

High Availability Application Note

Date: May 2013

Edgewater Networks, Inc.
2895 Northwestern Parkway
Santa Clara, CA 95051
USA
http://www.edgewaternetworks.com
Phone: 408 351-7200
Fax: 408 727-6430

© 2013, Edgewater Networks, Inc.

Edgewater Networks Confidential, All Rights Reserved.

Trademarks

EDGEWATER NETWORKS and Design, EDGEMARC, EDGEPROTECT, and EDGECONNECT are registered trademarks of Edgewater Networks, Inc. EDGEWATER NETWORKS and EDGEVIEW are trademarks of Edgewater Networks, Inc. Any other trademarks appearing in this manual are owned by their respective companies.

Copyright

This document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this document may be reproduced in any form by any means without prior written authorization of Edgewater Networks, Inc. Documentation is provided “as is” without warranty of any kind, either express or implied, including any kind of implied or express warranty of non-infringement of the implied warranties of merchantability or fitness for a particular purpose.

Export Notice

You (Purchaser) acknowledge that these products are subject to the U.S. Export Administration Regulations (the “EAR”) and that you will comply with the EAR. You are not located in Cuba, Iran, North Korea, Sudan, and Syria. You will not export or re-export this product, directly or indirectly, to: (1) any countries that are subject to US export restrictions (currently including, but not necessarily limited to, Cuba, Iran, North Korea, Sudan, and Syria); (2) any end user who you know or have reason to know will utilize them in the design, development or production of nuclear, chemical or biological weapons, or rocket systems, space launch vehicles, and sounding rockets, or unmanned air vehicle systems; or (3) any end user who has been prohibited from participating in US export transactions by any federal agency of the US government. In addition, you are responsible for complying with any local laws in your jurisdiction which may impact your right to import, export or use these products.

Licensing

Use of this product is subject to Edgewater Networks Software License Agreement.

Portions of this product include software sponsored by the Free Software Foundation and are covered by the GNU GENERAL PUBLIC LICENSE.

See the www.edgewaternetworks.com/licensing for more information regarding licenses.

Warranty

Product warranty does not apply to damage caused by lightning, power surges, or wrong voltage usage.

Typographical Errors

This publication could include technical inaccuracies or typographical errors, for which Edgewater Networks never can or shall be held liable. Changes are made periodically to the information herein; these changes will be incorporated in new releases of this publication. Edgewater Networks may make improvements or changes in the product or products described in this publication at any time, without prior notice.

Document Release Date

5/21/13

Contents

High Availability Application Note 4
Logical Designation and Naming Conventions Used in the This Document 4
Revertive Mode and Standard Operation (Non-Revertive) Mode 5
Revertive Mode 5
Standard Operation (Non-Revertive) 6
Network Topology Expectations 6
Status Transfer and State Exchange Messages 9
High Availability Operational States 10
Administratively Disabled Mode 11
Supported Functionalities and Known Limitations 11
4550/46xx Known Limitations 13
General High Availability Scenarios Either Unsupported or Not Tested/Certified 13
Support for HA State Exchange 14
Important Operational Guidelines 15
Compatibility 15
Configuration Changes on an HA Cluster 15
Typical Administrative Maintenance Procedure 15
High Availability Configuration 16
Tasks for Configuring High Availability 17
High Availability Debug and Diagnostics 23
High Availability SNMP Support 26

© 2013 Edgewater Networks, Inc. All rights reserved.

High Availability Application Note

The EdgeMarc VOS High Availability (HA) solution is designed to eliminate a single point of failure in a network configuration. High Availability uses the concept of a Primary and Secondary system. The High Availability solution accommodates outages and failures across several layers of the OSI reference model. These include:

- Physical (chassis or power failure)
- Data Link (Ethernet link failure)
- Network Routing (IP communication)
  - This does not constitute awareness of viability of upstream/downstream networking equipment past the directly connected switch or router (LAN and WAN)
- Application Layer Gateway
  - This does not constitute all ALG failures — only those that Edgewater Networks has empirical evidence of and those that cause a stoppage of the ALG process. In other words, the HA State Machine reads the process list and if it cannot find the pid, the ALG is restarted.

The High Availability solution described in this document references VOS firmware 11.6.5 and later. HA amplifies the EdgeMarc Site Survivability feature which accommodates failures across the upper layers of the OSI reference model. The High Availability Application supports on the following EdgeMarc hardware platforms:

- 5300LF2
- 46xx series platforms
- 4550
- 6400LF2
- 7300

Logical Designation and Naming Conventions Used in the This Document

High Availability (HA) Cluster	Two devices that are configured correctly for the High Availability feature.
High Availability (HA) Node	A single device configured as a part of an HA Cluster that is partnered with another device to form the HA Cluster. In later releases, this could include more than two devices, but currently is limited to a pair of devices. Also referred to as “HA Peer” or “HA Partner.”
Virtual IP (VIP)	The Layer 3 IP address that the High Availability cluster masquerades for the purpose of allowing LAN and WAN side elements to send requests and receive responses for a single IP address

© Edgewater Networks, Inc. All rights reserved.

Active	The node that is the acting agent for call control and responds to requests made on the VIP.
Standby	The node that is waiting to assume the role as the acting agent for call control. This node does not respond to requests made on the VIP.
Revertive Mode	The mode of operation where the configured Primary will always assume control of the HA cluster as the Active node.
Non Revertive Mode	The mode of operation where the last known HA node that has been instituted as Active as a result of a HA failover will remain the Active node.
Primary	In “Revertive” mode, the specified device that will attempt to assume Active status if all resources are available.
Secondary	In “Revertive” mode, the specified device that will maintain Standby status if the configured Primary is unavailable.
Failover	The action in an HA cluster where the currently Active unit experiences some degradation which causes the Standby to assume the Active role as part of normal State Exchange and status transfer. This could occur in both Revertive and NonRevertive mode.
Fallback	The action in an HA cluster that is configured for “Revertive” mode where the Active role will automatically “fall back” or be assumed by the configured Primary node, upon normalization of whatever state caused the configured Primary to relinquish the Active state previously, as part of the normal State Exchange and status transfer.



Note:

The Primary (P) can be configured to remain in Standby (P/S) mode or revert to Active (P/A) mode via HA cluster designation of Revertive or Non-Revertive mode.

Revertive Mode and Standard Operation (Non-Revertive) Mode

Revertive Mode

For the purposes of a “Revertive” cluster configuration, a node is designated (via configuration) as a Primary or Secondary. Consider the case where at startup, the Primary 5300LF2 (P) system is the Active peer (A); and the Secondary (S) system is the Standby peer (S). If the Primary (P/A) system fails because of a network or hardware failure, the Secondary (S/S) system will take over and become the new Active system (S/A). At this point, depending on the operational condition of the Primary, it might be the P/S.

Failover and fallback from the configured Primary to the configured Secondary (and vice versa) can occur under the scenarios below:

- P/A or S/A hardware platform fails
- P/A or S/A LAN link failure is detected (see caveat for 46xx and 4550 platforms)

© 2013 Edgewater Networks, Inc. All rights reserved.

- P/A or S/A WAN link failure is detected
- P/A or S/A is put into Administratively Disabled Mode
- Certain ALG failures
- Failure of the HA module to initialize normally in either the configured Primary or Secondary

Standard Operation (Non-Revertive)

In a Non-Revertive cluster, the node that starts and normalizes first will become Active regardless of whether it is the Primary or Secondary. Only if Revertive mode is enabled will the Primary always become Active. This applies to any event that could cause system restarts or reboots.

Failover and fallback from the Active to the Standby peer (and vice versa) can occur under the scenarios below:

- Active hardware platform fails
- Active LAN link failure is detected (see caveat for 46xx and 4550 platforms)
- Active WAN link failure is detected
- Active is put into Administratively Disabled Mode
- Certain ALG failures in the Active peer
- Failure of the HA module to initialize normally in the Active peer

Network Topology Expectations

Network Equipment Guidelines (refer to the diagrams below):

- It is advisable, but not required, to hard-code your network switching device Ethernet ports to a static setting (speed/duplex) for the connections between the EdgeMarc HA nodes and the switch ports. This is to avoid auto-negotiation mismatch states that have been observed with various 3rd party vendor switch models and versions.
  - See the appropriate product documentation for link configuration instructions.
  - If interfaces are hard-coded, ensure that each node's WAN Ethernet ports and LAN Ethernet ports have the identical speed/duplex setting.
- Power-over-Ethernet (PoE) SHOULD always be disabled on switch ports that the HA nodes are connected to.
- Spanning Tree Protocol implementations MUST not be configured on switch ports that the HA nodes are connected to.

HA peer network interfaces are expected to be on the same network, but node interfaces must be on different networks (usual EdgeMarc configuration). As an example

- HA Peer WANs are on 172.16.1.0/27

© 2013 Edgewater Networks, Inc. All rights reserved.

- HA Peer LANs are on 1.1.1.0/24
- Virtual IP addresses are assigned per network, so there needs to be at least three available IP addresses to apply to system configuration

WAN (HA Peer Network) Configuration

- Primary WAN: 172.16.1.2/27
- Secondary WAN: 172.16.1.3/27
- VIP WAN: 172.16.1.1

LAN (HA Peer Network) Configuration

- Primary LAN: 1.1.1.2/24
- Secondary LAN: 1.1.1.3/24
- VIP LAN: 1.1.1.1/24

Figure 1 illustrates the typical topology support for correct High Availability functionality.

Figure 1. **High Availability Topology Expectations**

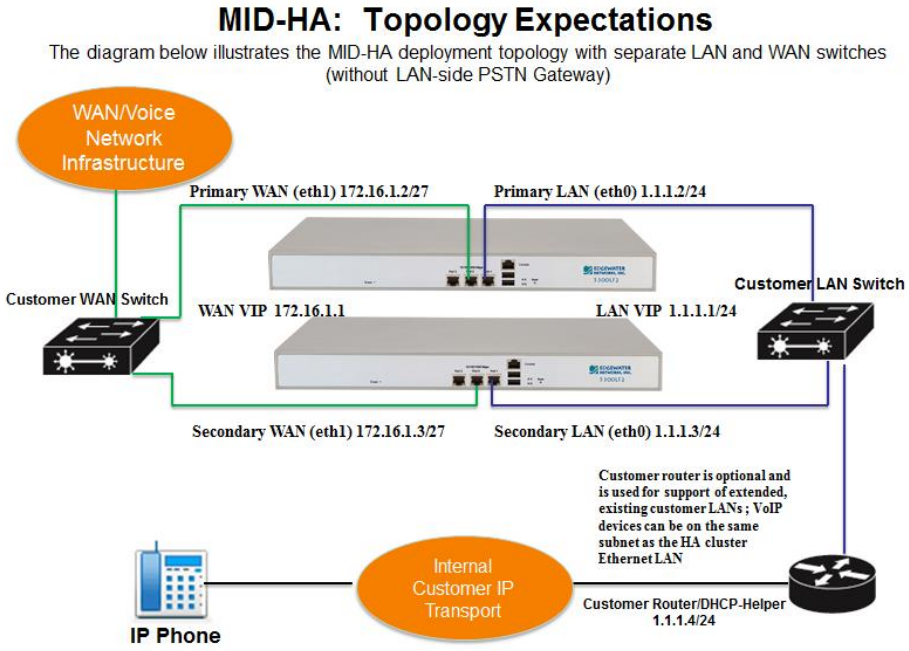


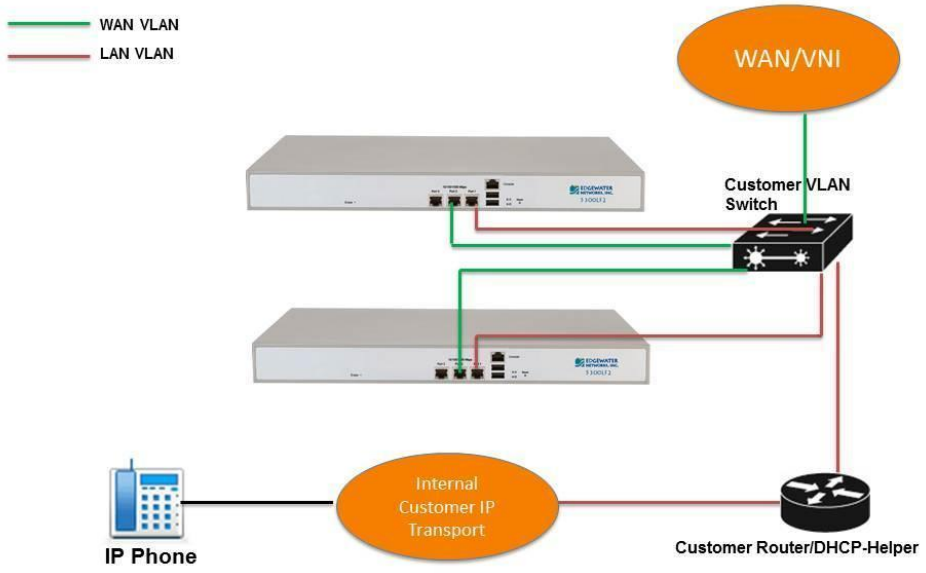
Figure 2 assumes the same IP addressing scheme, but with the use of VLANs for the purpose of segmenting networks via a robust switching fabric.

© 2013 Edgewater Networks, Inc. All rights reserved.

Figure 2. **High Availability Topology Expectations**

Figure 2: High Availability Topology Expectations

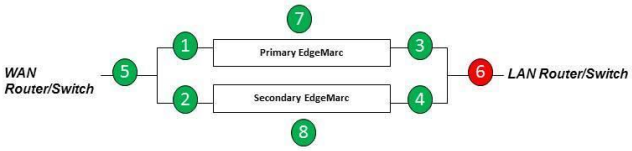
The diagram below illustrates the MID-HA deployment topology with VLANs (without LAN-side PSTN Gateway)



2

Figure 3 illustrates logical functionality and failure points and conditions with the 5300LF2 (the identification of failover states assumes that Revertive Mode is used).

Figure 3: Example Failure Points and Conditions for the 5300LF2



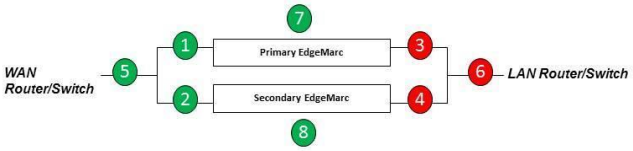
FAILURE POINT	RESULT
1	Secondary becomes 'Active' in remote control mode
2	Primary remains 'Active' in remote control mode
3	Secondary becomes 'Active' in remote control mode
4	Primary remains 'Active' in remote control mode
5	GR / Survivability settings apply
6	'url call disbled' or 'no service' condition for phones
7	Secondary becomes 'Active' in remote control mode
8	Primary remains 'Active' in remote control mode

3

Figure 4 illustrates logical functionality and failure points and conditions with the 4550 and 46xx series (the identification of failover states assumes that Revertive Mode is used).

© 2013 Edgewater Networks, Inc. All rights reserved.

Figure 4: Example Failure Points and Conditions for the 4608POE



FAILURE POINT	RESULT
1	Secondary becomes 'Active' in remote control mode
2	Primary remains 'Active' in remote control mode
3	'url call disbled' or 'no service' condition for phones
4	'url call disabled' or 'no service'' condition for phones
5	GR / Survivability settings apply
6	'url call disabled' or 'no service' condition for phones
7	Secondary becomes 'Active' in remote control mode
8	Primary remains 'Active' in remote control mode

4

Status Transfer and State Exchange Messages

High Availability is facilitated by intra-system communication using state transfer messages called the State Exchange Message (SEM). Each node within the High Availability cluster is configured to send State Exchange Messages (SEM) between the configured Primary (or Active) peer node to the configured Secondary (or Standby) node on a per-link basis (LAN, WAN, Management, if applicable). By using SEMs for all configured network links, you have a greater probability of successful SEM during link failures.

It is important that each node within the HA cluster is configured symmetrically. For example, if State Transfer is only enabled on the LAN interface but not the WAN interface of the Active System, and this link fails causing the Standby system to become once-Active (now Standby), the system will no longer be able to receive state updates. That means that if the link is restored (and the High Availability system is configured for “Revertive” behavior) the Primary system that is now in Standby status will assume the Active status without a full state transfer. To prevent this, configure each node symmetrically with State Transfer on the appropriate interfaces. There are additional caveats to State Transfer/State Exchange later in this document.

The following system configuration elements MUST BE symmetric across the HA peer nodes for correct functioning of HA:

© 2013 Edgewater Networks, Inc. All rights reserved.

- **License keys:** Each HA peer must have the same license key configuration.
- **ALG configuration:** All aspects of the ALG must be configured the same. This includes:
  - SIP Servers
  - Designation of transparent mode versus registrar mode
  - Certain hidden parameters that are used for AT&T's IPBE Geo-Redundancy feature; these are found in /etc/config/alg_defs.conf
  - SIP server list
  - Limit Inbound/Outbound Proxy checkboxes
- **Survivability configuration:**
  - Softswitch detection
  - Options Keepalives (if used)
  - Keepalive no-response backoff algorithm
  - SIP request monitor
  - SIP server redundancy
  - Registration control/rate pacing
- **Routing information:** (if downstream remote networks are used)
  - This concept does not affect High Availability functioning as much as it affects the user experience

Although the above list does not represent all permutations of configuration, it is especially important to ensure that there is no corruption of shared information between devices and also so that the expected post-failover behavior is the same among HA peers

High Availability Operational States

The High Availability application monitors certain configured resources and controls communication between the Active and Standby systems. If any of the configured resources on the Primary (P) are not available, the Secondary (S) becomes the Active via failover. The same is true for a fallback from Secondary to Primary.

The resources mentioned above include network resources (LAN interface, WAN interface, Management Interface), as well as the Application Layer Gateway process.

- If the ALG on the Active unit either manually or automatically restarts, the HA module will cause a failover. An automatic restart will occur if mand gets into a known error condition that causes the mand crash handler to restart mand and write out a tinfo file in /var/mand. Additionally, a syslog message will be generated.

The HA application will process its state machine once every pre-defined cycle time (the default is 200ms). Each cycle time, it will check all resources and if any resource

© 2013 Edgewater Networks, Inc. All rights reserved.

is deemed down, it will relinquish its Active state. SEM will be shared between each node in the configured pair. This message will either indicate that the system is operational and functioning or that it is down. Each system then decides, based on whether the other system is up and what the other system's priority is compared to its own (Active versus Standby), whether it should be the Active node in the HA cluster.

If the other system signals that it is down, it is obviously considered "Down."

However, if no messages are heard from the other system, it is considered "Absent." This determination is based on message tracking and a message timer. If the receipt of the last message is greater than the configured message lost threshold, then a system is determined "Absent".

There is no configurable setting for the number of messages lost. Rather, the determination is based on the configurable timeout in relation to the cycle time. So, if the timeout setting is 600ms and the cycle time is 200ms, the system waits for 3 messages before considering the other system lost. This parameter for establishing the cycle timeout is the following:

```
/etc/config/failover_defs/FAIL_MSG_TIMEOUT=<number of milliseconds>
```

Administratively Disabled Mode

A system can be taken offline (meaning, taken out of the High Availability mode) for routine maintenance, upgrades, reconfiguration, etc. This is known as "Administratively Disabled" mode. While this mode is enabled, the system under maintenance, though still online, will give up the Active status and will not attempt to regain the Active status. Additionally, it will not attempt to assume Active status if a state transfer request is received from the current Active system. This will persist across reboots so that a firmware upgrade can be tested before reactivating the system. This is enabled via a checkbox in the High Availability configuration page (to be discussed later in this document). Be aware that enabling this setting cause the Standby system to take over as the Active.

Caveat: this mode is only intended for routine maintenance and not for long-term operation. In other words, this mode should be employed if you are upgrading each node in an HA cluster.

Supported Functionalities and Known Limitations

As of VOS 11.6.5, the High Availability feature supports the following services and phone features:

- SIP Trunking – See Note 1 below
- Hosted SIP - See Note 2 below
- BLA, BLF, SCA – See Note 2 below

© 2013 Edgewater Networks, Inc. All rights reserved.



Note 1:

B2BUA usage with a SIP trunking scenario is not supported at this time.



Note 2: Phone Specific Features

- UDP-based BLA, SCA and BLF states are exchanged between the Primary and Secondary nodes.
- TCP-based SIP is supported, but the TCP connection state is not transferred between Primary and Secondary nodes. So the VoIP endpoint must reestablish a TCP connection to the remote server. However, once it does, the feature state has been transferred.
    a) Also, if a failover occurs due to power failure or network failure, the remote hosts will not be aware that the TCP connections have failed. This means that when a call is attempted after a failover, the remote hosts will attempt to use the old TCP connections first. Only after receiving a reset for this connection will they establish a new connection. This initial attempt, reset, reconnection cycle will be repeated up to four times; between a local phone and the EdgeMarc, between the EdgeMarc and the softswitch, the softswitch back to the EdgeMarc, and between the EdgeMarc and the called phone.
    b) Only the Primary/Active unit will have SIP TCP sockets open on port 5060.



Note 3: TCP Functionality

In order for a seamless failover between the Active and Standby devices to occur (as described in (a) above, the newly active EdgeMarc in the HA cluster must be able to send TCP reset packets to the phones and softswitch to trigger the remote systems to re-open the TCP connection.

a) For IP phones on the trusted LAN of the EdgeMarc, this would require that the customer-managed router / DHCP helper allows the TCP Reset.
b) For the WAN softswitch, a firewall rule might be required in the network to allow TCP resets to be sent to the softswitch/IPBE address.

Note that when the Active Node in the HA cluster exits Active mode (due to a resource availability change) thus transferring control to the Standby, the ALG is restarted on the device exiting Active/Master mode. This allows for the release of the server TCP sockets by the system exiting Active mode. If the ALG is a monitored resource, the HA resource monitor will determine that the node is not ready to assume Standby state, and will result in an ephemeral state where the HA node's state monitor will reflect "Failed". This state will prevent the node that is waiting to enter Standby status (and thus be a node that is ready to be failed over to) for up to ten seconds. It will also prevent the device that is currently the Active device to fail over to that node.

See the High Availability Debug and Diagnostics section this document for a complete description of the various transitional and permanent states that might be reflected by the state monitor.

© 2013 Edgewater Networks, Inc. All rights reserved.

As a rule of thumb, assuming default configuration, the following planning factors should be used:

- Failover time: Less than 1 second
- Fallback time: Up to 10 seconds (if the ALG is a monitored resource)

4550/46xx Known Limitations

High Availability does not apply to the following scenarios with the 4608POE:

- Non-Ethernet links (such as T1, MLPPP)
- Embedded POE switch ports cannot be used as a monitored link with which status of the POE interface is detected and shared (and thus used to force a failover if a link goes down between one 4608POE node in an HA cluster to another).
  - SEMs can be transmitted between the two POE switches, but neither HA peer will base a decision to fail over if the physical or data-link connection between the POE switches is lost
- FXS/FXO/PRI: Support for FXS/FXO/PRI requires either:
  - On-site switch of physical links from Active to Standby unit
  - 3rd-party intelligent channel bank to support some type of seamless switchover

General High Availability Scenarios Either Unsupported or Not Tested/Certified

- BLA/SCA presence change stress testing
  - Rapid toggling between lines on a phone, multiple phones or sidecars to generate a storm of Subscribe/Notify transactions during a HA state transition has not been tested.
- Ethernet link failure stress testing
  - Disconnecting/reconnecting Ethernet uplink (WAN to switch) / downlink (LAN to switch) at rapid, arbitrary rates that fall outside of 600ms (time needed for 3 State Exchange Messages to be swapped) and below 1 second (probable time for full state transition for Active system to be fully online and ready to process calls) in order to confuse the HA state machine and put the HA state machine into a race condition has not been tested.
  - This might cause a condition known as “State Shearing”, or a scenario where a switch is induced prior to the full transferal of state between the Active and Standby peers.

© 2013 Edgewater Networks, Inc. All rights reserved.

- Playing “Whack-a-Mole” with the Stateful ALG process
  - Purposefully killing the mand process on the Active and then within 1 second killing mand on the newly Active (previously Standby) and then killing the mand process on the subsequent newly Active device (the original Active) has not been tested.

Support for HA State Exchange

The following table lists the current support for High Availability State Exchange

Table 3: Support for High Availability State Exchange

SIP Trunking – Stateful (not using B2BUA)	FXS/FXO/PRI (see above description)
Hosted SIP telephony – Stateful (See Note above)	Physical link loss on 46xx switch ports (POE or 4-port)
UDP-based SIP BLA/SCA/BLF – Stateful (see note 2 above).	Serial WAN links (T1, MLPPP)
TCP-based SIP BLA/SCA/BLF (Stateful – see note 2 and 3 above)	IPSec VPN, SSL VPN
SIP Call Control Status (Stateful – see note 2 above)	Active wifi client state (4608POE only)
SIP Registration Status (Stateful – see note 2 above)	H.323
Remote/Local Mode status (Stateful)	Access Proxy
DNS reachability status for Session Border Controllers that are identified via FQDN (Stateful)	Proxy ARP
SIP Server reachability for Session Border Controllers identified by “List of SIP Servers” (Stateful)	Router-only Mode
Monitoring of SIP Options/Heartbeats for Survivability (Stateful)	Test UA
Monitoring of SIP requests for Survivability (Stateful)	
Monitoring of SIP Options/Heartbeats for Survivability (Stateful)	
DHCP Server Lease file – allows for unified DHCP server scopes across peers	

© 2013 Edgewater Networks, Inc. All rights reserved.

Important Operational Guidelines

Compatibility

- High Availability is not supported between two different hardware platforms (i.e., a 5300LF2 and 4608POE in the same HA cluster).
- High Availability is not supported with two different VOS versions running on each system (i.e., an Active unit running 11.6.5 and a standby unit running 11.11). This means that an upgrade of an HA cluster should be done either in a maintenance window when both units can be upgraded together, or by upgrading using the Administratively Down function described earlier in this document.

Configuration Changes on an HA Cluster

As a rule of thumb, when making configuration changes either via the GUI or the CLI to either system, the system should be in Administratively Disabled mode. Although there are certain configuration changes that have no applicability to the High Availability State Transfer service, there are several specific configuration changes that will always affect the State Transfer engine. These include, but are not limited to:

- Network
- VoIP ALG
- SIP Server
- SIP trunking
- Survivability
- Routing

Typical Administrative Maintenance Procedure

As of 11.6.5 and beyond, seamless maintenance without service interruption is not supported (see above description for known operational guidelines). Consequently, a routine maintenance procedure, such as a code upgrade or configuration change would include the following steps in order. For the purposes of this description, the terms “Node 1” and “Node 2” will reference the HA Cluster. Again, this should be done in a maintenance window as configuration changes and code upgrades are a service-interrupting procedure.

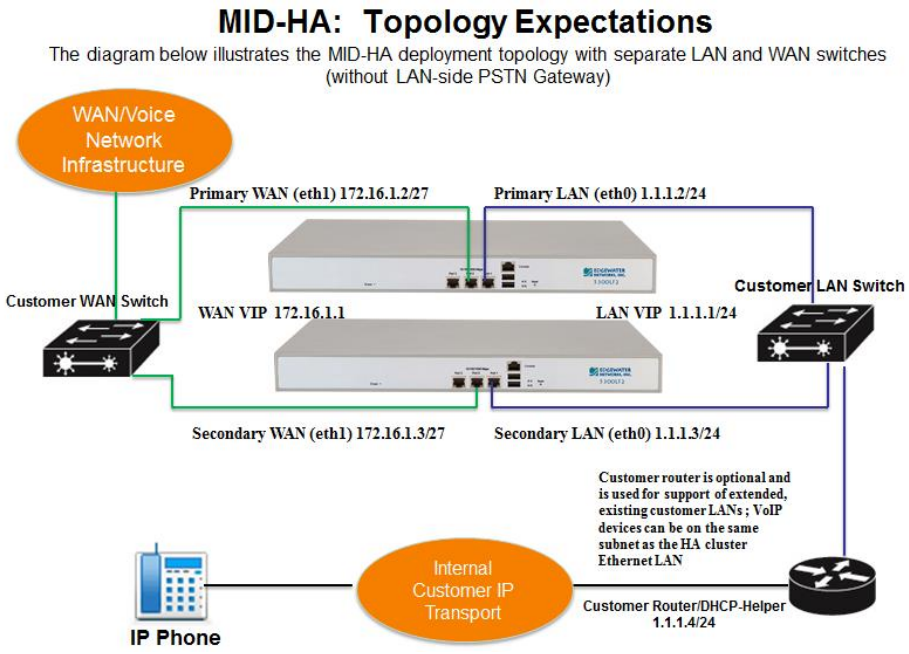
1. Identify the appropriate node in the HA cluster for the maintenance procedure. In this case, it is going to be the Standby (Node 2).
2. Set Node 2 in “Administratively Disabled” mode.
3. Verify that Node 1 remains the Active device.
4. Perform the maintenance operation on Node 2 (code upgrade, configuration change).
5. Set Node 1 in “Administratively Disabled” mode.

© 2013 Edgewater Networks, Inc. All rights reserved.

6. Disable “Administratively Disabled” mode for Node 2.
7. Verify Node 2 enters the Active state.
8. Perform the maintenance operation on Node 1 (code upgrade, configuration change).
9. Disable “Administratively Disabled” mode for Node 1.
10. If the HA cluster is configured for “Revertive”, verify that the appropriate node (configured for Primary) assumes the Active state. If not configured for “Revertive”, the currently Active device (Node 2) will remain the Active device.

High Availability Configuration

The first determination to make for High Availability is to establish if the HA cluster will be “Revertive” or “Non-Revertive”. Refer to the section Revertive Mode and Standard Operation (Non-Revertive) Mode for a definition of these modes. In the diagram below, Revertive Mode is used:



Attached LAN devices, as well as the far-end WAN devices, are unaware that two EdgeMarcs are installed because the pair of devices, although individually numbered with IP addresses on both the LAN and WAN, utilize the concept of a constant IP address called ***Virtual IP Address (VIP)***. A VIP instance is configured for both LAN and WAN. So, for example, LAN IP phones send and receive their signaling and RTP to/from the LAN VIP, and WAN VOIP network infrastructure send and receive information to/from the WAN VIP.

© 2013 Edgewater Networks, Inc. All rights reserved.

To configure the two EdgeMarcs, you must have a total of (at least) three IP addresses for the LAN, and three IP addresses for the WAN. For the purpose of this section, it is assumed that the user will select “Revertive Mode” operation. The address assignment is as follows:

Primary LAN: LAN address for the initially configured Primary node (eth0). This is the “real LAN IP,” or the IP address that is used to connect to the device for management purposes. From the diagram above, this is 1.1.1.2.

Secondary LAN: LAN address for the initially configured Secondary node (eth0). This is the “real LAN IP,” or the IP address that is used to connect to the device for management purposes. From the diagram above, this is 1.1.1.3.

VIP LAN: the VIP LAN address is what all LAN-side devices communicate with. For example, this becomes either the default gateway or outbound proxy address for IP endpoints such as phones on the same LAN or remote networks behind the LAN. From the diagram above, this is 1.1.1.1. Each of these addresses are required to be on the same LAN segment. For example:

Primary LAN: 1.1.1.2/24

Secondary LAN: 1.1.1.3/24

VIP LAN: 1.1.1.1/24

Primary WAN: WAN address for the initially configured Primary node (eth1). This is the “real WAN IP,” or the IP address that is used to connect to the device for management purposes. From the diagram above, this is 172.16.1.2.

Secondary WAN: WAN address for the initially configured Secondary node (eth1). This is the “real WAN IP,” or the IP address that is used to connect to the device for management purposes. From the diagram above, this is 172.16.1.3.

VIP WAN: the VIP WAN address is what all WAN-side devices communicate with. For example, this becomes the signaling address that the upstream IPBE communicates with, or the upstream Layer 3 router. From the diagram above, this is 172.16.1.1. Each of these addresses are required to be on the same WAN segment. For example:

Primary WAN: 172.16.1.2/27

Secondary WAN: 172.16.1.3/27

VIP WAN: 172.16.1.1/27

Tasks for Configuring High Availability

To configure High Availability with a redundant pair, complete the following prerequisite tasks on each HA Node:

1. Configure the LAN and WAN IP addresses (shown). See Task 1: Configure the LAN and WAN IP addresses on each HA Node.
2. Configure the Management Interface (Optional). See Task 2: Configure the Management Interface on each HA Node (Optional, only for the 5300LF2).
3. Configure the High Availability page (shown). See Task 3: Configure the High Availability page on each HA Node.

© 2013 Edgewater Networks, Inc. All rights reserved.

4. Configure identically the VOIP ALG page on each HA node that is part of the HA cluster (not shown).
5. Configure identically the Survivability Page on each HA node that is part of the HA cluster (not shown).
6. Configure identically the DHCP Configuration Page on each HA node that is part of the HA cluster (not shown)

Task 1: Configure the LAN and WAN IP addresses on each HA Node

Configure the LAN and WAN IP addresses. These are real, unique IP addresses. Refer to diagrams in this document for guidelines.

On each device, complete the following steps:

1. Choose **Network** from the EdgeMarc Configuration menu to open the Network Configuration page.
2. Enter the IP address for the LAN interface in the **IP Address** field under the LAN Interface Settings section.
3. Enter the IP address for the WAN interface and complete the following fields.
    a. Select the **Static IP Address** radio button (DHCP WAN and T1 WAN not supported with HA).
    b. Enter an IP address in the **IP Address** field under the WAN Interface Settings section.
    c. Complete any remaining fields (LAN subnet mask, WAN subnet mask, default gateway, DNS, etc).
4. Click **Submit** to save your results.

Task 2: Configure the Management Interface on each HA Node (Optional, only for the 5300LF2)

As demonstrated on the GUI configuration for High Availability, the HA state machine exchanges current state information across configured interfaces, to include LAN, WAN and MGMT (on the 5300LF and 5300LF2). If desired, configure the 5300 model's management interface.

On each device, complete the following steps:

1. Choose **System > Management Interface** from the primary web configuration menu.
2. Select the **Enable Management Interface** box.
3. Enter a unique management IP address in the **Management Interface IP Address** field. Each Management Interface IP address must be on the same subnet as the address of the HA peer node.
4. Enter the appropriate subnet mask in the **Subnet Mask** field.
5. Click **Submit** to save your results.

© 2013 Edgewater Networks, Inc. All rights reserved.

Task 3: Configure the High Availability page on each HA Node

Configure all parameters on the High Availability page for each device.

On each device, complete the following steps:

1. Choose **System > High Availability** from the primary web configuration menu.
2. **Enable High Availability**: Select the **Enable High Availability** check box to enable the feature.
3. **Enable HA DHCP**: Select **Enable HA DHCP**, if applicable (only in version 12.6.2 and beyond). This provides for stateful exchange of DHCP lease files.



Note:

The DHCP servers must be configured identically in each peer node.

4. **Enable Revertive Mode**: Select whether to designate **Enable Revertive Mode**.



Note:

If this checkbox is checked, it MUST be checked on the peer node's configuration as well. This designation means that the configured Primary will always assume Active state if systems and processes are normalized.

5. **Designation**: Specify the **Designation** field as **Primary** or **Secondary**.



Note:

One device must be designated as the Primary device and the other device must be designated as Secondary device.

6. **Instance ID**: Enter the appropriate **Instance ID**. The **Instance ID** is specific to two nodes in a Failover Group. Each of those nodes must be configured with the same Instance ID. This must also be unique to that Group so that multiple Failover Groups on the same network do not disrupt each other.
7. **Password**: Enter the password in the **Password** field. This is the credential that the High Availability pair will utilize to share information and also to ensure that invalid information is not passed via unauthorized system access. The password must be the same on both nodes within the same Failover Group.
8. **LAN IF Virtual IP address**: Enter the LAN Virtual IP address into the **LAN IF Virtual IP address** field. Remember, this address must be on the same subnet as the LAN IP of both HA peers, and must be different than each. Refer to conceptual diagrams in this document.
9. **WAN IF Virtual IP address**: Enter the WAN Virtual IP address into the **WAN IF Virtual IP address** field. Remember, this address must be on the same subnet as the WAN IP of both HA peers, and must be different than each. Refer to conceptual diagrams in this document.

© 2013 Edgewater Networks, Inc. All rights reserved.

10. **LAN IF Remote Address**: For each device, enter the actual configured (from the peer network page) LAN interface address of the other device in the **LAN IF Remote Address** field.
11. **WAN IF Remote Address**: For each device, enter the actual configured (from the peer network page) WAN interface address of the other device in the **WAN IF Remote Address** field.
12. **Management Remote Address**: If applicable, enter the actual configured (from the peer Management Interface configuration page) Management Interface address of the other device in the **Management Remote Address** field.
13. **Enable State Transfer**: For each link that will transfer state information, select its corresponding **Enable State Transfer** checkbox.
14. **Remote Name:** Configure the peer's **Remote Name** in this field, if desired. This is usually a hostname indicative of the peer. This will populate the appropriate SNMP trap with the hostname of the peer rather than the IP address of the peer.

You must enable state transfer for one or more of the IP interfaces. If you enable state transfer for multiple interfaces, you have a higher chance of successful state transfer during link failures. The purpose of the State Transfer, as previously discussed, is to facilitate a sharing of information so that when the Active EdgeMarc fails, or one of the links fails, the Secondary node will detect the failure and take over as Active.

Figure 3 shows an example of the initially designated Primary/Active node configuration.

© 2013 Edgewater Networks, Inc. All rights reserved.

Figure 3. **Primary/Active Node Configuration**

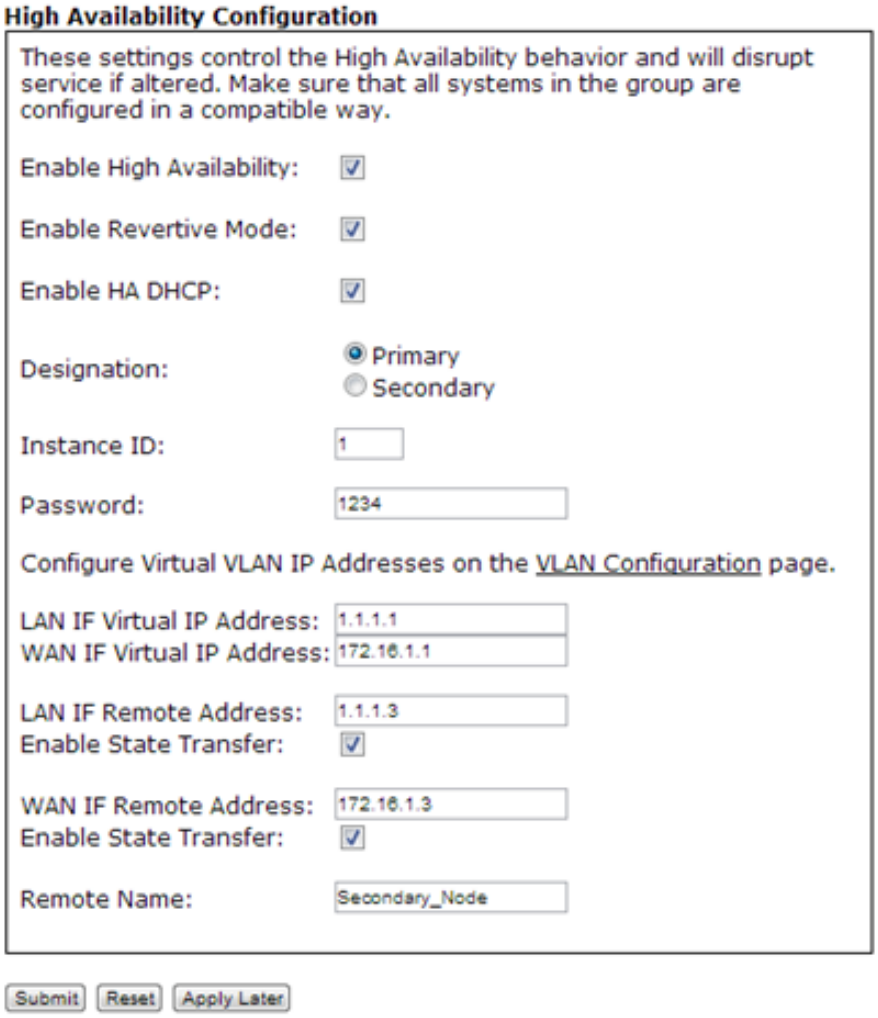
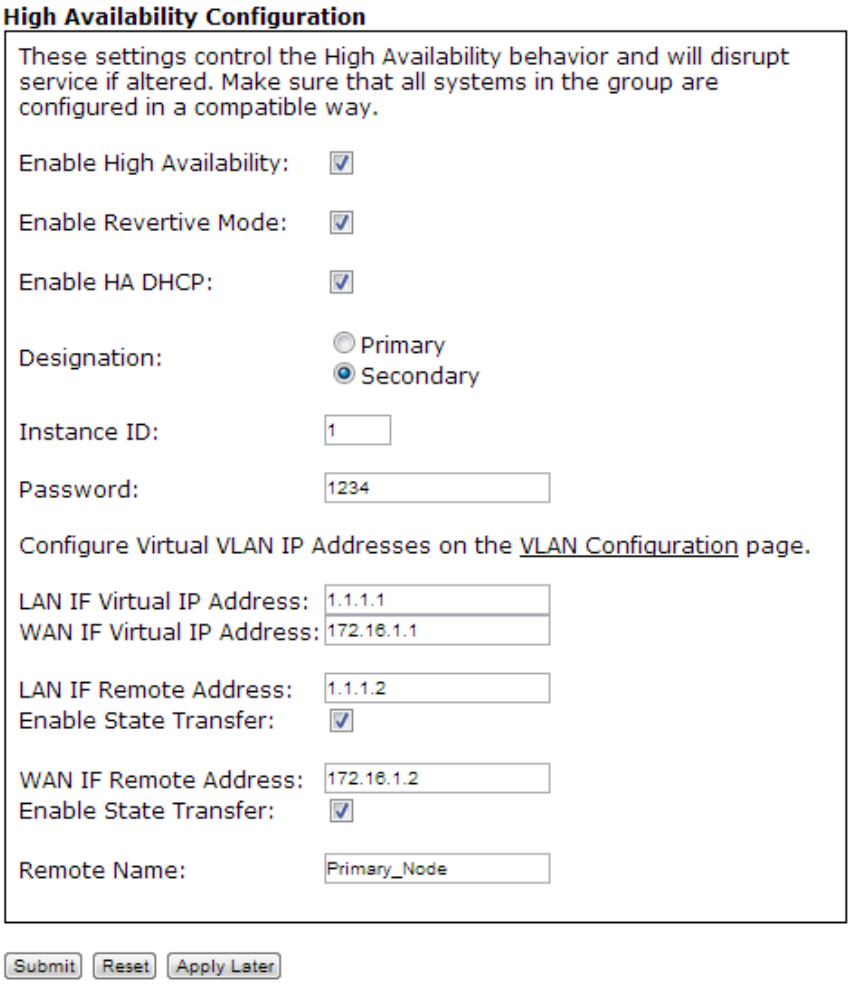


Figure 4 shows an example of the initially configured Secondary / Standby node configuration.

© 2013 Edgewater Networks, Inc. All rights reserved.

Figure 4. **Secondary/Standby Node Configuration**



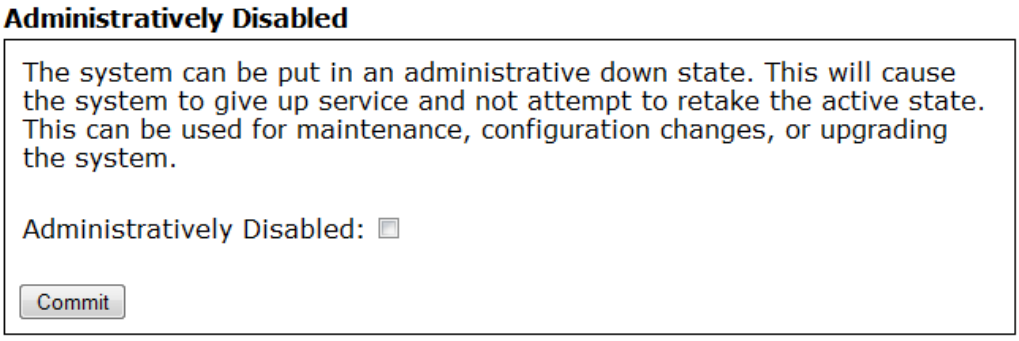
Note:

For redundant pairs that process high volumes of SIP signaling and RTP media, enable State Transfer of at least two, if not all three, IP interfaces (LAN and WAN, Management is optional). If one interface should fail, this increases the probability that the state will be successfully transferred between devices.

The portion of the Stateful Failover GUI shown in Figure 5 shows an example of how to set one of the nodes to “Administratively Disabled” status. This would be used in the “Typical Administrative Maintenance Procedure” described above.

© 2013 Edgewater Networks, Inc. All rights reserved.

Figure 5. **Setting a Node to Administratively Disabled Status**



High Availability Debug and Diagnostics

The following table defines the operational states reflected in either the GUI, syslog, or HA debug.

State	Description
Active	Active call processing node. All VoIP LAN→WAN and WAN→LAN traffic is passing through this node via the masqueraded VIP. A node reporting this indicates that the HA state machine is fully functional and is the Active state. “Master” is also reflected in logging.
Standby	Standby call processing node. The system is online but not processing calls or SIP state. “Backup” is also reflected in logging.
Initializing	A node reporting the initializing state indicates that the High Availability state machine is starting up/normalizing. This is an ephemeral status.
Down	A node reporting the Down state indicates that this state is Administratively down (a configuration/management option).
Failed	A node reporting the Failed state indicates that a failure has occurred that prevented the HA machine from initializing/normalizing. It can also indicate a lack of resources of the reporting node for an extended period (beyond the default 600ms). The Failed state is also reported on a node with HA enabled that cannot ‘find’ it’s peer.
Disabled	A node reporting the Disabled stated indicates that the High Availability feature is not enabled.
Unknown	A node reporting the Unknown state is in a state not reflected by either “Down” or “Failed”. The Unknown state is reported during the transitional time between ‘startup’ and Initializing state. It is an ephemeral status
Absent	A state characterizing one node in the HA pair, portrayed in the logs of the other node, if that node does not receive SEMs from the partner

© 2013 Edgewater Networks, Inc. All rights reserved.

Command Request Mode Usage

```
# failover <command>
```

To show the failover command mode help:

```
# failover help
```

State Changes

When a system changes state, the current status will be written to the status file: /var/run/failover.status

The content has the following format:

```
<numeric-state> <state-string> <other-hosts>
```

Where <numeric-state> is the numeric value of the current state, which can be one of the following values:

STATE_INIT = 0

STATE_DOWN = 1

STATE_FAILED = 2

STATE_BACKUP = 3

STATE_MASTER = 4

The <state-string> is the human readable state description

The <other-hosts> is a flag which is 1 if there are other hosts present, but 0 if there are no other hosts present.

High Availability Logging and Debug

```
# failover log 7 (turns debug on)
# failover log 0 (turns debug off)
```

To query the current instant state (if the configured instance ID is 1):

```
# failover instance 1 show
```

© 2013 Edgewater Networks, Inc. All rights reserved.

MID-HA: Debug using Failover Logging

The command group 'failover instance 1 show' displays certain state characteristics of the HA peer that you run the command on. This shows the Primary/Active system's state

```
# failover instance 1 show
[Sending request 'instance 1 show'...]

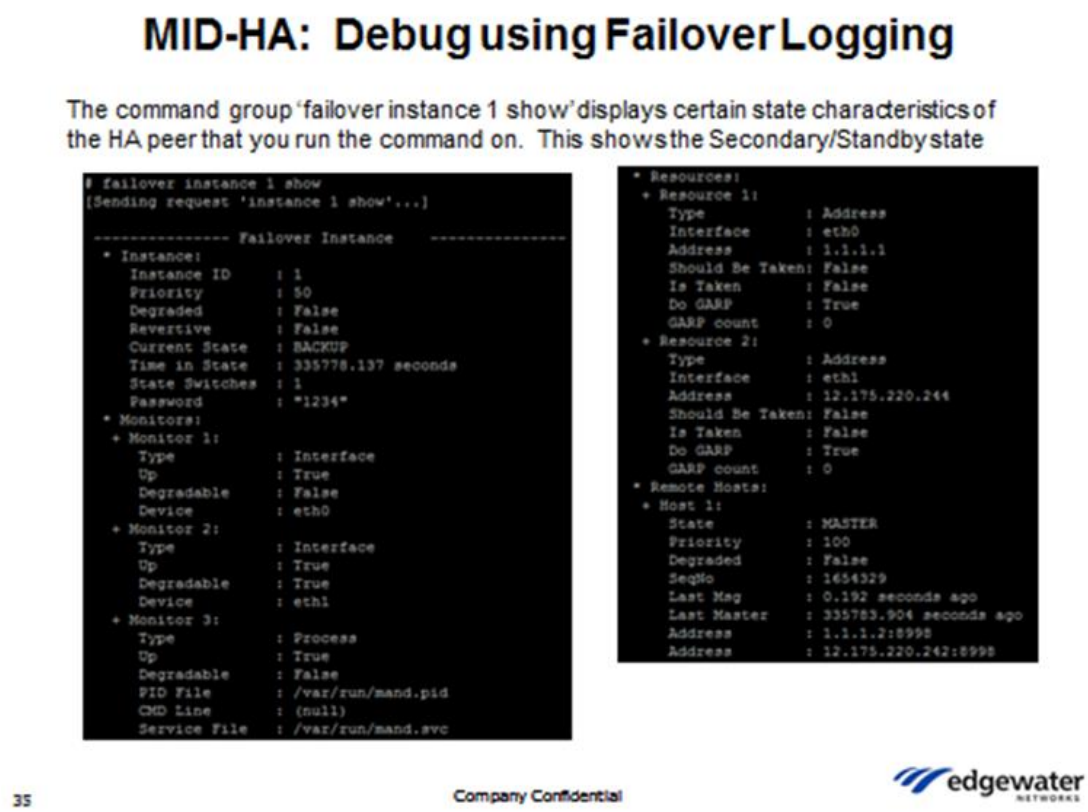
--------------- Failover Instance    ---------------
 * Instance:
    Instance ID      : 1
    Priority         : 100
    Degraded         : False
    Revertive        : False
    Current State    : MASTER
    Time in State    : 334974.086 seconds
    State Switches   : 2
    Password         : "1234"
 * Monitors:
  + Monitor 1:
    Type             : Interface
    Up               : True
    Degradable       : False
    Device           : eth0
  + Monitor 2:
    Type             : Interface
    Up               : True
    Degradable       : True
    Device           : eth1
  + Monitor 3:
    Type             : Process
    Up               : True
    Degradable       : False
    PID File         : /var/run/mand.pid
    CMD Line         : (null)
    Service File     : /var/run/mand.svc
```

```
 * Resources:
  + Resource 1:
    Type             : Address
    Interface        : eth0
    Address          : 1.1.1.1
    Should Be Taken: True
    Is Taken         : True
    Do GARP          : True
    GARP count       : 0
  + Resource 2:
    Type             : Address
    Interface        : eth1
    Address          : 12.175.220.244
    Should Be Taken: True
    Is Taken         : True
    Do GARP          : True
    GARP count       : 0
 * Remote Hosts:
  + Host 1:
    State            : BACKUP
    Priority         : 50
    Degraded         : False
    SeqNo            : 1650621
    Last Msg         : 0.087 seconds ago
    Last Master      : 220989.981 seconds ago
    Address          : 1.1.1.3:8998
    Address          : 12.175.220.243:8998
```


Company Confidential

edgewater NETWORKS

© 2013 Edgewater Networks, Inc. All rights reserved.



To show the internal syslog back-log (saved internally to prevent syslogs from disappearing on a system with a small log size):

```
   # failover syslog
--------------- Recent Failover Logs ---------------
Apr 15 17:02:05 INF monitor.c:242       mon_check_status():
MON[1:Interface:eth0]->UP, reason="Interface up"
```

High Availability SNMP Support

The High Availability feature has a MIB and supports trap notification to upstream management servers. Below is a subset of the Edgewater Networks Fault Tolerance MIB relevant only to High Availability functionality. Please note that the SNMP support detailed below is only available in code versions 11.12.x and above (as of the writing of this document).

© 2013 Edgewater Networks, Inc. All rights reserved.

```
-- EDGEWATER-EMCPE-FAULT-TOLERANCE-MIB.my
-- MIB generated by MG-SOFT Visual MIB Builder Version 7.0  Build 220
-- Tuesday, May 8, 2012 at 13:00:00
EDGEWATER-EMCPE-FAULT-TOLERANCE-MIB DEFINITIONS ::= BEGIN

IMPORTS
       ewnEdgeMarcCPE
FROM EDGEWATER-SMI-MIB
       OBJECT-GROUP, NOTIFICATION-GROUP
FROM SNMPv2-CONF
       OBJECT-TYPE, MODULE-IDENTITY, NOTIFICATION-TYPE
FROM SNMPv2-SMI;


-- 1.3.6.1.4.1.19139.100.200
emcpeFaultTolerance MODULE-IDENTITY
       LAST-UPDATED "201205081300Z"-- May 08, 2012 at 13:00 GMT
       ORGANIZATION
"Edgewater Networks, Inc."
       CONTACT-INFO
"Edgewater Networks, Inc.
Phone: 408.351.7200
Web: www.edgewaternetworks.com
E-mail: support@edgewaternetworks.com"
       DESCRIPTION
"Edgewater Networks, Inc. Fault Tolerance MIB module."
       ::= { ewnEdgeMarcCPE 200 }


--
-- Node definitions
--

-- 1.3.6.1.4.1.19139.100.200.4
emcpeFTHighAvailEnable OBJECT-TYPE
       SYNTAX INTEGER
{
disable(0),
enable(1)
}
       MAX-ACCESS read-only
       STATUS current
       DESCRIPTION
"Flag indicating whether the High Availability
feature is enbabled."
       ::= { emcpeFaultTolerance 4 }

-- 1.3.6.1.4.1.19139.100.200.8
emcpeFTHighAvailState OBJECT-TYPE
       SYNTAX INTEGER
{
initializing(0),
down(1),
failed(2),
backup(3),
master(4),
disabled(5),
unknown(6)
}
       MAX-ACCESS read-only
       STATUS current
```

© Edgewater Networks, Inc. All rights reserved.

```
       DESCRIPTION
"Current operating state of the High Availability feature."
       ::= { emcpeFaultTolerance 8 }

-- 1.3.6.1.4.1.19139.100.200.10
emcpeFTHighAvailPartner OBJECT-TYPE
       SYNTAX OCTET STRING (SIZE (0..80))
       MAX-ACCESS read-only
       STATUS current
       DESCRIPTION
"Hostname of the High Availability partner device."
       ::= { emcpeFaultTolerance 10 }
** NOTE – if the Hostname of peer device is not configured in the High Availability
configuration (either GUI page or CLI file), this object defaults back to the IP address
of the partner.
-- 1.3.6.1.4.1.19139.100.200.14
emcpeFTSurviveLanEnable OBJECT-TYPE
       SYNTAX INTEGER
{
disable(0),
enable(1)
}
       MAX-ACCESS read-only
       STATUS current
       DESCRIPTION
"Flag indicating whether the LAN-side Survivability
feature is enbabled."
       ::= { emcpeFaultTolerance 14 }

-- 1.3.6.1.4.1.19139.100.200.18
emcpeFTSurviveLanState OBJECT-TYPE
       SYNTAX INTEGER
{
lost(0),
acquired(1)
}
       MAX-ACCESS read-only
       STATUS current
       DESCRIPTION
"Current operating state of the LAN-side Survivability
feature."
       ::= { emcpeFaultTolerance 18 }

-- 1.3.6.1.4.1.19139.100.200.20
emcpeFTSurviveLanPartner OBJECT-TYPE
       SYNTAX OCTET STRING (SIZE (0..80))
       MAX-ACCESS read-only
       STATUS current
       DESCRIPTION
"Hostname of the Survivability trunking device."
       ::= { emcpeFaultTolerance 20 }

-- 1.3.6.1.4.1.19139.100.200.24
emcpeFTSurviveWanEnable OBJECT-TYPE
       SYNTAX INTEGER
{
disable(0),
enable(1)
}
       MAX-ACCESS read-only
```

© 2013 Edgewater Networks, Inc. All rights reserved.

```
       STATUS current
       DESCRIPTION
"Flag indicating whether the WAN-side Survivability
feature is enbabled."
       ::= { emcpeFaultTolerance 24 }

-- 1.3.6.1.4.1.19139.100.200.28
emcpeFTSurviveWanState OBJECT-TYPE
       SYNTAX INTEGER
{
local(0),
remote(1)
}
       MAX-ACCESS read-only
       STATUS current
       DESCRIPTION
"Current operating state of the WAN-side Survivability
feature."
       ::= { emcpeFaultTolerance 28 }

-- 1.3.6.1.4.1.19139.100.200.50
emcpeFaultToleranceTraps OBJECT IDENTIFIER ::= { emcpeFaultTolerance 50 }

-- 1.3.6.1.4.1.19139.100.200.50.4
emcpeFTHighAvailActive NOTIFICATION-TYPE
       OBJECTS { emcpeFTHighAvailState, emcpeFTHighAvailPartner }
       STATUS current
       DESCRIPTION
"High Availability trap indicating that the feature is Active.
The availability state and the hostname of the standby
device are included in the trap."
       ::= { emcpeFaultToleranceTraps 4 }

-- 1.3.6.1.4.1.19139.100.200.50.8
emcpeFTHighAvailStandby NOTIFICATION-TYPE
       OBJECTS { emcpeFTHighAvailState, emcpeFTHighAvailPartner }
       STATUS current
       DESCRIPTION
"High Availability trap indicating that the feature is in Standby.
The availability state and the hostname of the primary device
are included in the trap."
       ::= { emcpeFaultToleranceTraps 8 }

-- 1.3.6.1.4.1.19139.100.200.50.12
emcpeFTSurviveLanAquired NOTIFICATION-TYPE
       OBJECTS { emcpeFTSurviveLanState, emcpeFTSurviveLanPartner }
       STATUS current
       DESCRIPTION
"Survivability trap indicating that the LAN-side keep-alive
message has been aquired. The survivability state and
the hostname of the trunking device are included in the trap."
       ::= { emcpeFaultToleranceTraps 12 }

-- 1.3.6.1.4.1.19139.100.200.50.16
emcpeFTSurviveLanLost NOTIFICATION-TYPE
       OBJECTS { emcpeFTSurviveLanState, emcpeFTSurviveLanPartner }
       STATUS current
       DESCRIPTION
"Survivability trap indicating that the LAN-side keep-alive
message has been lost. The survivability state and the
```

© 2013 Edgewater Networks, Inc. All rights reserved.

```
hostname of the trunking device are included in the trap."
       ::= { emcpeFaultToleranceTraps 16 }

-- 1.3.6.1.4.1.19139.100.200.50.20
emcpeFTSurviveWanRemote NOTIFICATION-TYPE
       STATUS current
       DESCRIPTION
"Survivability trap indicating that the system has entered
Remote mode."
       ::= { emcpeFaultToleranceTraps 20 }

-- 1.3.6.1.4.1.19139.100.200.50.24
emcpeFTSurviveWanLocal NOTIFICATION-TYPE
       STATUS current
       DESCRIPTION
"Survivability trap indicating that the system has entered
Local mode."
       ::= { emcpeFaultToleranceTraps 24 }

-- 1.3.6.1.4.1.19139.100.200.50.100
emcpeFaultToleranceTrapsGroups OBJECT IDENTIFIER ::= { emcpeFaultToleranceTraps 100 }

-- 1.3.6.1.4.1.19139.100.200.50.100.1
emcpeFaultToleranceTrapsGroup NOTIFICATION-GROUP
       NOTIFICATIONS { emcpeFTHighAvailActive, emcpeFTHighAvailStandby,
emcpeFTSurviveLanAquired, emcpeFTSurviveLanLost, emcpeFTSurviveWanRemote,
emcpeFTSurviveWanLocal }
       STATUS current
       DESCRIPTION
"Description."
       ::= { emcpeFaultToleranceTrapsGroups 1 }

-- 1.3.6.1.4.1.19139.100.200.100
emcpeFaultToleranceGroups OBJECT IDENTIFIER ::= { emcpeFaultTolerance 100 }

-- 1.3.6.1.4.1.19139.100.200.100.1
emcpeFaultToleranceGroup OBJECT-GROUP
       OBJECTS { emcpeFTHighAvailEnable, emcpeFTHighAvailState, emcpeFTSurviveLanEnable,
emcpeFTSurviveLanState, emcpeFTSurviveWanEnable,
emcpeFTSurviveWanState, emcpeFTHighAvailPartner, emcpeFTSurviveLanPartner }
       STATUS current
       DESCRIPTION
"Description."
       ::= { emcpeFaultToleranceGroups 1 }


       END
--
-- EDGEWATER-EMCPE-FAULT-TOLERANCE-MIB.my
```

© 2013 Edgewater Networks, Inc. All rights reserved.

Edgewater Networks, Inc.
2895 Northwestern Parkway
Santa Clara, CA 95051
Phone: (408) 351-7200
info@edgewaternetworks.com



© Edgewater Networks, Inc. All rights reserved.