SBC Redundancy

In this section:

Related articles:

- Link Detection Group - Link Monitor - EMA
- Link Detection Group - CLI
- Admin - CLI
- Enabling SBC Core Geographical Redundancy HA Mode

Platform Redundancy

The SBC 5000/7000 series platforms support 1:1 box level redundancy through the communication paths provided by the external High Availability (HA) Ethernet ports.

Types of platform redundancy include:

- Server—The Server modules provide 1:1 redundancy with automatic switchover and no interruption to stable calls.
- Power—The SBC has no common power supplies; each server module is connected to both redundant AC or DC power feeds and handles its own DC conversion.
- Cooling—The SBC 5000 series has three fan modules and the SBC 7000 four fan modules. The SBC can operate indefinitely if one fan module fails.

For more information on installing the HA pairs, see Installing SBC Application.

Geographical Redundancy High Availability (GRHA)

The SBC Core includes a CLI action command to support Geographical Redundancy High Availability (GRHA) mode providing enhanced support against network degradation. GRHA mode supports active and standby servers located in two different data-centers. This is performed by changing the bond device monitoring from MII to ARP. ARP monitoring is used to detect issues when SBCs are connected to switches, which are connected to each other through a network. GRHA mode protects against data center and network failures. In addition to switching the bond monitoring, the leadership algorithm changes the decision on which SBC survives split-brain recovery. The change is included in a GRHA situation since any HA link loss is primarily due to data center isolation.

This feature is not configurable during software installation nor is it changeable during an upgrade.

Bond Link Monitoring

The SBC 5000/7000 series platforms support Bond Monitoring which is configurable using the CLI `setHaConfig bondMonitoring` command (described below) to change the bond device monitoring from MII to ARP. ARP monitoring is used to detect issues when SBCs are connected to switches which are connected to each other through a network. GRHA mode protects against data center and network failures.

Note
Bond link monitoring is not applicable to the SBC SWe platform.

Enhanced Leadership Algorithm

When network issues prevent the SBCs from communicating, both nodes become active (split-brain recovery). Once the communication is re-established, one of the nodes must be restarted again to become the standby. The SBC Core supports an enhanced leadership algorithm which is configurable using the CLI `setHaConfig leaderElection` command. This algorithm changes the decision on which the SBC survives split-brain recovery. This functionality is included in a GRHA situation since any HA link loss is primarily due to data center isolation.

When network issues prevent the SBCs from communicating, both nodes become active (split-brain recovery). Once the communication is reestablished, one of the nodes must be restarted again to become the standby. The enhanced Leadership Algorithm changes how the SBC decides which node will survive as active by choosing the node that was promoted to active during the split-brain recovery. The algorithm also performs additional checks to handle situations where a node may restart or start for the first time while communication is interrupted.

> **Note**
> All nodes must use the same algorithm. The default algorithm is used until new peer election algorithm is configured and available on both the nodes.

HA Network Requirements

GRHA mode network requirements are listed below:

- 10G/1G bandwidth for direct and not direct cable connections.
- The network connecting to the SBCs must support ARP monitoring.

HA Link Delays per HA Pair

The SBC HA servers can be deployed to different geographical locations to meet various disaster recovery requirements. The following table lists High Availability link delays per HA pair for each SBC Core platform.

Table 1: SBC Core HA Link Delays

SBC Platform	Maximum HA Link Round-Trip Delay (ms)	Call Rate (cps)	Call Capacity	Delay After Switchover	
				Mgmt Connectivity Restoration Time (mm:ss)	HA Restoration Time (mm:ss)
51x0	110	150	10,000	00:35	05:45
52x0	60	450	64,000	00:45	08:55
5400					
7000	20	1350	150,000	00:28	05:25
SWe	110	25	4,250	00:20	06:56

Non-Disruptive Fail-Over

In a High Availability configuration with an active and standby (redundant) server where active server fails, switch-over is completely automatic preserving the integrity of stable calls.

A switch-over from active to redundant server can result in packet loss. Fax (and modem) calls are generally not tolerant to media interruptions. Despite the fact that some fax and modem calls may be preserved during a switch-over, it is not uncommon for fax machines and modems to terminate their transmissions as a result of a server switch-over.

SBC 7000 Port Redundancy and Link Detection

Port Redundancy Model

Each SBC 7000 supports two primary (active) and two secondary (standby) 10 GigE media interfaces (packet ports). The standby port functionality provides redundant port protection for each of the active media interfaces. In an HA scenario, the backup CE has its own primary and secondary packet ports. See Figure 1 for a depiction of the HA port redundancy configuration.

For a depiction of media port inter-connectivity in an HA configuration, see the Management and HA Port Connections diagram on page Connecting SBC 7000 Series Ethernet and Data Cables.

Figure 1: SBC 7000 HA Port Redundancy Model

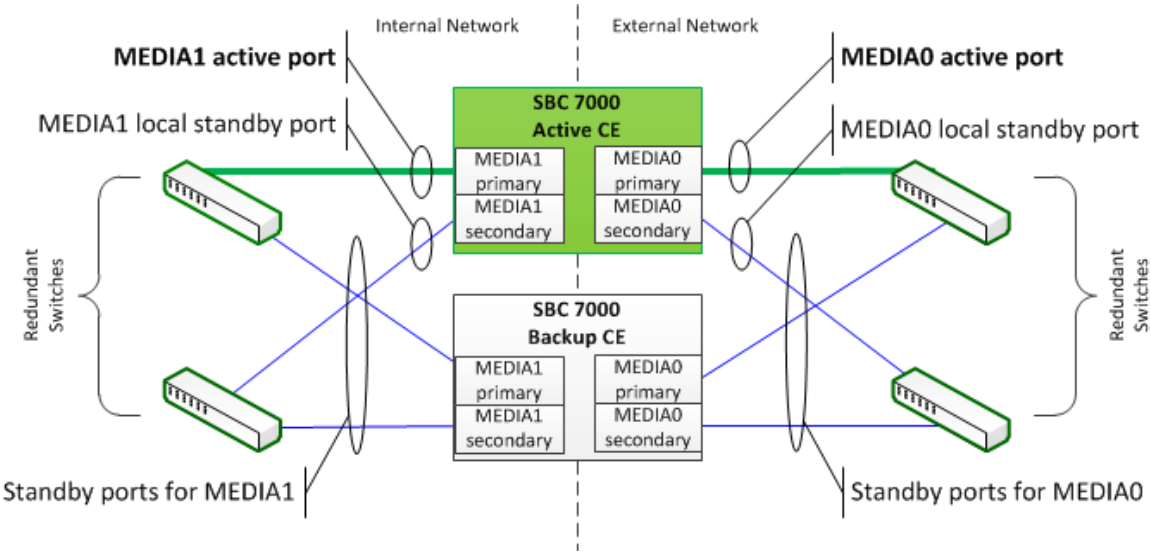

Terminology:

- **Primary port:** An Ethernet packet port that will attempt to become active if on an active CE. All four packet ports on the 52x0, the two packet ports on a 51x0 and two dedicated packet ports on the 7000 (MEDIA 0_P, MEDIA 1_P) are considered primary ports.
- **Secondary port:** An Ethernet packet port designated as an alternate for a specific on-board primary port. The SBC 7000 contains one secondary port for each 10 GigE primary port (SBC 5xx0 servers do not have secondary ports). The primary and secondary port roles are static and not modifiable by the user.
- **Active port:** An Ethernet packet port that is currently selected for use (e.g. for signaling, media, etc.); either a primary or secondary port on an active CE.

  ⚠ A port which is in the active state does not necessarily imply that is "up".

- **Local standby port:** A standby port on an active SBC 7000 CE providing redundancy protection to the currently active port.
- **Standby port:** A collective term for a local standby port on an active CE or any packet port on an inactive CE. Standby ports can provide protection for active ports.
- **Enabled or Disabled ports:** A packet port may be administratively enabled or disabled. A port that is disabled cannot be an active port. Packet ports on an inactive CE do not have their own distinct administrative state. They share this configured element with their counterpart on the active CE.

⚠ A port's role (Primary/Secondary) is independent of the port's state (Active/Standby).

Link Detection Support

The SBC Core supports the capability to perform link detection on standby and active Ethernet ports to facilitate determining the health of standby port before initiating a switchover/failover. The intent is to allow simple connectivity checking to test the ability of SBC to send/receive Ethernet frames, connectivity to the adjacent switch/router, and the ability of the switches/router to do basic layer 2 receiving/forwarding/sending.

The following probing mechanisms are available on the SBC platforms:

Table 2: SBC Probing Mechanism Types

Probing Mechanism	SBC Platforms	Affected Ports	Purpose
Physical link detection	All	All ports (active/standby CEs)	• Detects the presence of the port cable and that the adjacent device is powered on (enabled by default on all physical ports configured by Link Monitor except for any ports administratively disabled or set out-of-service). • Reports any hardware failures to NRS and Link Verification Manager (LVM) tasks.
ICMP ping	All	Active ports only	Checks two-way connectivity between the SBC port and the configured destination (adjacent router) by sending ICMP Ping messages at configured intervals to the destination. ⚠ When destination IP address is configured in a Link Monitor, ICMP Ping is enabled along with physical link detection. By setting the destination IP address to NULL (0.0.0.0), only physical link detection can be enabled.
ARP ACD/ICMPv6 NUD*	SBC 7000 only	Standby ports only	Verifies physical media by checking two-way traffic through at least the local Ethernet interface, the cable, and the adjacent layer 2/3 switching function. Layer 3 verification is accomplished using ARP ACD Probes (for IPv4) or Neighbor Discovery (for IPv6) mechanisms to probe an arbitrary, operator-specified target IP address on a local IP subnet, typically the address of the next-hop router (Gateway IP address). Depending on the address family (IPv4/IPv6) of the gateway IP address configured, either ARP ACD or ICMPv6 NUD probing messages are sent in such a way that explicit assignments of IP addresses to the standby ports are not required. See below for specifics on IPv4 ARP ACD requirements. When IP Target is to 0.0.0.0 and/or “probeOnStandby” is disabled, only the physical link state between the active/standby SBC port and the adjacent router is monitored.

* Address Resolution Protocol - Address Conflict Detection / Internet Control Message Protocol Version 6 – Neighbor Unreachability Detection

ARP ACD/ICMPv6 NUD Methods for Standby Ports

IPv4 ARP ACD

If the destination address configured is an IPv4 address, then IPv4 probing is initiated by sending ARP Probe requests and listening for the responses.

ARP Request probes are sent with:

- **Sender IP address** of 0.0.0.0. The use of 0.0.0.0 is compatible with rfc 5227 on IPv4 Address Conflict Detection. This is convenient to use on standby ports since IP addresses are not assigned for standby ports.
- **Sender hardware address** containing the current local MAC address assigned to the sending port.
- **Target IP address** containing the configured target IP address to be probed.
- **Target hardware address** containing all zeros. The ARP request is sent on the LAN using L2 broadcast.

The target is required to respond to the ARP probe with an ARP Response having an L2 unicast MAC as the DESTINATION and SOURCE. If the target replies with a GARP or ARP request in the form of a broadcast, the SBC drops these requests due to DDOS functionality enabled in the application code.

Refer to Link Detection Group - CLI for command to disable probe functionality on the Standby port if router can not reply to the ARP probe with a

unicast destination MAC address.

IPv6 ICMPv6 NUD

If the destination address configured is an IPv6 address, then IPv6 probing would be initiated using Neighbor Unreachability Detection mechanism (RFC 4861 section 7). This is based on Neighbor Solicitation and Neighbor Advertisement ICMPv6 messages.

Because these are IP packets, the SBC needs IP addresses to send/receive them. The SBC uses auto-generated link local IPv6 address from the current local MAC address.

Neighbor Solicitation messages are sent with:

- **IP source address** containing auto-generated link local IP address
- **IP destination address** containing configured target IP address
- **ICMP layer target address** containing configured target IP address
- **ICMP layer source link layer address** This field is left blank to prevent the target from learning our L2: L3 address binding from these probes.

The Neighbor Solicitation message is sent on the LAN via L2 unicast to the system with the target IP address.

The target can be expected to respond with a Neighbor Advertisement using L2 unicast. Received messages are validated per RFC 4861 section 7.1.2: Check that the S bit = 1 (solicited) and that the target address = our configured target IP address.

> ⚠ The SBC 7000 may reduce the call accept rate when syncing from the active to the standby CE under full load causing some calls to get rejected with a 503 message even when the applied load is below the specified maximum call rate. This condition clears once the synchronization to the standby completes. Additionally, some calls may get rejected with a 503 message when synchronization occurs while the applied load is near the maximum specified.

The impact of a link or switch failure on SBC 7000 is depicted in the diagram below.

Figure 2: SBC 7000 HA Port Status After a Link or Switch Failure

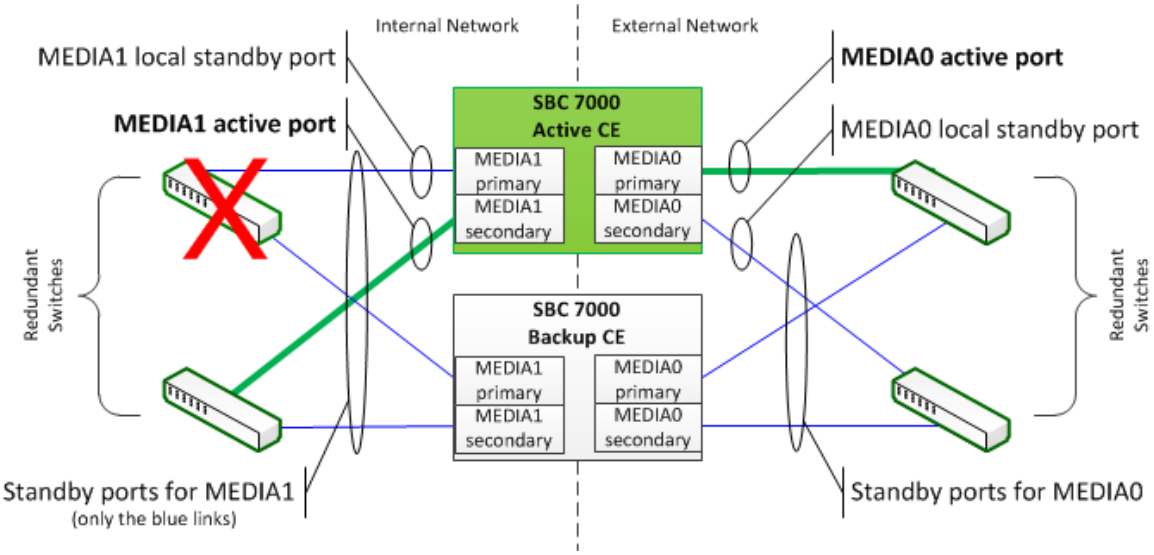

ARP Implementation

Most switches, in their default behavior, forward ARP Probes without issue; however, if a switch has ARP inspection/filtering functionality enabled, that feature must not discard RFC5227 ARP Probes as “invalid” or it cannot be configured on ports connected to SBC 7000 Media ports.

The SBC 7000 Standby ports do not have IP addresses while in standby mode, so they cannot generate ICMP Echo Requests for Sonus Link Detection purposes. If configured to provide a similar logical connectivity check, the SBC 7000 standby ports instead send an ARP Probe to the target IP address (see Address Context - Link Detection Group (EMA) or Link Detection Group - CLI for details with configuring Link Detection).

ARP Probes are forwarded like any other traffic by most switches in their default behavior. Some switches have features (for example, Dynamic ARP Inspection, Dynamic ARP Protection, etc.) that, when enabled, discard “invalid” ARP packets. These features may incorrectly consider ARP

Probes per RFC 5227 to be “invalid” ARP packets. If such a feature is enabled, the feature must not discard RFC5227 ARP Probes as “invalid” or it cannot be configured on the switch ports connected to SBC 7000 Media ports.

Disabling ARP Probing on Standby Ports

The SBC 7000 includes the flag, probeOnStandby, for use in disabling ARP/NUD probing by Link Monitors on standby packet ports in case routers in your network do not respond correctly to the ARP probes. This scenario can lead to Link Monitor declaring itself as failed. The CLI syntax is shown below (default value of `probeOnStandby` is 'enabled').

> ⚠ Disabling ARP/NUD probing can possibly lead to a toggling situation since we rely only on the physical port health on the standby packet port.

When the `probeOnStandby` flag is disabled, there is a possibility for a toggling situation when the Link Monitor on the active port detects a failure via ICMP ping to a destination, while on the standby packet port it can only use the physical health of the port. Therefore, once the port becomes standby, it can look healthy if the physical port is up; however, when it becomes active and fails to reach the configured destination, it looks unhealthy.

> ⚠ The Port Redundancy Group includes a mechanism to detect a scenario where link failures begin rapidly toggling between active and standby packet port. If this scenario occurs, packet port redundancy continues for physical port failures, but not for link failures reported by the Link Monitors.

```
% set addressContext <addressContext_name> linkDetectionGroup <LDG_name> linkMonitor <name>
probeOnStandby <disabled | enabled>
```

SBC 7000 Redundancy Performance

The different aspects of SBC 7000 redundancy performance depends on the size of the configuration. The following configuration profiles are defined to reduce the number of test combinations.

- 1K profile – Comprises 1,000 instances each of address contexts, zones, interface groups, interfaces, signaling ports, and trunk groups.
- 4K profile – Comprises 4,000 instances each of address contexts, zones, interface groups, interfaces, signaling ports, and trunk groups.
- 4/40K profile – Comprises 4,000 instances each of address contexts, zones, interface groups, interfaces and signaling ports; and 40,000 trunk groups.

HA Restoration

The SBC 7000 platform supports 1:1 box level redundancy. The full HA protection can be restored after switchover and virgin standby start states.

Table 3: SBC 7000 Full HA After Switchover State

Configuration Profile	Maximum Time
1K Profile	10 minutes
4K Profile	35 minutes
4/40K Profile	60 minutes

Table 4: SBC 7000 Full HA After Virgin Standby Start

Configuration profile	Maximum Time

1K Profile	90 minutes
4K Profile	120 minutes
4/40K Profile	150 minutes

The above time to HA protection only applies when replacing a failed SBC 7000 chassis with completely new hardware.

HA Connectivity Requirements

To meet the redundancy performance requirements, HA connectivity between the active and standby nodes in a SBC 7000 HA pair must meet certain delay and packet loss metrics. These metrics are the same as for the SBC 5000 series platform and are summarized in the following table.

Table 5: SBC 5000/7000 Packet Loss Percentage per HA Pair Delay

Delay (ms)	Loss (%)
5	1
10	0.3
30	0.2
50	0.1
75	0.001
100	0.0005

Media Signaling and Management Recovery

The SBC 7000 platform uses health-checking and hot-standby techniques to efficiently detect faults and to recover media, signaling, and management connectivity with minimal external effect regardless of the current rate or capacity loading of the system.

Signaling

After a fault is detected and the system switches over, the behavior of the SBC 7000 with respect to call and registration signaling is as follows:

- Call and registration signaling are accepted and responded to within five seconds of switch-over. Note that this does not indicate that calls will be accepted at this time. However, within this time window they are cleanly rejected so upstream nodes are not unduly burdened with retransmissions, and are then able to quickly re-route the calls to alternate servers.
- Most new calls and registrations are accepted within 15 seconds of switch-over. Not all calls may be accepted at this point. Depending on the load at the time of the switch-over, the box may control congestion by throttling the call accepts until the switch-over cleanup is complete.
- Full call and registration acceptance is achieved within two minutes of switch-over.

Management Interfaces

Management interfaces are available within two seconds after a switch-over. Specifically, it is possible to login through the management interfaces on the activated standby within two seconds.

Media Recovery

The media recovery time depends on the failure mode and whether the calls are pass-through or transcoded. The following table shows the worst-case SBC 7000 media recovery time for different conditions:

Table 6: SBC 7000 Media Recovery Time

Condition	Media Recovery Time	
	No Transcoding (ms)	Transcoding Enabled (ms)
Software Failure	50	700
Operator Initiated Switchover	50	700
Packet Port Failure	250	900
Reset Active System via CLI/EMA	50	700
Hard Power Failure (power-off)	1200	1800

The SBC 7000 may reduce the call acceptance rate when syncing from the active to the standby CE under full load causing some calls to get rejected with a 503 message even when the applied load is below the specified maximum call rate. This condition clears once the synchronization to the standby completes. Additionally, some calls may get rejected with a 503 message when synchronization occurs while the applied load is near the maximum specified.

Port Switchover Latency After Hardware or Link Detection Failure

The following table lists the maximum time for the SBC 7000 to fail over to Local-Standby Port after a hardware port or link detection failure.

Table 7: Approximate Time for Fail-over to Local-Standby Port After a Failure

Failure Type	Max Time to Fail Over to Local-Standby Port (ms)
Hardware port failure	200
Link detection failure	250