Avaya Aura® Application Enablement (AE) Services High Availability (HA) White Paper

June 10, 2014

Table of Contents

1 Introduction

This white paper is intended for those responsible for architecting, designing, and/or deploying an application or an Avaya Aura® Application Enablement Services server in a High Availability (HA) configuration.

Uninterrupted telephony is important for many enterprises, especially for mission critical applications. Avaya Aura® Application Enablement (AE) Services on System Platform (SP) supports a high availability (HA) cluster of two nodes. The active server node automatically fails over to the standby node in the event of a hardware failure. Client applications are able to reestablish communication with the AE Services cluster when the failover is complete.

The AE Services HA solution is not supported on the AE Services Software-Only and Bundled offers.

AE Services provide support for 3 High Availability (HA) solutions:

- Geo Redundant High Availability (GRHA)
- Machine Preserving High Availability (MPHA)
- Fast Reboot High Availability (FRHA)

GRHA is not a state preserving HA in the AE Services 6.3.1 Release. However starting in the AE Services 6.3.3 release, GRHA is a partial state preserving HA. The state associated with the AE Services service DMCC is preserved. When a controlled failover occurs, AE Services are stopped on the current active VM and AE Services are started on the new active VM (previously the standby VM). During this phase the DMCC service will load its preserved state data. GRHA allows two AE Services Virtual Machines (VMs) to be placed in two datacenters that are separated by a LAN/WAN. The VM host can be either Avaya Aura System Platform or VMware.

MPHA is a state preserving HA based on check pointing a running VM at frequent intervals (e.g., 50milli seconds). At each check point, memory (including CPU registers) and disk state of a protected VM are synchronized with the standby server. In the case of a failover, the standby server becomes active and in the process activates the replicated (synchronized at the latest check point) VM.

FRHA is a partial state preserving HA similar to GRHA offered in the AE Service 6.3.3 release.

MPHA and FRHA are offered via System Platform where the active and standby System Platform servers hosting the AE Services servers are connected via a crossover cable.

FRHA/MPHA/GRHA comparison table

HA Solution	State Preserving Failover	Cross-Over Cable	Impacts AES capacity	Single AE Services Server IP Address	One AE Services License File	Host VM
FRHA	**Partial (DMCC only)**	**Yes (1 Gb NIC)**	**No**	**Yes**	**Yes**	**SP**
MPHA	**Yes**	**Yes (10 Gb NIC)**	**Yes (~50%)**	**Yes**	**Yes**	**SP**
GRHA	• **No (AES 6.3.1)** • **Partial (AES 6.3.3 or later for DMCC only)**	**No (over LAN)**	**No**	• **No (AES 6.3.1 or later)** • **Yes (AES 6.3.3 or latter)**	**Yes**	**SP or VMware**

This white paper also focuses on the AE Services interactions with Avaya Aura® Communication Manager (CM) 6.0 (or later) in a HA configuration for the survivable core server (also known as Enterprise Survivable Server or ESS) and the survivable remote server (also known as Local Survivable Processor or LSP).

2 Geo Redundant High Availability (GRHA)

2.1 Overview

AE Services Release 6.3.1 introduced the Geo Redundant High Availability (GRHA) option that allows active and standby AE Services servers to be located in two different datacenters. The VM host can be either Avaya Aura System Platform or VMware.

With the Machine Preserving High Availability (MPHA) and Fast Reboot High Availability (FRHA) solutions that are offered via Avaya Aura System Platform, active and standby AE Services servers are connected via a crossover cable. As per the CAT5 and CAT6 Ethernet cable specification, the cable between the servers should not be longer than 100 meters. The GRHA solution removes this limitation.

GRHA allows two AE Services Virtual Machines (VMs) to be placed in two datacenters that are separated by a LAN/WAN with round trip time within 100 milliseconds. To ensure AE Services

server does not failover due to hardware failure, the GRHA offer when used with System Platform can leverage the MPHA technology to provide hardware protection for AE Services VM in each datacenter. **Note**: MPHA is not supported on VMware.

For more information on MPHA, please refer to the MPHA section in this white paper.

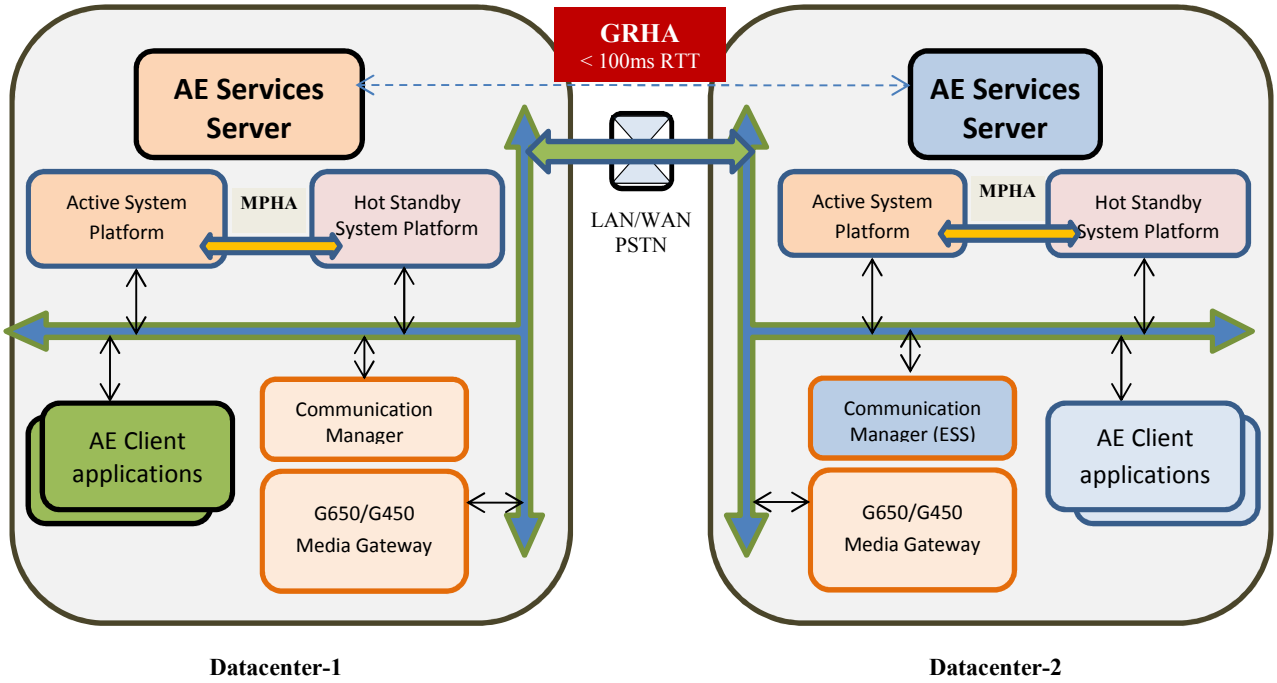

Note: MPHA provides hardware protection in each datacenter

In this document the term “controlled failover” refers to a failover requested by either an administrator or by software logic when it detects degradation in the state of health of the current active server. The term “uncontrolled failover” refers to a failover which occurs because the current active server is not reachable from the current standby server.

GRHA is not a state preserving HA in the AE Services 6.3.1 Release. In the AE Services 6.3.3 and later release, GRHA is a partial state preserving HA solution for the AE Service DMCC only.

When a controlled failover occurs, AE Services are stopped on the current active VM and AE Services are started on the new active VM (previously the standby VM). In case of an uncontrolled failover, AE Services are started on the new active (previously standby) VM. Depending on the reason for the uncontrolled failover, the previous active VM could be in an isolated network or it could be in the shutdown state.

2.2 Key features

Once GRHA is configured via the AE Services management console, the active AE Services configuration data is kept in sync with the AE Services VM in the standby datacenter. An arbiter daemon process running on each AE Services VM ensures that only one AE Services server provides service at a time.

Controlled failover requests can be made by the system administrator (from the AE Services Management Console) or can be requested by the Arbiter (running on each AE Services server) if it detects health deterioration on the active server. The Arbiter daemon will not request a failover if the standby server health is the same as the active server or worse.

2.2.1 Controlled failover of AE Services to standby datacenter AE Services VM

A controlled failover may occur for any of the following scenarios:

- The administrator requested a failover via the AE Services management console.
- The Arbiter daemon process periodically checks to ensure that the root file system is read-write enabled and that the AE Services VM has at least 6% of free memory. A controlled failover is initiated if either the root file system becomes read-only or if the active AE Services VM has less than 6% of free memory.
- The Arbiter daemon process periodically checks to ensure that it can reach an administered network destination (via ICMP ping). The default frequency and the network destination can be changed when HA is enabled. By default it takes about 10 seconds to detect network failure. If any of the administered network destinations are not reachable from the active AE Services VM, but are reachable from the standby AE Service VM, a controlled failover is initiated.
  - Note: if it is not desirable for the standby AE Services VM to become active (in the standby datacenter) due to a network ping failure in the active datacenter, then the “ping targets” fields on the GRHA configuration page may be left empty.

2.2.2 Automatic activation of AE Services on standby datacenter

The GRHA configuration page allows customers to administer a “failure detection” interval. This interval can be anywhere between 1 second and 1 hour. If the active AE Services VM is not reachable from the AE Services server running in the standby datacenter for a duration of the administered “failure detection” interval, then the standby AE Services VM activates AE Services and therefore becomes active.

2.2.3 Automatic recovery from split brain condition

When the active and standby servers fail to reach each other over the configured paths, each server assumes that the other server is dead. If both servers are powered and healthy, both will be running in active mode. This is known as a split brain condition. In a split brain condition, GRHA ensures that neither server can reach the other over a switched IP network; therefore only one server can be connected to the network. And therefore only one AE Services VM can provide service. Once the network heals, the Arbiter will choose which AE Services VM to run in active mode. If a "preferred node" is administered, then that node is made active. If a "preferred node" is not administered, then the server that became active most recently is chosen to be the active AE Services server.

Note that the current active AE Services VM will remain active in any of the following scenarios, even if the current standby AE Services VM is administered as the preferred node:

- If the standby AE Services server is rebooted.
- If the Arbiter daemon process on the standby AE Services server is restarted.
- If GRHA is stopped (disabled) and then started (enabled) on the standby AE Services server.
- If the Active AE Services VM is rebooted and the administered "failure detection" interval is greater than 2 minutes. If the administered "failure detection" interval is less than 2 minutes, then the standby AE Services VM may become active.

2.3 Benefits of GRHA

- Protects against datacenter failures.
- Protects against network failures (if configured).
- A crossover cable is not required. Both of the AE Services VM servers are separated by the LAN/WAN. It is recommended that both of the AE Services VMs have the same CPU and memory configuration.
- Only one set of AE Services licenses is needed for both AE Services servers (active and standby).
- Allows the AE Services VM in each datacenter to be in different networks or subnets where the static address of each server is used. GRHA is also supported where both AE Services servers share a single virtual IP address in the same subnet as both AE Services servers.
- GRHA in and of itself utilizes very little CPU resources, and therefore does not impact AE Services server capacities. However, note that MPHA does have an impact on AE Services server capacities. MPHA is required as part of the GRHA offer to provide hardware protection in the AE Services 6.3.1 release. MPHA is optional for GRHA in the AE Services 6.3.3 and later release.

- Three levels of GRHA licenses: SMALL, MEDIUM and LARGE. Please refer to the ***Avaya Aura® Application Enablement Services 6.3.1 Administration and Maintenance Guide***, section "Administering the Geo Redundant High Availability feature" for more information.

2.4 Effect of a controlled/uncontrolled failover on AE Services clients

AE Services applications must have the ability to connect to two AE Services IP addresses if a virtual IP address is not used. When a failover occurs, the application will detect a session (or socket) drop. The application should try to get service from the same AE Services server first. If a session cannot be reestablished within 2 minutes, the application must try the "other" AE Services server IP address to get service. Subsequent attempts to connect to either IP address should be made every 10 seconds and up to 6 times before moving on to the "other" AE Services server IP address.

After the application connects using a new session, it must re-establish all monitors and register all the endpoints as if the AE Services server came out of a reboot.

The time it takes for the application to start receiving service would depend on the total time associated with following activities:

- Time it takes for the standby to detect a failure of the active AE Services VM. This depends on administered "failure detection" interval and applies only in case of uncontrolled failover. For a controlled failover, this time is almost 0.
- Time it takes for AE Services to be activated on the new active server. Currently this time is approximately 1 minute.
- Time it takes for the application to connect to the new AE Services server and to recreate all its monitors/registrations/associations. This time depends on the number of devices the application is trying to monitor/register.

2.5 Limitations:

- GRHA is not supported on the Software-Only and Bundled AE Services offers.
- GRHA is not supported on System Platform if the AE Services VM is not protected using MPHA technology for the AE Services 6.3.1 release. Note, the use of MPHA with GRHA is optional for the AE Services 6.3.3 and later release.
- GRHA is supported only on IPV4 networks.
- GRHA does not protect against AE Services software failures

3 Machine Preserving High Availability (MPHA)

3.1 Overview

In the System Platform 6.2.1 release, Avaya Aura System Platform introduced two additional high availability solutions as part of Locally Redundant High Availability (LRHA) configurations. LRHA employ two servers connected over a crossover cable in active standby mode. The two servers (active and standby) must have the same System Platform 6.2.1 or later software installed.

LRHA configurations supported in System Platform 6.2.1 and later include:

- Fast Reboot HA (FRHA): when enabled for a VM, the VM is rebooted on the new server every time a failover occurs, for both controlled and uncontrolled failovers.
- Live Migration HA (LMHA): when enabled for a VM, the VM is Live Migrated to the new server when a controlled failover occurs. For uncontrolled failovers the VM will be rebooted on the new server (previously standby).
- Machine Preserving HA (MPHA): when enabled for a VM, the VM is activated on the new (previously standby) server for both controlled and uncontrolled failovers.

This section focuses on the MPHA feature. For more formation on LRHA configurations please see "**Installing and Configuring Avaya Aura™ System Platform" and "Administering Avaya Aura™ System Platform**" documents available at http://support.avaya.com.

In System Platform 6.2.1 and later, if MPHA is selected for a VM, the remaining VMs are automatically set for LMHA. LMHA is only available in the context of a VM enabled with MPHA.

In this document the term "controlled" failover refers to a failover requested by either an administrator or by software logic, when it detects degradation in the state of health of the current active server. The term "uncontrolled" failover refers to a failover which occurs because the current active server is not reachable from the current standby server.

MPHA is a state preserving HA solution based on check pointing a running VM at frequent intervals (e.g. 50 milliseconds). At each check point, memory (including CPU registers) and disk state of a protected VM are synchronized with the standby server. In case of a failover (e.g. the active server dies abruptly or administrator requested a failover), the standby server becomes active and in the process activates the replicated (synchronized at the latest check point) VM. Note, System Platform will restart the VM in which AE Services is running if it stops responding as a result of a software fault condition. This failure will not cause a failover.

Memory replication at its core is based on the Xen hypervisor and its Live-migration technology. The Disk replication is based on the open source application Distributed Replicated Block Device (DRBD).

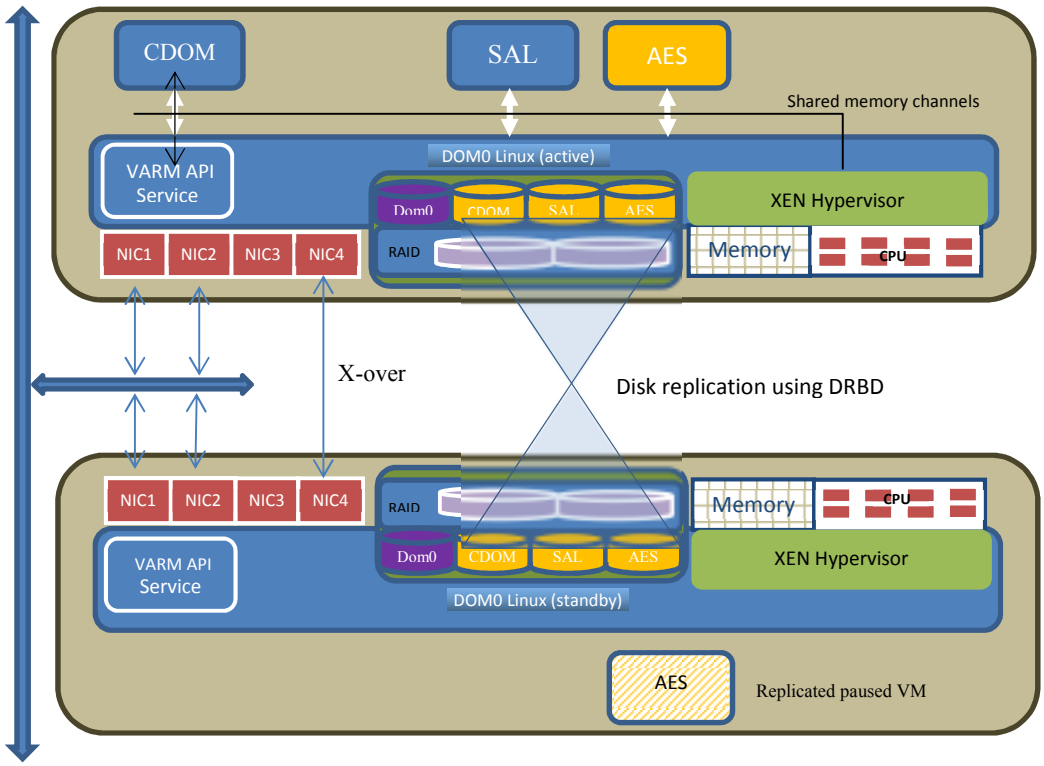

VARM=Virtual Appliance Resource Management

3.2 Key features

3.2.1 Seamless failover in the event of controlled failover requests

For controlled failover requests, MPHA ensures completion of the last check-point before activating the MPHA protected VM. All LMHA protected VMs are live migrated to the new active server. Controlled failover requests can be made by the system administrator (from CDOM Web console) or can be requested by the State of Health monitoring daemon (running on dom0 of each server), if it detects hardware health deterioration on the active server. State of Health monitoring daemon would not request a failover if the standby server health is the same as the active server or worse. Controlled failover would happen in any of the following cases:

- Administrator requested failover via CDOM web-console

- If a faulty hardware is detected by the State of Health daemon. The State of Health daemon checks the health of the following hardware components:
  - FANs on the motherboard
  - Motherboard Temperature sensors
  - Motherboard Voltage sensors
  - Motherboard Current sensors
  - Presence of power source
  - Health of Hard drives
  - RAID controller status
  - RAID controller battery status
- The State of Health daemon also checks if the "dom0" (aka host) root files system is read-write enabled and if "dom0" has at least 6% of free memory. A controlled failover is initiated if the root file system becomes read-only or if the active dom0 has less than 6% of free memory.
- If "dom0" cannot reach an administered network destination (via an ICMP ping). The default frequency and the network destination can be changed when HA is enabled. By default it takes about 10 seconds to detect network failure. The 10 second delay in detecting the network failure may cause sockets to drop or lose messages over the network. Applications running on the guest VM are responsible for recovering from such data loss.
- For controlled failover, CDOM and the Services VM (if configured) will be migrated to the new active server.
- On a controlled failover, a memory protected VM will be available for service on the new server within 500 milliseconds.

Health objects monitored by the State of Health daemon are listed in the following table.

Health object Type	Frequency (seconds)	Alarm threshold	FAFF* (seconds)	Comments
Temperature	60	3	20	
Fans (motherboard)	60	3	20	
Voltage	5	3	1	
Current	5	3	1	
Power source	60	3	20	
Hard drive	5	3	1	
Host's root file system	5	3	1	Check if root file system is writable
Host's free memory	5	3	1	Check if there is at least 6% free memory available
Raid Status	5	3	1	

Raid Battery Status	5	3	1	
Network Status**	2	5	2	Pings host's gateway for eth0 network by default.

*FAFF=Frequency After First Failure

3.2.2 Almost seamless failover in the event of sudden failures for MPHA protected VMs

Uncontrolled failover can happen if the current active server fails suddenly or is not reachable over the crossover link or via the switched IP network. The previously check-pointed (50-100ms old) VM is then activated on the new active server. The VM may have to recover from its lost state/sockets.

3.2.3 Automatic recovery from split brain condition

When active and standby servers fail to reach each other over all configured paths, each server assumes the other server is dead. And thus, if both servers are powered and healthy, both will be running in active mode and will be operating in a split brain condition. Since MPHA ensures each server cannot reach each other while in the split brain condition, only one server can be connected to the network and therefore only one VM can be on the network. Once the network heals, the Arbiter will choose the server that was connected to the network (over all network interfaces) for the longest period of time, to stay active and the other server will be forced to become the standby server.

3.2.4 Adaptive check-pointing

MPHA protected VMs are tested for a given capacity. If the advertised capacity of a VM is exceeded, it is possible that the VM could be suspended for a longer period of time in order to copy the higher number of dirty pages a higher load would create. This in turn can affect the way the VM provides service. In such cases, memory replication will be abandoned temporarily and an error is logged. If the overload condition persists, an alarm will be generated.

- Replication will disengage temporarily if replication interferes with the VM's ability to provide service or if the VM is used over its capacity.
- In some cases certain requests related to memory replication may timeout and therefore abandon replication temporarily. Memory replication will re-engage within 20 seconds automatically. An error will be logged when this happens.
- If an uncontrolled failover occurs when replication is not fully engaged, the protected VM will be rebooted on the new active server.

3.2.5 No limit on CPUs allocated to Application Virtual Machine

Unlike other HA solutions in the market MPHA does not limit the number CPUs that can be allocated to the VM.

3.3 What does MPHA provide?

MPHA mostly provides protection against hardware failures and to some extent network failures. It does not protect against software failures within the protected VM.

- Protection against hardware failure
  - In case of sudden hardware failure (uncontrolled) the protected VM will be running on the new active server from its last check-pointed state within ~1 second
  - In case of controlled failover or if hardware health degrades, the protected VM will be running on the new active server, without loss of any data within ~500ms
- Protection against network failure to some extent, as detection for network failure can take longer (default 10 seconds).
- Does not protect against VM software failures
- Costs approximately 30-50% of the CPU used by the protected application VM

3.4 Power requirements for System Platform HA systems

Each server in the HA pair must have two power modules. Each of these power modules must be connected to a power source. This configuration ensures there is no single point of failure in the power infrastructure powering these servers. Valid configurations are:

	Server1	Server2
PM1	PS1 (with UPS)	PS1 (with UPS)
PM2	PS2 (with UPS)	PS2 (with UPS)

	Server1	Server2
PM1	PS1 (with UPS)	PS1 (with UPS)
PM2	PS1 (with UPS)	PS1 (with UPS)

	Server1	Server2
PM1	PS1	PS1
PM2	PS2	PS2

PM1/PM2: Power module 1 and 2 respectively, representing the two power modules installed on each server.

PS1/PS2: Power Source1 and 2 respectively. Two power sources mean each power source is connected to a different power grid or generator.

UPS: Uninterrupted power supply. Expectation is that the UPS kicks in when needed in such a way that there is power to the server continuously without a power glitch. The UPS does not have to be per-server basis. An UPS servicing an entire datacenter will suffice as long as it meets "without a glitch" requirement.

3.5 Effect of uncontrolled failover on Application VM clients

During an uncontrolled failover, the standby server takes about 450-500 milliseconds to detect that the active server is not reachable. Once the standby server determines that the active server is not reachable, it activates the replicated (check-pointed) VM within another ~500 milliseconds. The MPHA protected VM will be running within 1 second. Since this newly activated VM could be about ~100 milliseconds behind in time with respect to the entities the VM was communicating with, the VM and its clients are responsible for recovering from the following:

- Potential loss of TCP sockets
- The VM and its clients are responsible for recovering from any lost state
- The VM may have lost data committed to disk since the last checkpoint.
- The VM and its clients may lose some events and some requests may timeout.

Note that the above scenario could happen only in case of "Uncontrolled" failover. The following are some of the conditions where an uncontrolled failover may occur:

- Memory chips gone bad on the current active server
- All disks have failed (Note SP has RAID configured)
- Motherboard failure
- CPU failure
- Kernel panics (a kernel bug)
- Both power supplies/sources go bad
- Other causes that could cause the active server unusable

3.6 Effect of Uncontrolled failover on Application Enablement (AE) Services Clients

When an AE Services VM that is ~100 milliseconds old, is activated on a new active server, the AE Services VM will be ready for service within 1 second of the previous active server disappearing. The ability for the clients to continue to receive service seamlessly depends on whether clients (and the SDKs they depend on) can reconnect to the AE Services VM if sockets drop.

- Clients and Communication Manager are ~100 milliseconds ahead of time with respect to the newly active AE Services VM.

- If there was any TCP traffic during the last incomplete check point (~100 milliseconds) the TCP sockets between the AE Services VM and its clients may drop.
- Clients may lose some events.
- Data written to the disk since the last checkpoint may be lost.
- To preserve the Transport link, Communication Manger must have CM6.2 Service Pack 2 (or newer) installed.

Effect on DMCC and TR87 clients (minimal)

- DMCC clients can re-connect to the AE Services VM and resume the session. DMCC clients may lose events that were generated during failover. For first party call control, Communication Manager will refresh the current state of the display and lamp state associated with various buttons.
- For TR87 clients, the SIP dialogs can survive a socket drop. Therefore all association created will remain intact after failover. SIP Dialogs that were in transient state in the past 100 milliseconds could be affected.

Effect on TSAPI, CVLAN and DLG clients

- If TSAPI, CVLAN and DLG client sockets drop, the client applications have to reestablish all associations.
- In the future, the TSAPI and JTAPI SDKs will re-establish these sockets upon a socket failure and preserve the previous session. It will also launch an audit and recover from the lost state in the AE Services VM.

4 Fast Reboot High Availability (FRHA)

4.1 Overview

FRHA is similar to GRHA with the limitation of a crossover cable and the use of a single IP address on System Platform. The AE Services 6.1 and later release on System Platform provides HA relative to earlier releases of AE Services, as well as the AE Services 6.1 Software-Only and Bundled offers. This configuration monitors the server nodes for loss of network connectivity and hardware failure events. This information is used to detect faults and decide when to failover from the active node to the standby node in the server cluster. The AE Services VM on the standby node is restarted when a failover event occurs. This feature enables AE Services to continue to provide service to client applications with reduced downtime when a System Platform hardware failure event occurs. In addition to this, System Platform will restart the

virtual machine in which AE Services is running if it stops responding as a result of a software fault condition.

This figure illustrates an HA cluster of AE Services on System Platform using FRHA communicating to Communication Manager through the PE interface.

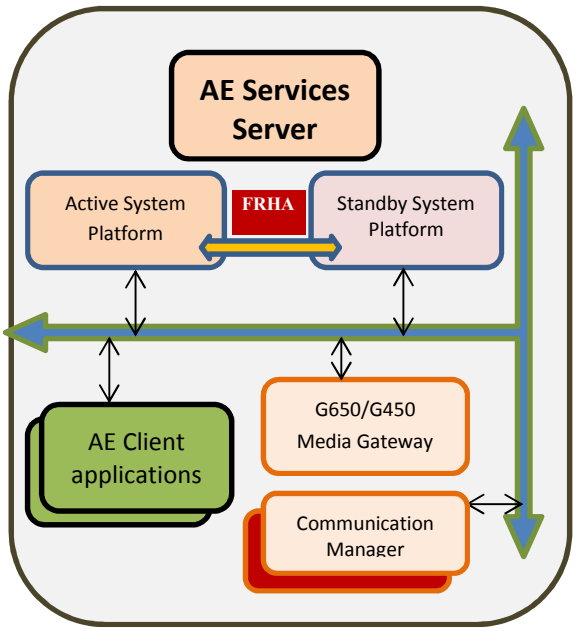

5 The AE Services and High Availability

5.1 Device, Media and Call Control Service (DMCC)

Avaya recommends that applications reestablish DMCC sessions and verify that all associations (monitors, registrations) are still active after a network interruption. This is true whether or not the interruption was caused by an AE Services server failover.

In addition to the AE Services server failover feature, the DMCC service provides recovery from a software fault or a shutdown that does not allow the DMCC Java Virtual Machine (JVM) process to exit normally. The DMCC Service Recovery feature is available on all AE Services offers: Software-Only, Bundled, VMware and System Platform (SP). When the DMCC JVM process is restarted after an abnormal exit, the DMCC service is initialized from persisted state information on the hard disk. This persisted state information is saved during normal operation, and represents the last known state of the DMCC service prior to a JVM abnormal exit. The

state information includes session, device, device/call monitoring and H.323 registration data. Following the restart of the DMCC JVM, the persisted information is used to re-create the sessions, device IDs, monitors and H.323 registrations that existed just prior to the software fault. Note that only H.323 registrations that use the Time-To-Service feature on Communication Manager can be recovered.

From a client application's point of view, the DMCC recovery appears as a temporary network interruption that requires the client to reestablish any disconnected sessions. When the client application reestablishes the session, the DMCC service will send events to the client for any resources that could not be recovered. These events may include "monitor stopped" and/or "terminal unregistered" event messages, and will enable the client to determine what, if anything, needs to be restored (using new service requests). Otherwise, the client application may continue to operate as usual.

Note that an AE Services server failover or a restart of the DMCC JVM results in the teardown and re-creation of the H.323 endpoint registrations, but is limited only to the AE Services server. The Communication Manager is unaware that this is taking place, and sees the endpoints as still registered. Be aware that this may have an effect on any calls in progress for these endpoints. If the client application specified "server-media" dependency mode for a call, then the call (and any recording on the call) will be terminated. Alternatively, if the client application specified "client-media" dependency mode for the call, then the call should survive the failover (or restart). However, it is possible that some state changes for the endpoint and its call may have been missed by the AE Services server during the failover. For example, if the far-end hangs up the call at the same moment that a failover occurs, then it is possible that any, or all, of the resultant "HookswitchEvent", "MediaStopEvent", "DisplayUpdatedEvent" and "LampModeEvent" messages could be lost.

5.2 TSAPI, CVLAN, DLG and Transport Services

No runtime state information is persisted for these services. The client application must restore any state that existed before the service was restarted.

5.3 Recommendations

The following items are recommended:

- Communication Manager should be configured for H.323 registration using the Time-To-Service feature.
- AE Services 6.1 and later should use the PE interface – even in survivable server environments. PE connections offer reduced complexity:

CLAN Path: AES server → CLAN → TDM bus → IPSI → CM Server

PE Path: AES server → CM Server

- A local HA cluster of AE Services 6.1 and later on System Platform or VMware servers is used.
- An application that uses the Device, Media and Call Control (DMCC) service should keep trying to reestablish the DMCC session when it loses its socket communication link to the AE Services server because the DMCC runtime state is preserved during a failover. This applies to all AE Services configurations.
- An application that uses the CallVisor Local Area Network (CVLAN), Definity Local Area Network Gateway (DLG) or Telephony Service Application Programming Interface (TSAPI) service should reestablish its socket connections and its monitors/associations if it loses the socket connection to the service on the AE Services server because no runtime state is preserved for these services. TSAPI applications also need to reestablish route registrations.

Avaya recommends that all applications in a survivable server configuration connect to a local AE Services 6.1 (or later) server that, in turn, is connected to either the media server at the main site or a media gateway/survivable server at the remote site. In this configuration, the applications and associated AE Services servers at the remote sites are always active and are supplying functionality for the local resources at the remote site. This type of configuration ensures the most seamless survivability in an enterprise survivable configuration.

For more information about Communication Manager's availability assessment and methodologies, see the white paper, ***Avaya Communication Manager Software Based Platforms: High Availability Solutions, Avaya Media Servers and Gateways***, available on the Avaya Support Web site, http://support.avaya.com

6 Communication Manager and High Availability

6.1 Overview

Avaya Aura® Communication Manager (CM) provides survivable core servers (also known as ESS) and survivable remote servers (also known as LSP) for failover from the primary media server. This feature provides the ability for media gateways, endpoints, application servers like AE Services and its applications to continue their operations without a major outage. Survivable server configurations have been supported, with some limitations, since the AE Services 3.0 release. However, important enhancements were added to both the AE Services 6.1 release and the Communication Manager 6.0 release to better support these configurations.

Since the AE Services 6.1 and Communication Manager 6.0 release, switch connections on both the Control Local Area Network interface cards (CLANs) and Processor Ethernet (PE) connections are fully supported in all survivable server configurations. Additionally, any Device, Media and Call Control (DMCC) endpoints registered to the primary switch using the Time-To-Service feature (TTS) will automatically re-register to a survivable server. DMCC endpoint registrations not using the Time-To-Service feature will be unregistered when the Communication Manager fails over to a survivable server.

AE Services 6.1 and later allows the Communication Manager survivable server nodes to be administered within a switch connection, in a priority order, along with the PE IP address of those nodes. When used in conjunction with Communication Manager 6.0 or later, it provides the means to deterministically control the AE Services server connectivity in failover situations.

6.2 Communication Manager Survivable Servers

Communication Manager offers two survivability options: survivable core and survivable remote. Survivable core servers ensure business continuity in the event of connection failure or events leading to total failure of the main server complex, such as natural disaster. Survivable remote servers enhance redundancy for branch gateways within networks. Survivable remote servers take over segments that have been disconnected from their primary call server and provide those segments with Communication Manager operation until the outage is resolved.

For more information about Communication Manager's survivability deployment, see: ***Avaya Aura® Solution Deployment,*** available on the Avaya Support Web site, http://support.avaya.com

6.2.1 Survivable Core Server (ESS)

The survivable core server (ESS) provides survivability to an Avaya configuration by allowing backup servers to be placed in various locations in the customer's network. AE Services can utilize either CLAN or PE connectivity to a survivable core server node. There are some basic and fundamental differences between CLAN and PE connectivity. First, a CLAN always resides in a port network. If AE Services is able to establish a session with a Communication Manager node via a CLAN connection, then AE Services can always conclude on its own that that Communication Manager node is active (by definition, if a port network is registered to a Communication Manager node, it is active). It does not have to rely on Communication Manager to send that information. On the other hand, a survivable server node will always accept a PE connection from AE Services, even if it is not active. In this case, AE Services must rely on Communication Manager to send its server role and server state information, which requires Communication Manager 6.0 or later, to know if that node is active (see section 6.7.1).

Second, the **port network** controls to which Communication Manager node AE Services is connected via a CLAN. In disaster recovery scenarios, AE Services has no control over which Communication Manager node will provide service to that port network, and therefore has no control over which Communication Manager node will provide service to AE Services (over that CLAN AEP connection). In this case, Communication Manager's administration controls under what conditions and to which Communication Manager node the port network will connect. Conversely, PE connections are made directly to Communication Manager nodes, and AE Services can therefore directly control (via its own administration) from which Communication Manager node it requests service. Depending on the system topology and disaster recovery requirements, either type of connectivity can be successfully utilized.

6.2.2 Survivable Remote Server (LSP)

A media gateway like the G250, G350, G430, G450 or G700 can be controlled by a survivable remote server (LSP) running Communication Manager if the primary Communication Manager media server is unavailable or down. Typically, survivable remote servers are configured on (remote) media gateways at branch offices so that those media gateways can get service in case the connectivity to the main site is down (e.g. WAN connectivity failure as shown in Figure 2 or site destruction as shown in Figure 3). Once the local media gateways detect a failure of connectivity to the primary media server, they register with the Communication Manager running on that survivable remote server, which becomes active.

6.2.3 Survivable server – PE Connectivity

PE connectivity can be used with both survivable core and survivable remote servers, and the behavior is same for both. PE connectivity is only supported with survivable server configurations when AE Services 6.1 or later servers are connected to Communication Manager 6.0 or greater nodes. Any other combination of AE Services and Communication Manager results in pre AES 6.1 behavior.[1] Section 6.7.1 provides the rational for this restriction.

The list below describes the behavior of the AE Services:

a. **DMCC (Device and Media Control) Service**
   As long as the application is configured to connect to media gateways that are local to the AE Services server's site, recovery with a survivable server should be very straightforward. For DMCC endpoints registered using the Time-To-Service feature, the endpoints will be automatically re-registered to the survivable server, with no notification to the application necessary. However, the application will receive an unregistered event for each DMCC endpoint that is not using the Time-To-Service

[1] A work around is provided for AE Services 6.1 with non-CM 6.0 or greater switches when PE is used with a single ESS or LSP node. See ***PSN 3156u - PE support for ESS and LSP scenarios*** for more details.

feature when connectivity is lost to the primary Communication Manager media server. There is a slight possibility that an endpoint using the Time-To-Service feature will fail to re-register, in which case an unregistered event for that endpoint will also be sent to the client application.

At this point, the application should begin attempts to re-register the DMCC endpoints (that failed the automatic re-registration) with the same IP address(es) it was using before. Note that it takes a little over 3 minutes for the media gateway to connect to a survivable server. For this reason, it is recommended that the application keep trying to register with the same media gateway (through the AE Services server) for that amount of time before it tries to register with a survivable remote server (if one exists). When the media gateways connect with the survivable server, the registration attempts will begin to succeed. After the application has successfully re-registered all DMCC endpoints, it should reestablish its previous state and resume operation.

b. CallInformation Services within DMCC, Call Control Services within DMCC, and all other CTI services

The CallInformation and Call Control services within DMCC and all other CTI Services (TSAPI, CVLAN, DLG and JTAPI) use the Transport (AEP) link to communicate with Communication Manager. The transport links (Switch Connections) on each AE Services server should be administered to communicate only with PEs for Communication Manager media servers that are local to the AE Services server's site. If the system is configured in this fashion, the application/AE Services server will not have to take any unusual action to recover in the event that a gateway loses connectivity to the primary Communication Manager node and transitions to a survivable server.

If a media gateway loses connectivity to the primary server for an extended period of time (configurable on Communication Manager), it will register with the local survivable server. Within 5 seconds of that registration, that survivable server will inform AE Services 6.1 (or later) that is has transitioned from idle to active. AE Services 6.1 (or later) will re-evaluate its current session. If this survivable server node has a higher precedence than the current Communication Manager node in use (or if there is no current session), it will be used. If an AE Services server changes Communication Manager nodes, it will notify any connected applications of this event via a LinkDownEvent, for DMCC CallInformationServices, or a CTI link down indication, for CTI services. For Call Control Services within DMCC, Avaya recommends that applications add a CallInformationListener and look for a LinkDownEvent for indication that connectivity to the primary site is down. (In future releases, Call Control Services clients will receive a MonitorStop request for all call control monitors if the link is lost

to the main site.) Depending on the CTI application programming interface (API), clients will receive an appropriate event when the connectivity to the primary site is down. CVLAN clients will receive an "abort" for each association. TSAPI clients will receive a CSTAMonitorEnded event if the client is monitoring a device and/or a CSTASysStatEvent with a link down indication if the client is monitoring system status. TSAPI clients will also receive a CSTARouteEnd event for any active routing dialogs, and a CSTARouteRegisterAbort event for any registered routing devices. Avaya JTAPI 5.2 and later clients will receive a PROVIDER_OUT_OF_SERVICE event if the client has ProviderListeners. Otherwise, a ProvOutOfServiceEv event will be received if the client has ProviderObservers. DLG clients will receive a link status event with a link down indication and a cause value.

The AE Services server will then automatically notify the application that the CTI link is back up, and the application can begin to resume normal operations. Since there is no run-time state preserved on a transition from a primary Communication Manager media server to a survivable server (as there is with an interchange on a duplicated Communication Manager media server pair) all application state must be reestablished. Note that, from the AE Services server's and application's perspectives, the failure scenario and recovery actions appear ***exactly*** the same as a long network outage between the AE Services server and the gateways.

6.2.4 Survivable core server (ESS) – Non PE Connectivity

This case is essentially the same as in previous AE Services releases, with one exception: fragmentation (see section 6.7). Since the AE Services server is connected through a CLAN, it will automatically move to the new Communication Manager node when the gateway (port network) in which that CLAN resides moves to the new Communication Manager node. In the event that the AE Services server is connected through multiple CLANs to more than one Communication Manager node, the rules described in the fragmentation section (section 6.7) govern which Communication Manager node will be used for the active session.

The list below describes the behavior of the AE Services:

a. **DMCC (Device and Media Control) Service**
   As long as the application is configured to connect to CLANs in the local media gateways, recovery with a survivable core server should be very straightforward. For DMCC endpoints registered using the Time-To-Service feature, the endpoints will be automatically re-registered to the survivable core server (ESS), with no notification to the application necessary. However, the application will receive an unregistered event for each DMCC endpoint that is not using the Time-To-Service feature when

connectivity is lost from the local media gateways (like G650) to the primary Communication Manager media server. There is a slight possibility that an endpoint using the Time-To-Service feature will fail to re-register, in which case an unregistered event for that endpoint will also be sent to the client application.

At this point, the application should begin attempts to re-register the DMCC endpoints (that failed the automatic re-registration) with the same IP address(es) it was using before. Note that it takes a little over 3 minutes for the media gateway (like G650) to connect to a survivable core server (ESS). For this reason, it is recommended that the application keep trying to register with the same CLAN (through the AE Services server) for that amount of time before it tries to register with a survivable remote server (if one exists). When the media gateways (like G650) connect with the survivable core server (ESS), the registration attempts will begin to succeed. After the application has successfully re-registered all DMCC endpoints, it should reestablish its previous state and resume operation.

b. CallInformation Services within DMCC, Call Control Services within DMCC, and all other CTI services

The CallInformation and Call Control services within DMCC and all other CTI Services (TSAPI, CVLAN, DLG and JTAPI) use the Transport (AEP) link to communicate with Communication Manager. The transport links (Switch Connections) on each AE Services server should be administered to communicate only with CLANs in the media gateways that are local to the AE Services server's site. If the system is configured in this fashion, the application/AE Services server will not have to take any unusual action to recover in the event that a gateway loses connectivity to the primary Communication Manager media server and transitions to a survivable core server (ESS).

If a media gateway loses connectivity to the primary Communication Manager server for an extended period of time (more than 10 seconds), all AEP sockets that are established through CLANs resident in that media gateway will drop. If an AE Services server loses all of its AEP connections for more than 30 seconds, its session will drop, and the CTI links will go down. AE Services will notify any connected applications of this event via a LinkDownEvent, for DMCC CallInformationServices, or a CTI link down indication, for CTI services. For Call Control Services within DMCC, Avaya recommends that applications add a CallInformationListener and look for a LinkDownEvent for indication that connectivity to the primary site is down. (In future releases, Call Control Services clients will receive a MonitorStop request for all call control monitors if the link is lost to the main site.) Depending on the CTI API, clients will receive an appropriate indication when the connectivity to the primary site is down. CVLAN clients will receive

an “abort” for each association. TSAPI clients will receive a CSTAMonitorEnded event if the client is monitoring a device and/or a CSTASysStatEvent with a link down indication if the client is monitoring system status. TSAPI clients will also receive a CSTARouteEnd event for any active routing dialogs, and a CSTARouteRegisterAbort event for any registered routing devices. Avaya JTAPI 5.2 and later clients will receive a PROVIDER_OUT_OF_SERVICE event if the client has ProviderListeners. Otherwise, a ProvOutOfServiceEv event will be received if the client has ProviderObservers. DLG clients will receive a link status event with a link down indication and a cause value.

The AE Services server will then automatically attempt to reestablish the AEP links. Note that it takes a little over 3 minutes for the media gateway (like G650) to connect to a survivable core server (ESS). Once the media gateway has registered with the survivable core server (ESS), the AE Services server will succeed in establishing its AEP links very soon thereafter (after around 30 seconds). As soon as an AEP link is established, the application will be notified that the CTI link is back up, and the application can begin to resume normal operations. Since there is no run-time state preserved on a transition from a primary Communication Manager media server to a survivable core server (as there is with an interchange on a duplicated Communication Manager media server pair) all application state must be reestablished. Note that, from the AE Services server’s and application’s perspectives, the failure scenario and recovery actions appear ***exactly*** the same as a long network outage between the AE Services server and the gateways.

6.3 Communication Manager survivable server considerations for AE Services deployment

Both CLAN and PE connectivity can be used in survivable environments. To an extent, by using CLAN connectivity, primary control resides with the Communication Manager administration, since Communication Manager controls the order in which the media gateways containing those CLANs will seek service from the various configured survivable servers. Therefore, AE Services will only be able to choose from the set of survivable servers that are providing service to those CLANs. However, if AE Services uses PE connectivity to the survivable servers, then it will know the state of all of the survivable servers to which it can connect, and, therefore, its survivability administration provides the primary control for which survivable server will be used during an outage. Note, however, that Communication Manager administration and media gateway administration still control which survivable servers provide service during outages.

6.3.1 Communication Manager media server resource recovery

Communication Manager provides a number of options for the primary media server to recover port networks and/or media gateways after an outage. These resources can be reclaimed on demand or automatically (based on numerous criteria like time of day or usage). For example, a survivable remote server (LSP) is typically left controlling local media gateways as long as there are active calls on those media gateways, even after the primary Communication Manager media server comes back on line. In this case, it would be detrimental if a local AE Services 6.1 (or later) server connected to that survivable remote server (LSP) moved back to the primary Communication Manager node while that survivable remote server (LSP) was still controlling the local media gateways. Therefore, AE Services 6.1 (or later) provides flexibility for determining when to return back to the primary (main) Communication Manager server after an outage. A checkbox is provided on the Survivability Hierarchy OAM screen that will control AE Services 6.1 (or later) behavior when the primary Communication Manager 6.0 (or later) node comes back online while AE Services has an active session established with a Communication Manager 6.0 (or later) survivable server node. By default, AE Services 6.1 (or later) will remain talking to that survivable server node as long as it is active. As soon as that node becomes idle (either as a result of all port networks and media gateways unregistering from that node or the session to that node is lost), AE Services 6.1 (or later) will, once again, automatically return to the primary Communication Manager node. All previous versions of AE Services automatically and immediately returned to the primary Communication Manager server as soon as it could connect to it, which often required manual intervention (to disable the connection to the primary Communication Manager node) to prevent unexpected loss of service. Also note that if AE Services 6.1 (or later) is connected to a pre Communication Manager 6.0 server, it will automatically and immediately return to the primary Communication Manager server as soon as it can connect to it (AE Services 6.1 (or later) cannot tell when a pre Communication Manager 6.0 survivable server node is idle or active).

6.4 Recovery Scenario Examples

The following example configuration shows an enterprise implementation over four locations utilizing both CLAN and PE connectivity with survivable core and survivable remote servers. Note that this example has been constructed to illustrate various failover scenarios within a single example, and therefore is not intended to (and does not) represent a typical real life installation. However, the principles of recovery discussed between any two sites do apply to typical configurations, and these principles can be extended to suite actual needs. In the figures below, green lines represent active communication paths (i.e., currently being used to provide service), yellow lines represent standby communication paths (i.e., connection is established, but not being used to provide service yet), and red lines represent broken

communication paths (i.e., previously established communication path that has been lost as a result of an outage).

6.4.1 Normal Operation

Figure 1 is an illustration of a sample enterprise network configuration.

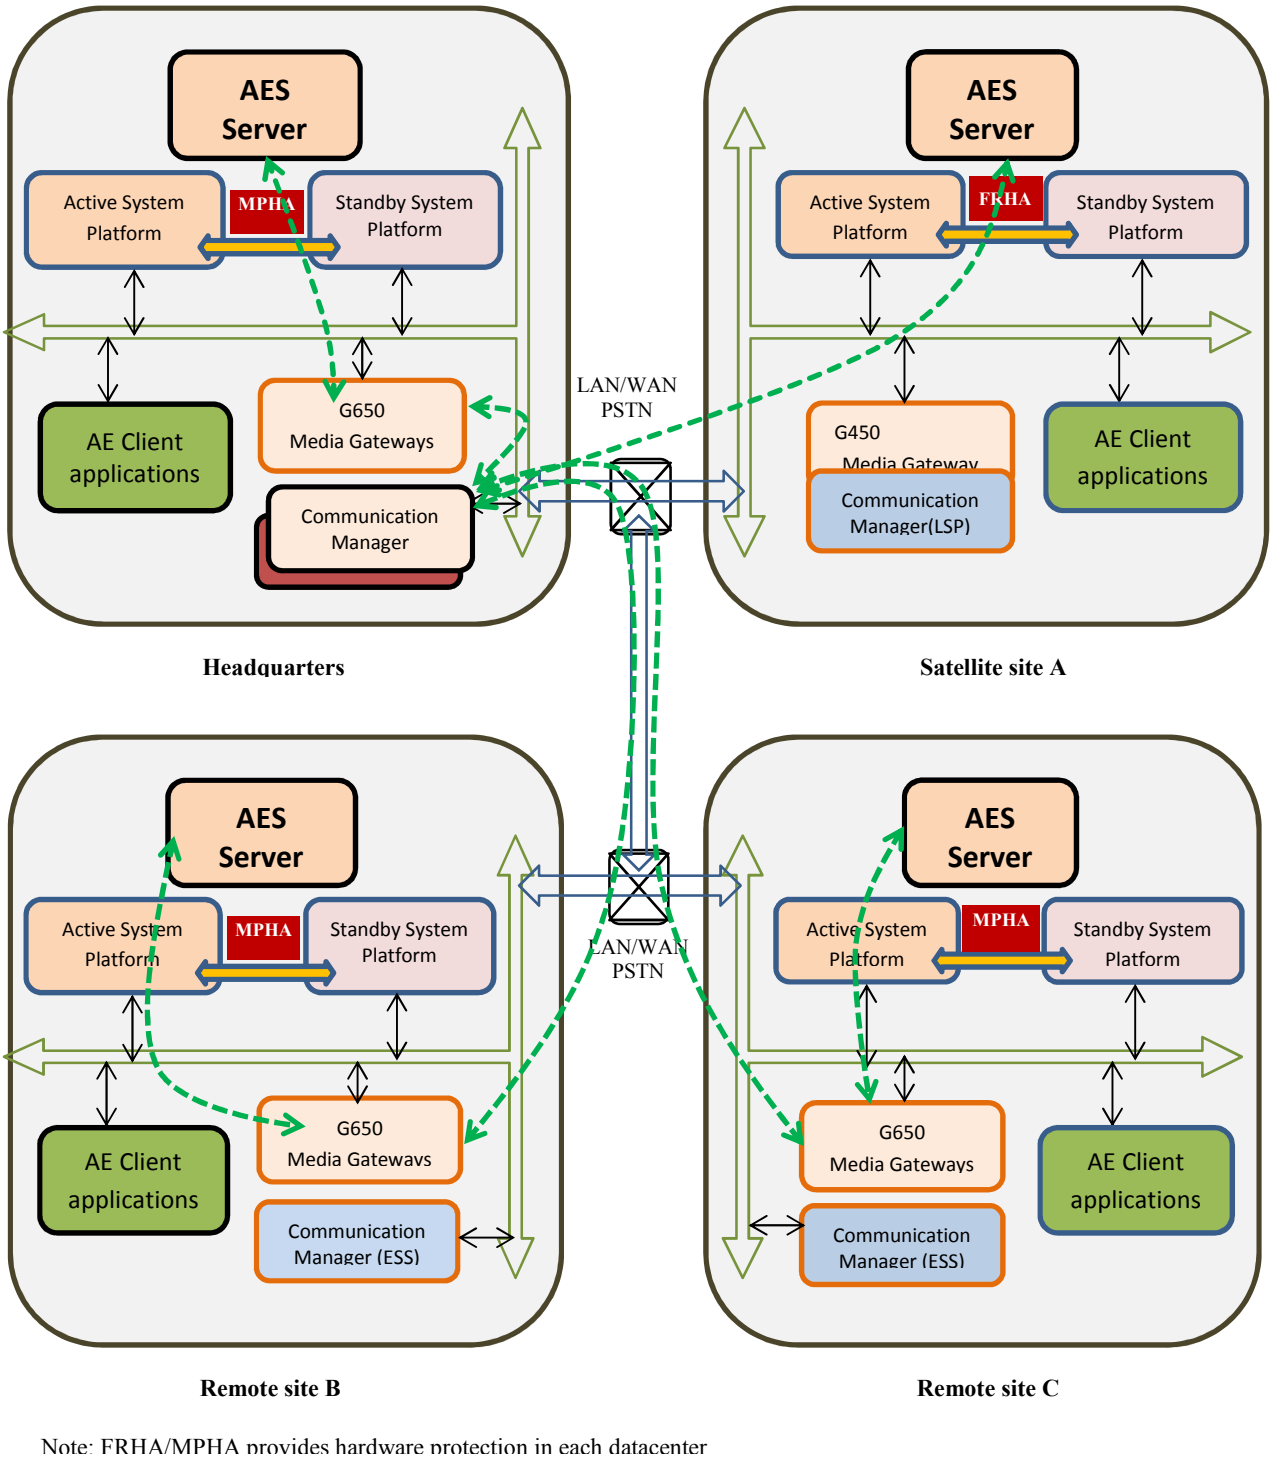

Figure 1: Normal Operation

Each site has a high availability (HA) AE Services cluster and associated application. The headquarter site has the primary S8800 media server pair and G650 media gateways. The AE

Services server at the main site is connected via CLANs to Communication Manager. Satellite site A (e.g., branch office) has a G450 media gateway with a survivable remote server (LSP). The AE Services server at satellite site A is connected to both the primary Communication Manager server and the survivable remote server (LSP) via PE connections. Remote site B has an S8800 survivable core server (ESS) and G650 media gateways. The AE Services server at remote site B is connected via CLANs to Communication Manager. Remote site C has an S8800 survivable core server (ESS) and G650 media gateways. The AE Services server at remote site C is connected to the primary Communication Manager and the survivable core servers at remote sites B and C via PE connections. Furthermore, all AE Services servers are configured to stay on the survivable server as long as they are providing service.

Avaya recommends that all applications have a local AE Services server. In this configuration, the applications and associated AE Services server at the remote sites are always active and are supplying functionality for the local resources at the remote site. As described in this document, this type of configuration ensures the most seamless service in a survivable configuration.

6.4.2 Wide Area Network (WAN) Outage

In case of a WAN outage (as shown in Figure 2), each remote site becomes independent and provides service without major interruption to endpoints and applications.

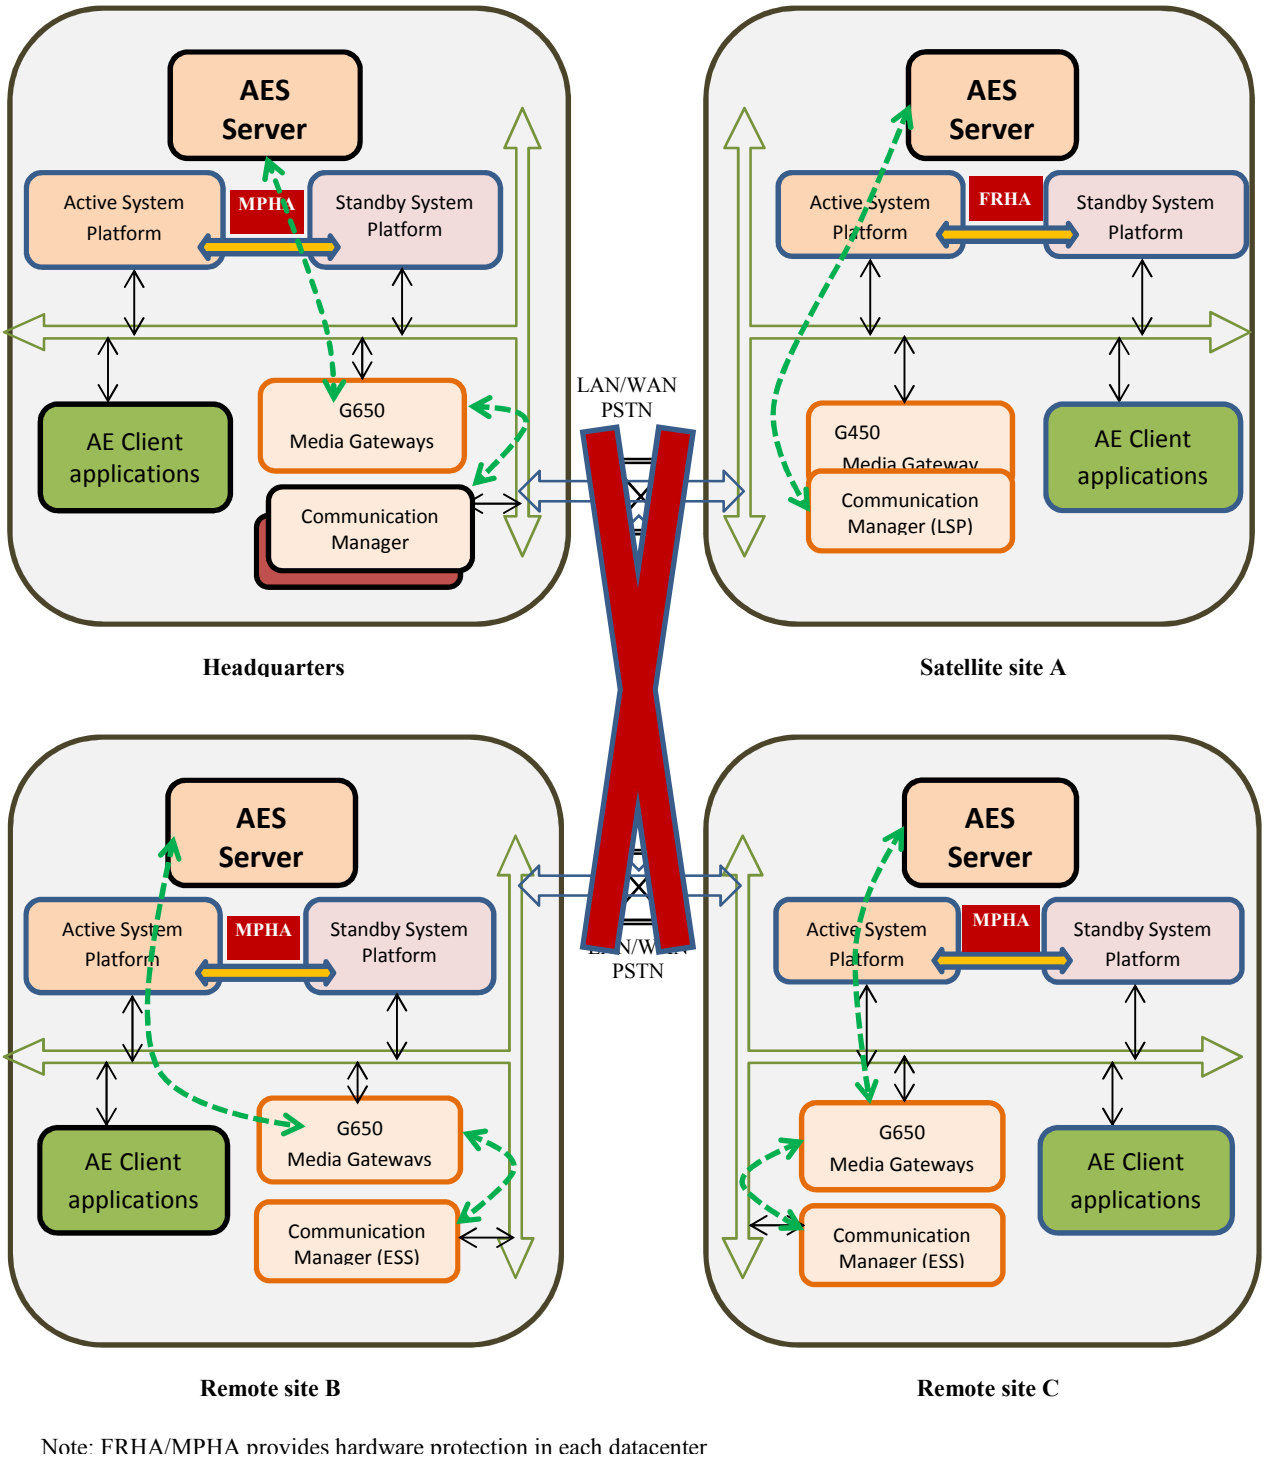

Figure 2: Wide Area Network (WAN) Outage

Satellite site A with a G450 media gateway will have the survivable remote server (LSP) go online and the G450 media gateway will connect to that local survivable server. It is

recommended to configure the primary search list of the G450 media gateway such that it contains CLANs (or PE) of only one site (i.e. headquarters in this case). The secondary search list should contain the survivable remote server (LSP) at the local site. The AE Services server will detect connectivity failure with the primary site (headquarters) and will automatically start using the survivable remote server (LSP) to provide service.

The G650 media gateways at remote site B will connect to the local survivable core server (ESS) in case of a WAN outage. The AE Services server at this site will automatically connect with the survivable core server (ESS) through the G650 media gateways. All of this will be transparent to the AE Services server and its applications except for what will appear to be a brief network outage.

The G650 media gateways at remote site C will connect to the local survivable core server (ESS) in case of a WAN outage. The AE Services server at this site will detect connectivity failure with the primary site (headquarters) and will automatically start using the survivable core server (ESS) to provide service.

The site at the headquarters will continue to function as it did previously in case of a WAN outage.

Note: Each of the remote sites and the headquarter site will not be able to access each other's resources during a WAN outage. Also, at each of the remote sites, this will be transparent to the AE Services applications except for what will appear to be a brief network outage (described in detail section 6.4.4).

6.4.3 Primary Site Destruction

If the main headquarters site is completely down but the WAN is functional, (as shown in Figure 3 below), the remote sites will behave similar to the WAN outage scenario described above, but with one important exception. With the survivable core server (ESS) feature, the system will attempt to stay as "whole" as possible. Since the WAN is still intact, all of the G650 gateways end up being controlled by the same survivable core server (ESS) at Remote Site B. Since the applications and AE Services servers were configured to support only the local resources at the remote sites, the applications continue to function the same whether the sites operate independently (WAN failure) or jointly (normal operation or site destruction at headquarters). Also note that satellite site A's behavior remains unchanged as well (i.e. it still uses its own survivable local server (LSP), instead of the survivable remote server (ESS) at remote site B, that is because the assigned CLAN for G450 is in the main site).

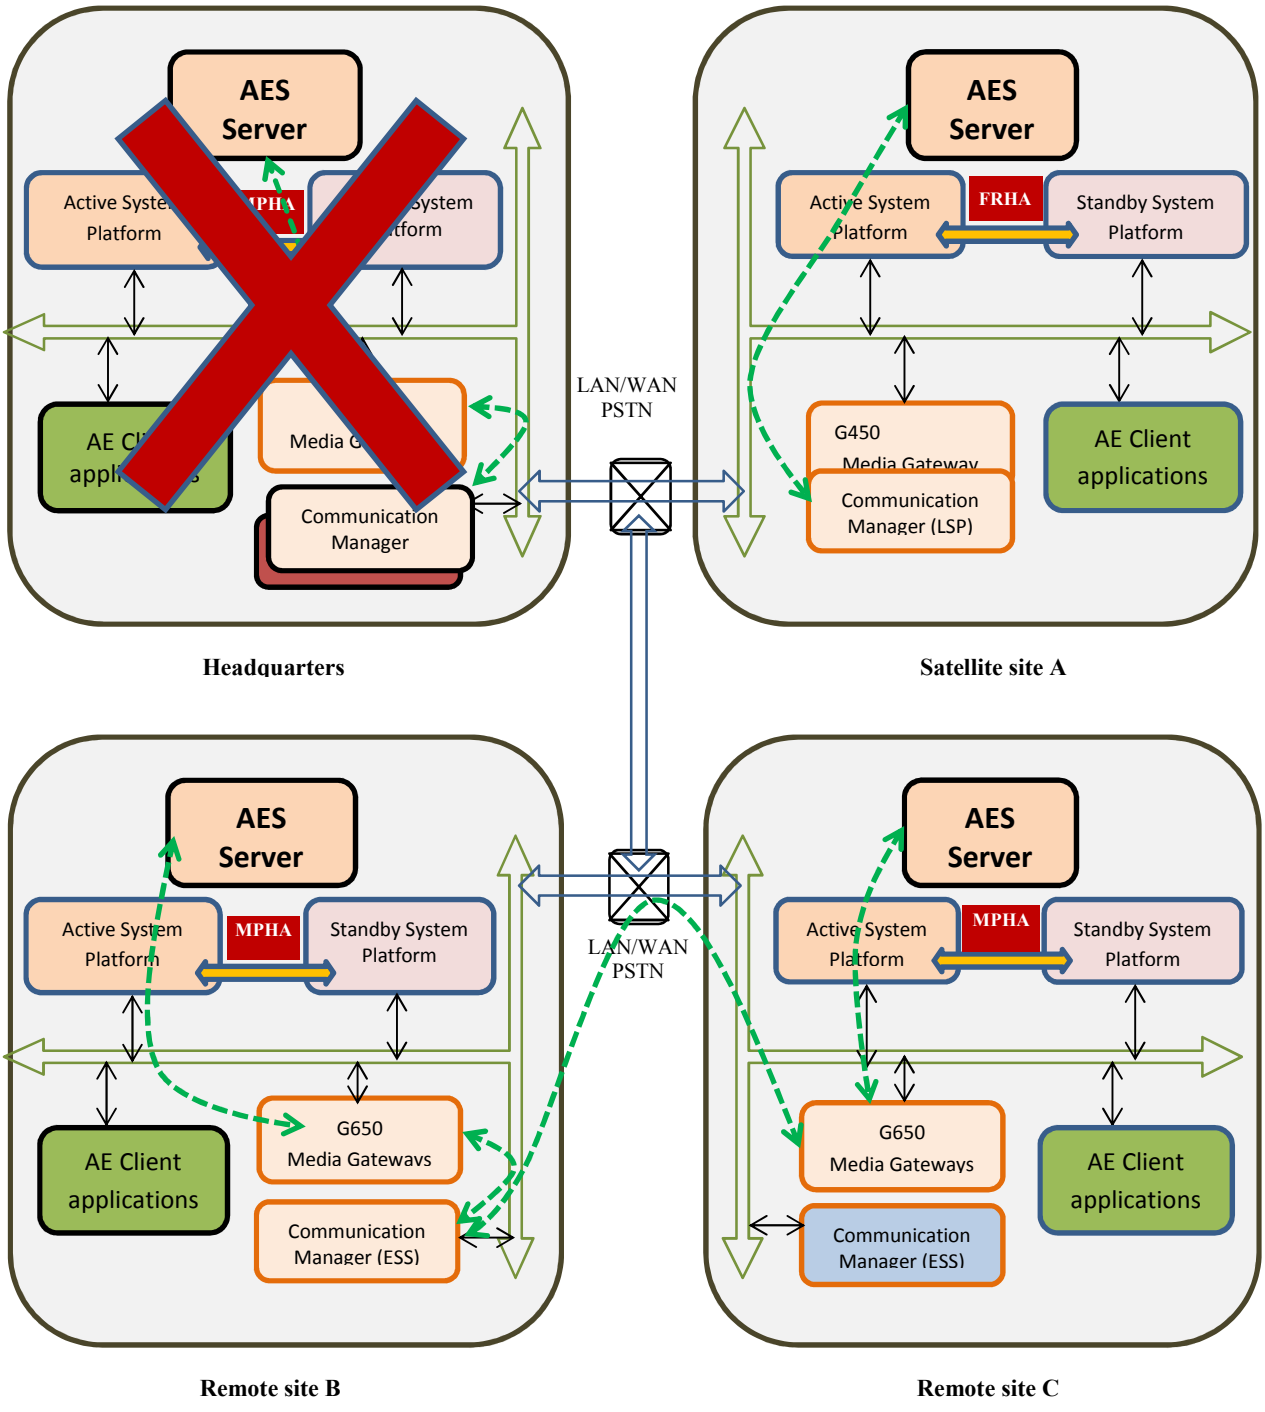

Figure 3: Site Destruction

6.4.4 Secondary Failure

If the survivable remote server (ESS) at remote site B fails (secondary failure) while the main headquarters site is still completely down but the WAN is functional, (as shown in Figure 4 below), the remote sites will behave **exactly** the same as the site destruction outage scenario described above. The only difference is which survivable core server (ESS) is providing service.

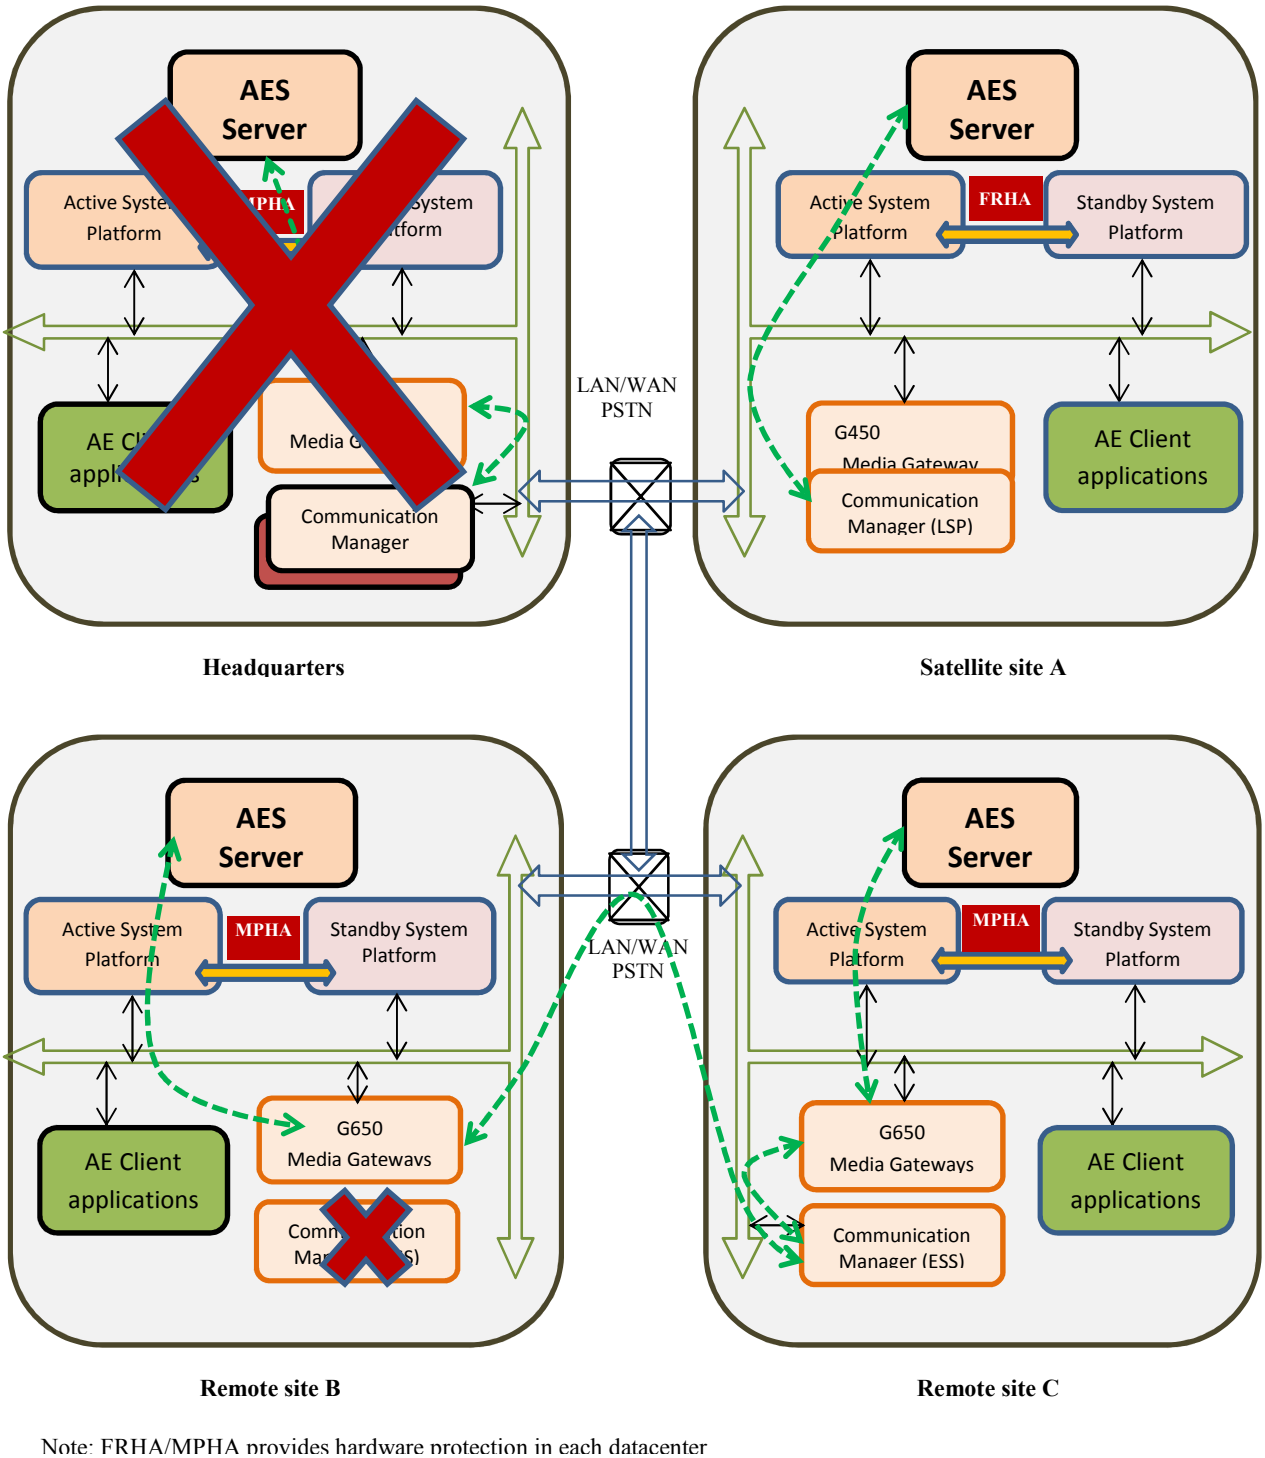

Figure 4: Secondary Failure

6.5 AE Services Behavior when a new Communication Manager Node is selected

The list below describes the general behavior of AE Services when the AE Services session moves from one Communication Manager node to another:

a. **DMCC (Device and Media Control) Service**
For DMCC endpoints registered using the Time-To-Service feature, the endpoints will be automatically re-registered to the new Communication Manager node, with no notification to the application necessary. However, the application will receive an unregistered event for each DMCC endpoint that is not using the Time-To-Service feature when connectivity is lost. There is a slight possibility that an endpoint using the Time-To-Service feature will fail to re-register, in which case an unregistered event for that endpoint will also be sent to the client application.

At this point, the application should begin attempts to re-register the DMCC endpoints (that failed the automatic re-registration). After the application has successfully re-registered all DMCC endpoints, it should reestablish its previous state and resume operation.

b. **CallInformation Services within DMCC, Call Control Services within DMCC, and all other CTI services**
The CallInformation and Call Control services within DMCC and all other CTI Services (TSAPI, CVLAN, DLG and JTAPI) use the Transport (AEP) link to communicate with Communication Manager.

If AE Services moves off of a Communication Manager node, it will notify any connected applications via a LinkDownEvent (DMCC CallInformationServices) or a CTI link down notification (CTI services). For Call Control Services within DMCC, Avaya recommends that applications add a CallInformationListener and look for a LinkDownEvent for indication that connectivity to the main site is down. (In future releases, Call Control Services clients will receive a MonitorStop request for all call control monitors if the link is lost to the main site.) Depending on the CTI API, clients will receive an appropriate event when the connectivity to the current Communication Manager node is lost. CVLAN clients will receive an “abort” for each association. TSAPI clients will receive a CSTAMonitorEnded event if the client is monitoring a device and/or a CSTASysStatEvent with a link down indication if the client is monitoring system status. TSAPI clients will also receive a CSTARouteEnd event for any active routing dialogs, and a CSTARouteRegisterAbort event for any registered routing devices. Avaya JTAPI 5.2 and later clients will receive a PROVIDER_OUT_OF_SERVICE event if the client has ProviderListeners. Otherwise, a ProvOutOfServiceEv event will

be received if the client has ProviderObservers. DLG clients will receive a link status event with a link down indication and a cause value.

As soon as AE Services connects to the new Communication Manager node, the application will be notified that the CTI link is back up, and the application can begin to resume normal operations. Since there is no run-time state preserved on a transition between Communication Manager nodes (as there is with an interchange on a duplicated Communication Manager media server) all application state must be reestablished. Note that, from the AE Services server's and application's perspectives, the failure scenario and recovery actions appear ***exactly*** the same as a long network outage between the AE Services server and the primary Communication Manager media server.

6.6 Differences between AE Services 6.1 and previous AE Services releases in Communication Manager Survivable Configurations

- Processor Ethernet is now fully supported in multi node survivable server environments (requires Communication Manager 6.0 or later).
- AEP connections to inactive survivable core server (ESS) and survivable remote server (LSP) nodes remain up (requires Communication Manager 6.0 or later). Previously, these connections would be established and immediately dropped (approximately every 40 seconds).
- Communication Manager 6.0 (or later) survivable server nodes inform AE Services 6.1 (or later) whenever they transition to active (i.e., a least one port network or media gateway is registered) or idle (i.e., there are no port networks or media gateways registered).
- AE Services 6.1 (or later) will not allow a session to an idle Communication Manager 6.0 (or later) survivable server node. In this case, all Computer Telephony Integration (CTI) links will remain down until that node becomes active.
- AE Services 6.1 (or later) can be configured to stay connected to an active survivable server node even after the primary (main) Communication Manager node comes back online. Previously, AE Services would instantly and automatically switch back to the primary Communication Manager server as soon as it connected to it, often further disrupting service.
- DMCC endpoints registered to the primary switch (using the Time-To-Service feature) will automatically re-register to the survivable server node. Similarly, automatic re-registration back to the primary Communication Manager node will occur at the appropriate time.

Note: This feature is not intended to be used in a geographically redundant configuration, in which the AE Services server is a "hot standby" at a remote site, connected to a "standby" Communication Manager survivable server node. In fact, it prevents establishment of CTI links to idle Communication Manager 6.0 (or later) survivable server nodes. It does allow an AE

Services server to be a geographically redundant standby server. In this case, the session will become active as soon as the survivable server node becomes active.

6.7 Communication Manager Fragmentation

Communication Manager's fragmentation occurs when AE Services is able to communicate to multiple active Communication Manager nodes in the same system at the same time. In this case, AE Services evaluates each node based on administered survivability parameters to determine which one is the best candidate to use.

6.7.1 Communication Manager 6.0 (or later) State Information

When an AE Services 6.1 (or later) server connects to a Communication Manager 6.0 (or later) media server, the Communication Manager's media server will inform the AE Services server of its server role (Main, ESS, or LSP) and its server state (active or idle). By definition, the main (or primary) Communication Manager server is always active when it is running. A Communication Manager survivable server is active when there is at least one port network **or** media gateway registered. A Communication Manager survivable server is idle when there is **no** port network or media gateway registered. This information allows AE Services 6.1 (or later) to determine which Communication Manager nodes are viable for use. If multiple nodes in the same system are viable at the same time, the cluster IDs/MIDs are used to determine from which node service will be provided, based on the priority as administered on the Survivability Hierarchy OAM screen. Every time a Communication Manager node transitions its state (either to active or idle), AE Services will evaluate which node should be used based on the survivability hierarchy settings.

Survivable servers are always ready to provide service, and wait for port networks or media gateways to register to them. Within 5 seconds of that registration, that survivable server will inform AE Services 6.1 (or later) that it has transitioned from idle to active, which allows AE Services 6.1 (or later) to use that node if necessary. Likewise, as soon as the last port network or media gateway unregisters, it will inform AE Services 6.1 (or later) that it has transitioned back to idle, which will stop AE Services 6.1 (or later) from using it.[2]

If PE connectivity is used to connect to multiple pre Communication Manager 6.0 nodes, then AE Services 6.1 (or later) will not be able to tell which nodes are actually active, and therefore it may choose (based on the priority administered on the survivability hierarchy OAM screen) to use an idle Communication Manager node, which would not be able to provide any reasonable service to the end user applications. This is not an issue with CLAN connectivity. Since CLANs

[2] Communication Manager maintenance can take 30-60 seconds to decide that the last media gateway has unregistered after it actually has unregistered, so AE Services will be notified within 5 seconds after that delay. There is no such maintenance delay for port networks.

reside in port networks, any Communication Manager node to which AE Services 6.1 (or later) connects via CLAN is active by definition (since the port network must be registered to that Communication Manager node to provide the connectivity).

6.7.2 Communication Manager media server node priority

Every Communication Manager node[3] in a system has a unique cluster ID/MID. The primary, or main, Communication Manager node always has a cluster ID/MID of 1. Survivable server nodes have unique cluster IDs/MIDs greater than 1. AE services 6.1 (or later) provides the ability to administer the cluster IDs/MIDs of survivable server nodes in priority order on a per switch connection basis with the Survivability Hierarchy OAM screen. The main Communication Manager node always has the highest precedence. Cluster IDs/MIDs that are received by AE Services 6.1 (or later), but are not listed on the Survivability Hierarchy OAM screen, will have a lower precedence than all IDs that are listed there. Within the set of cluster IDs/MIDs that are not listed on that screen, the lower number IDs will have higher precedence. If, your example, cluster IDs/MIDs 7, 3, and 5 where administered on the Survivability Hierarchy OAM screen (in that priority order), and AE Services 6.1 (or later) connected to four different survivable server nodes that returned cluster IDs/MIDs of 2, 4, 5 and 7, these nodes would be evaluated in the following priority order: 7, 5, 2, and 4.

[3] Note: A Communication Manager node can be either a simplex media server or a duplicated media server pair. Both servers on a duplicated system have the same cluster ID/MID.

7 Terminology and Acronyms

Term	Meaning
AEP	Application Enablement Protocol
AE Services	Application Enablement Services
API	Application Programming Interface
ASAI	Adjunct Switch Application Interface
CLAN	Control Local Area Network interface card
CM	Communication Manager
CTI	Computer Telephony Integration
CVLAN	CallVisor LAN
DLG	Definity LAN Gateway
DMCC	Device, Media and Call Control
ESS	Enterprise Survivable Server, now called survivable core server
HA	High Availability
JTAPI	Java Telephony API
LAN	Local Area Network
LSP	Local Survivable Server, now called survivable remote server
MID	Module ID
OAM	Operations Administration and Maintenance
PE	Processor Ethernet, also referred to as procr
SP	System Platform
TSAPI	Telephony Server API
TTS	Time-To-Service

WAN	Wide Area Network