating IP Support

Azure requires the NSv XS appliances to have permission to modify network interface IP assignments during HA failover.

Enable Managed Identity

For both NSv virtual machines:

1. In the Azure Portal, go to the **Virtual Machine** settings.
2. Navigate to **Identity**.
3. Set **System Assigned Managed Identity** to **On**.

Assign Resource Group Permissions

1. Navigate to the Azure **Resource Group** containing both NSv instances.
2. Open **Access Control (IAM)**.
3. Add a role assignment:
    - **Role**: Contributor
    - **Members**: Managed identities of both NSv VMs

ⓘ **IMPORTANT:** Without these permissions, the NSv firewalls cannot move floating IP addresses during an HA failover.

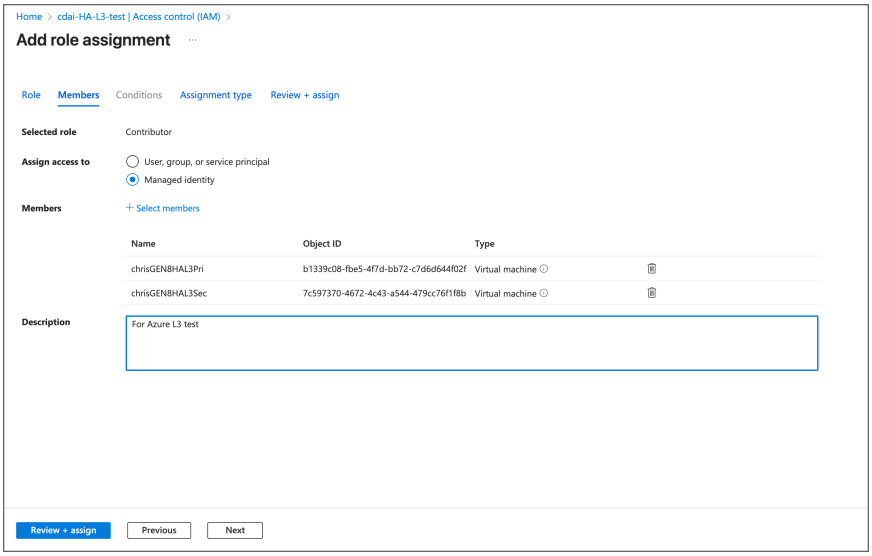

Monitoring IP Configuration

Monitoring IPs provide persistent management access to each NSv regardless of HA state. Configure Monitoring IPs under **Device** > **High Availability** > **Monitoring**:

X0 Monitoring IPs

- Primary: 192.168.180.27
- Secondary: 192.168.180.38

X1 Monitoring IPs

- Primary: 192.168.1.35
- Secondary: 192.168.1.53

Monitoring IPs remain bound to each firewall's network interface and do not move during an HA failover.

Configuring Active NSv Firewall Using the Floating Public IP

On the Active firewall, log in with the Floating Public IP address. To get the IP address:

1. On the Azure portal, navigate to **Resource Groups**.
2. Select the **VM ha-HA1**.
3. On the left panel, click **Networking**.
4. Select the network interface **ha-HA1-Interface-X1**.
5. On the left panel, click **IP configurations**.
6. Copy the **Floating Public IP** address and paste it into the address bar of your 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. Navigate to **DEVICE | High Availability > Settings** 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 not needed if the two HA interfaces are in the same subnet. If they are in different subnets, enter a gateway address. The Azure default is x.x.x.1.

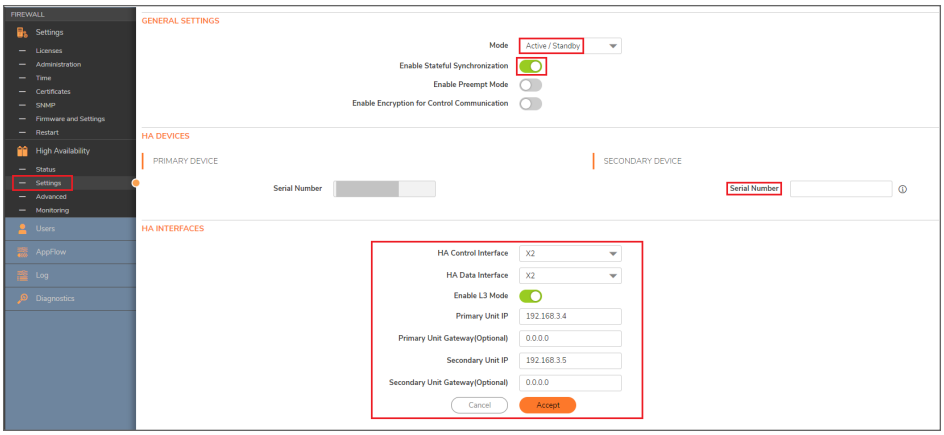

3. Add monitoring IPs: **X0** is optional and **X1** is mandatory.

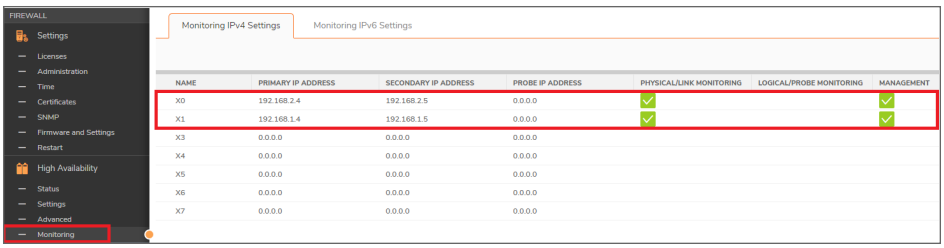

4. Navigate to **DEVICE | High Availability > Status** and verify that the pair forms. The Secondary firewall reboots. The pair forms in a few minutes, and the status shows **Primary Active** and **Secondary Standby**.

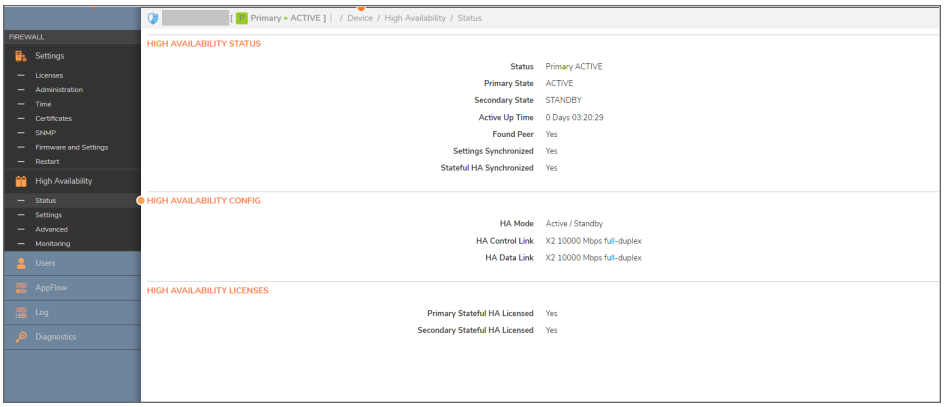

ⓘ | **NOTE:** Multiple floating IP address support in Azure requires SonicOS 8.2.3 or later.

Adding Additional Floating Public IP

Adding an additional Floating Public IP address is optional.

To add an additional Floating Public IP address:

1. On the Azure portal, navigate to **Resource Groups**.
2. Select the **VM ha-HA1**.
3. On the left panel, click **Networking**.
4. Select the network interface **ha-HA1-Interface-X1**.
5. On the left panel, click **IP configurations**.
6. Click **Add** to Add IP configuration and specify the name and IP allocation method.

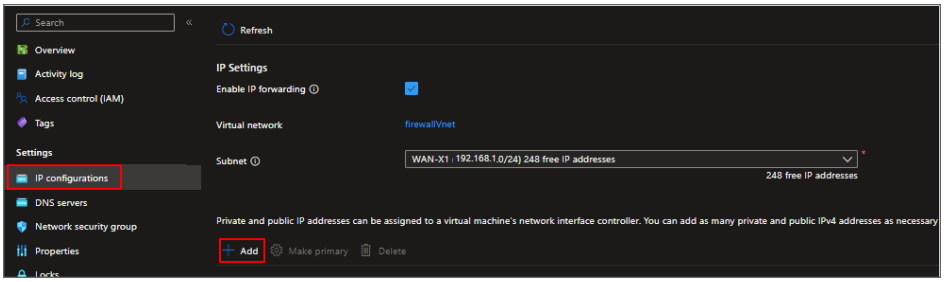

7. **Associate the public IP** by using an existing public IP or create a new public IP address.

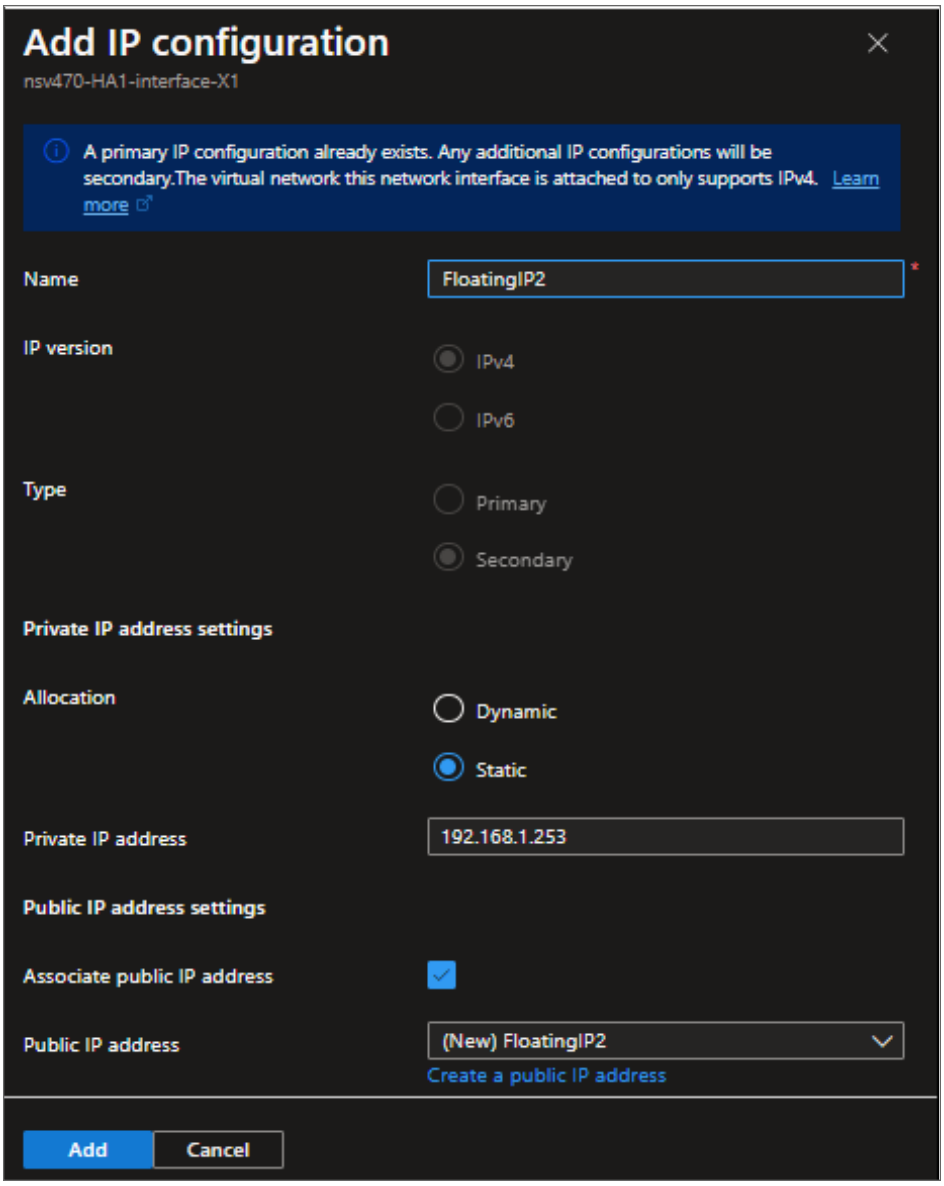

After you add the IP address in Azure, configure it on the firewall, for example, in a NAT policy. See the SonicOS NAT documentation.

6

Fine Tuning High Availability

Topics:

- Advanced Settings
- Configuring Advanced High Availability Settings

Advanced Settings

The **Advanced Settings** provides the ability to fine-tune the High Availability configuration as well as synchronize settings and firmware between the High Availability firewalls .

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 and Active/Active DPI Prerequisites

Configuring Advanced High Availability Settings

To configure advanced settings:

1. Log in as an administrator to the SonicOS Management Interface on the Active firewall.
   ⓘ **NOTE:** You can configure the settings only on the Active firewall. The Standby firewall operates in read-only mode.
2. Navigate to **DEVICE | High Availability > Advanced**.

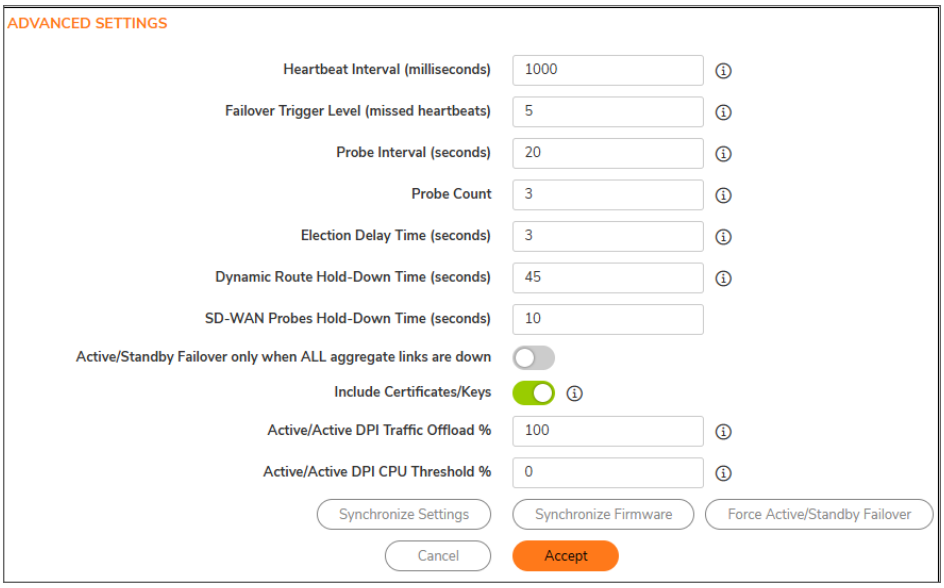

3. In the **Heartbeat Interval (milliseconds)** field, specify how frequently the firewalls in the High Availability pair exchange heartbeat messages. This setting applies across both units in the HA pair and helps determine peer health and responsiveness.

    - **Default:** 1,000 milliseconds (1 second)
    - **Range:** 1,000–300,000 milliseconds

    NOTE: .Use higher intervals in high-traffic environments to reduce processing overhead. Lower intervals increase sensitivity and may cause unnecessary failovers during periods of temporary network congestion. This timer works in conjunction with the Failover Trigger Level (missed heartbeats).

4. In the **Failover Trigger Level (missed heartbeats)** field, define how many consecutive heartbeats can be missed before initiating a failover. This setting applies to both units in the HA pair.

    This setting works together with the Heartbeat Interval timer. For example, if the **Failover Trigger Level** is set to 5 and the **Heartbeat Interval** is 10,000 milliseconds (10 seconds), the system waits 50 seconds without receiving a heartbeat before triggering failover.

    - **Default:** 5
    - **Range:** 4–99

5. In the **Probe Interval (seconds)** field, enter the number of seconds between probes sent to the specified IP address(es) to verify that the network's critical path remains reachable. Logical monitoring for the HA pair uses this interval.

    - **Default:** 20 seconds
    - **Range:** 5–255 seconds

6. In the **Probe Count** field, specify how many consecutive missed probes must occur before SonicOS determines that the critical path or probe target is unreachable. Logical monitoring uses this count for HA failover decisions.

- **Default:** 3
- **Range:** 3–10

7. In the **Election Delay Time (seconds)** field, specify how long the Active firewall waits before considering an interface “up” and stable. It is useful when switch ports have spanning-tree delays.
   - **Default:** 3 seconds
   - **Range:** 3–255 seconds
8. In the **Dynamic Route Hold-Down Time (seconds)** field, enter the duration for which the new Active firewall retains old dynamic routing entries after failover. This allows time to relearn routes from dynamic routing protocols (RIP, OSPF, BGP) before replacing old entries.
   - **Default:** 45 seconds
   - **Range:** 0–1200 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 uses a dynamic routing protocol. During this time, the new Active firewall 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.
9. In the **SD-WAN Probes Hold-Down Time (seconds)** field, specify how long SonicOS waits before it restores a previously failed WAN path to the active state.
   - **Default:** 30 seconds
   - **Range:** 0–3600 seconds
10. Select **Active/Standby Failover only when ALL aggregate links are down** to fail over only when every aggregate link is down. An aggregate link is a group of physical interfaces that SonicOS combines as one logical link. This option is not selected by default.
11. Select **Include Certificates/Keys** to synchronize all certificates and keys within the HA pair. This option is selected by default.

    ⓘ **NOTE:** Synchronization copies private keys from one firewall to the other. Use a secure connection between the firewalls for the HA control and data links.
12. In the **Active/Active DPI Traffic Offload %** field, enter the percentage of DPI-eligible traffic that the Active firewall forwards to the Standby firewall when the CPU threshold is met.
    - **Default:** 60%
    - **Range:** 0–100%

    For example, a value of 60% specifies that 60% of eligible DPI traffic is to be offloaded after the CPU threshold is reached.

    ⓘ **NOTE:** This setting applies to DPI-eligible traffic, not to all traffic passing through the firewall. A value of 100% does not mean that all firewall traffic or all firewall processing is transferred to the Standby firewall.

13. In the **Active/Active DPI CPU Threshold %** field, enter the CPU utilization level on the Active firewall that triggers DPI offloading.

    - **Default:** 60%
    - **Range:** 0–100%

    When CPU utilization reaches or exceeds the configured threshold, the Active firewall begins forwarding the percentage of eligible DPI traffic specified in **Active/Active DPI Traffic Offload %** to the Standby firewall.

    Setting the CPU threshold to **0%** allows offloading to begin without waiting for CPU utilization to increase.

    For example, setting **Active/Active DPI Traffic Offload %** to **100%** and **Active/Active DPI CPU Threshold %** to **0%** allows eligible DPI traffic to be offloaded as soon as it is available.

    ⓘ **NOTE:** .These settings control inspection offloading. They do not change which appliance forwards traffic between the LAN and WAN.

14. (Optional) To force synchronize the SonicOS preference settings between the Primary and Secondary firewalls, click **Synchronize Settings**.
15. (Optional) Click **Synchronize Firmware** to synchronize the firmware version between the Primary and Secondary firewalls.
16. (Optional) Click **Force Active/Standby Failover** to test that failover works. The firewalls switch roles, which briefly interrupts traffic.

    △ **CAUTION: Run this test during a maintenance window. The failover interrupts traffic.**
17. Click **Accept** to complete the High Availability configuration.
    SonicOS synchronizes all settings to the Secondary firewall.

7

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. The Active/Active DPI offloading test uses the CLI of the Standby firewall and the **DEVICE | High Availability > Advanced** page.

Topics:

- Configuring Active/Standby High Availability Monitoring
- IPv6 High Availability Monitoring
- IPv6 HA Monitoring Considerations
- Verifying Active/Active DPI Traffic Offloading

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 firewall.
2. Navigate to **DEVICE | High Availability > Monitoring**.

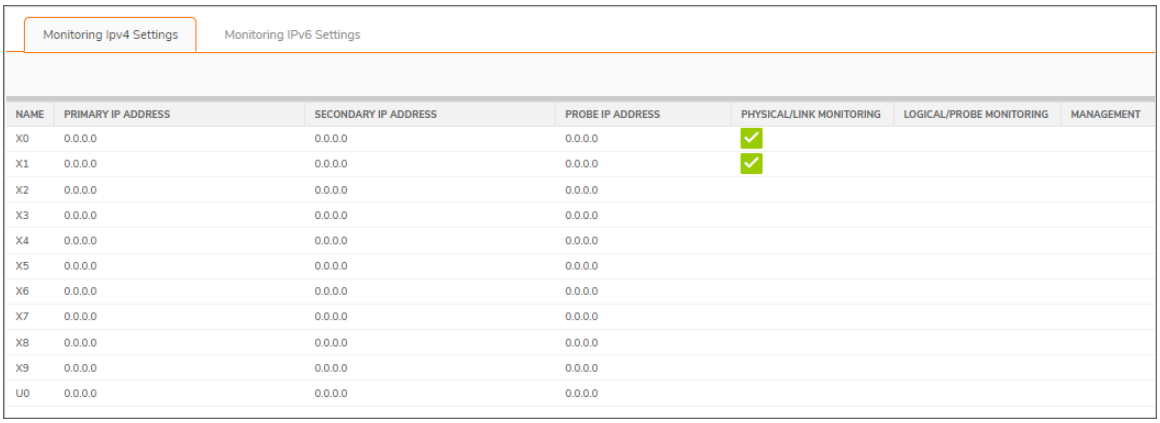

3. Click the **Edit** icon for an interface on the LAN, such as X0. The **Interface Monitoring Settings** dialog is displayed.
4. To enable link detection between the designated HA interfaces on the Primary and Secondary units, leave **Physical/Link Monitoring** selected. This option is selected by default.
5. In the Primary IPv4/v6 Address field, enter the unique LAN management IP address of the Primary unit. This address must be in the same subnet as the 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. This address must be in the same subnet as the 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 firewalls regularly ping this probe IP address. If both firewalls ping the target successfully, no failover occurs. If neither firewall pings the target successfully, no failover occurs, because SonicOS assumes that the problem is with the target. If one firewall can ping the target and the other cannot, failover occurs to the firewall that can ping the target. For more information, see About HA Monitoring.

   The Primary IPv4/v6 Address and Secondary IPv4/v6 Address fields must be configured with independent IP addresses on 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 reduces 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 Secondary firewalls 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. Select IPv6 and refer to IPv6 HA Monitoring Considerations for configuration details. For the IPv4 procedure, see Configuring Active/Standby High Availability Monitoring.

IPv6 HA Monitoring Considerations

Consider the following when configuring IPv6 HA Monitoring:

- In the **Interface Monitoring Settings** dialog box, the **Physical/Link Monitoring** and **Override Virtual MAC** options 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 as the interface, and it cannot be the same as the global IPv6 address or the link-local address of the Primary or Secondary firewall.
- If the monitoring IP address is set to anything other than ::, the two addresses 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.

Verifying Active/Active DPI Traffic Offloading

Use the CLI of the Standby firewall to verify that it processes offloaded DPI traffic. This test uses the **DEVICE | High Availability > Advanced** page, which differs from the monitoring settings on this page.

To verify Active/Active DPI offloading:

1. Verify that the **HA Mode** field on **DEVICE | High Availability > Status** shows **Active/Active DPI** and that the Active/Active DPI link is connected.
2. Verify that Gateway Anti-Virus, Gateway Anti-Spyware, IPS, Application Control, or Capture ATP is enabled and applied to the test traffic. For an HTTPS test, enable DPI-SSL. For an SSH test, enable DPI-SSH.
3. On the active appliance, navigate to **DEVICE | High Availability > Advanced** and set:
   a. **Active/Active DPI Traffic Offload %**: 100%
   b. **Active/Active DPI CPU Threshold %**: 0%
4. Click **Accept**.

   CAUTION: These settings send all eligible DPI traffic to the Standby firewall and can increase its load. Run this test during a maintenance window.
5. Access the CLI of the Standby firewall. In the normal Primary Active configuration, this is the Secondary firewall. Use the management IP address of the Secondary firewall, which you set on **DEVICE | High Availability > Monitoring**.
6. Run the following command to establish a baseline. Note the DPI packet-processing counters in the output.

   ```
   show tech-support-report active-utm
   ```

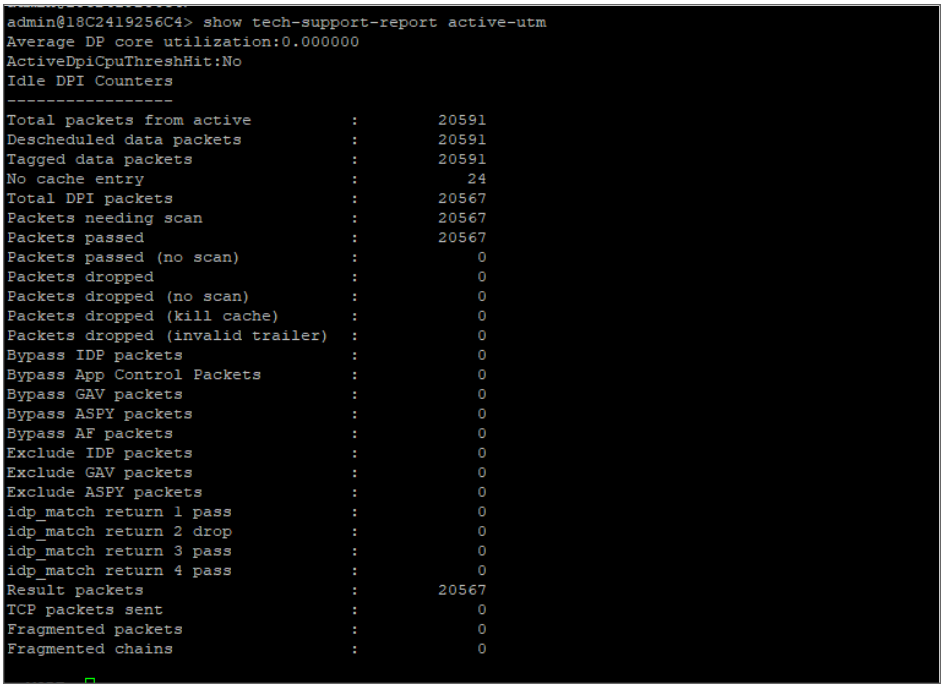

7. Generate eligible HTTPS or SSH traffic from a LAN client to a WAN destination. For example, download a large file (about 100 MB or more) over HTTPS. Eligible traffic is TCP traffic that uses a supported protocol

and that a DPI security service inspects.

8. While the traffic flows, run the command again. Compare the DPI packet-processing counters with the baseline. Increased counters confirm that the Standby firewall handles DPI packets.
9. After the test, restore the offload percentage and CPU threshold to the values that the deployment requires. The defaults are 60% for both settings. Click **Accept** to save the values.

   NOTE: Verify offloading with the reported DPI packet-processing activity. Do not use the dashboard bandwidth reading of the Standby firewall as the only verification method.

8

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

Use Case 1: Manage Azure HA Firewall

The following table lists the IP address names that this guide uses.

Name in Azure	Name in SonicOS	Purpose
Floating IP	Interface IP	Manages the Active firewall. The address moves to the new Active firewall after a failover.
Primary IP	Monitor IP	Manages each firewall independently, in the Active or Standby state.

Use the floating IP address to manage the Active firewall.

Use the two HA monitor IP addresses to manage each firewall independently.

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

In the Azure HA environment, the floating IP address (the interface IP on SonicOS) is the default gateway of the LAN subnet. Traffic from the LAN to the external network goes to the Active firewall, because the floating IP address moves with the Active role. After an HA failover, the former Active firewall becomes the Standby firewall, and traffic goes to the new Active firewall.

Use Case 3: Configure DNAT on Azure HA Firewall

This use case shows an example of Destination Network Address Translation (DNAT). The DNAT policy remaps traffic from the Internet that is sent to the floating IP address (interface IP) 192.168.1.254 to a specific machine (192.168.2.200) on the LAN. The following figures show the example before and after an HA failover.

THE FLOATING IP ADDRESS ALWAYS STAYS ON THE ACTIVE FIREWALL

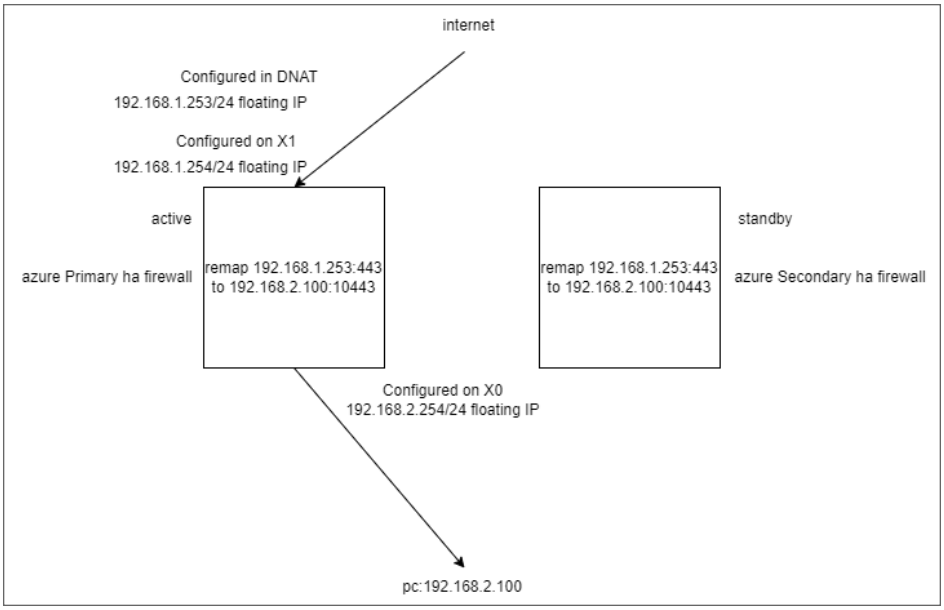

AFTER THE HA FAILOVER

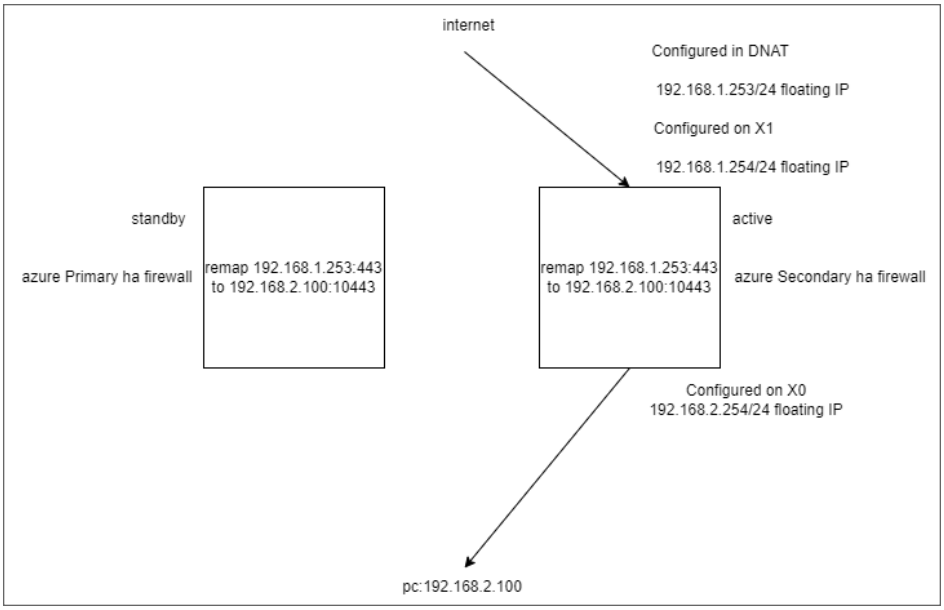

Use Case 4: Configure DNAT on Azure HA Firewall

This use case shows an example of Destination Network Address Translation (DNAT) that uses an additional floating IP address. The DNAT policy remaps traffic from the Internet that is sent to the additional floating IP address 192.168.1.253 (not the interface IP) to a specific machine (192.168.2.100) on the LAN. The following figures show the example before and after an HA failover.

THE FLOATING IP ADDRESS ALWAYS STAYS ON THE ACTIVE FIREWALL

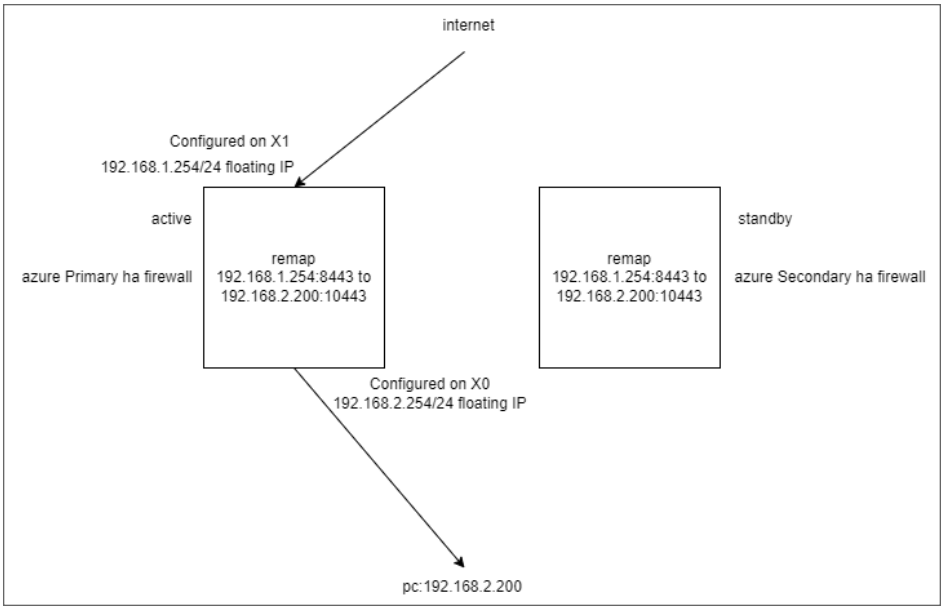

AFTER THE HA FAILOVER

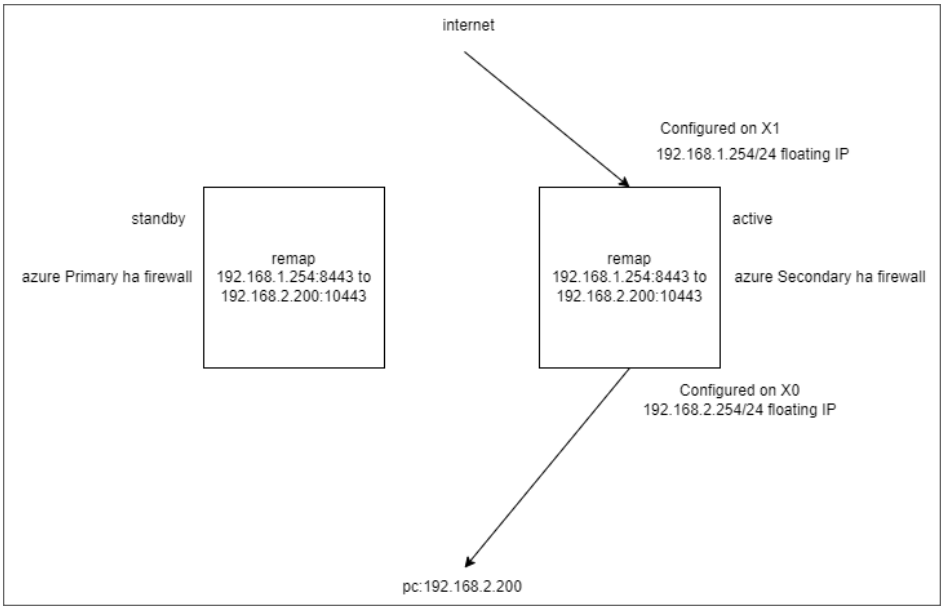

9

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 Technical Certifications
    - For Training and Assets

About This Document

SonicOS High Availability Administration Guide
Updated - October 2026
Software Version - 8
232-006186-00 Rev C

Copyright © 2026 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.

Use of Artificial Intelligence

Portions of this document were developed with artificial intelligence (AI) assistance. All AI-assisted content has been reviewed, verified, and modified as needed by qualified SonicWall technical staff to ensure accuracy and completeness. SonicWall assumes full responsibility for the accuracy of this document.

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