White Paper

Communication Manager Survivability in an Environment with Media Servers

Abstract

This white paper discusses a variety of topics pertaining to adding Media Servers (MS) to a Communication Manager (CM) configuration that includes Survivable Processors (SP). The intended audience is CM planners and administrators. The reader is assumed to already be familiar with how survivability is used with Port Networks (PN) and Media Gateways (MG).

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

"Documentation" means information published by Avaya in varying mediums which may include product information, operating instructions and performance specifications that Avaya may generally make available to users of its products and Hosted Services. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of documentation unless such modifications, additions, or deletions were performed by Avaya. End User agrees to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked websites referenced within this site or documentation provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on Avaya hardware and software. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya’s standard warranty language, as well as information regarding support for this product while under warranty is available to Avaya customers and other parties through the Avaya Support website: http://support.avaya.com or such successor site as designated by Avaya. Please note that if You acquired the product(s) from an authorized Avaya Channel Partner outside of the United States and Canada, the warranty is provided to You by said Avaya Channel Partner and not by Avaya.

"Hosted Service" means a hosted service subscription that You acquire from either Avaya or an authorized Avaya Channel Partner (as applicable) and which is described further in Hosted SAS or other service description documentation regarding the applicable hosted service. If You purchase a Hosted Service subscription, the foregoing limited warranty may not apply but You may be entitled to support services in connection with the Hosted Service as described further in your service description documents for the applicable Hosted Service. Contact Avaya or Avaya Channel Partner (as applicable) for more information.

Hosted Service

THE FOLLOWING APPLIES IF YOU PURCHASE A HOSTED SERVICE SUBSCRIPTION FROM AVAYA OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE), THE TERMS OF USE FOR HOSTED SERVICES ARE AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO UNDER THE LINK "Avaya Terms of Use for Hosted Services" OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, AND ARE APPLICABLE TO ANYONE WHO ACCESSES OR USES THE HOSTED SERVICE. BY ACCESSING OR USING THE HOSTED SERVICE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE DOING SO (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THE TERMS OF USE. IF YOU ARE ACCEPTING THE TERMS OF USE ON BEHALF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE AUTHORITY TO BIND SUCH ENTITY TO THESE TERMS OF USE. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO ACCEPT THESE TERMS OF USE, YOU MUST NOT ACCESS OR USE THE HOSTED SERVICE OR AUTHORIZE ANYONE TO ACCESS OR USE THE HOSTED SERVICE. YOUR USE OF THE HOSTED SERVICE SHALL BE LIMITED BY THE NUMBER AND TYPE OF LICENSES PURCHASED UNDER YOUR CONTRACT FOR THE HOSTED SERVICE, PROVIDED, HOWEVER, THAT FOR CERTAIN HOSTED SERVICES IF APPLICABLE, YOU MAY HAVE THE OPPORTUNITY TO USE FLEX LICENSES, WHICH WILL BE INVOICED ACCORDING TO ACTUAL USAGE ABOVE THE CONTRACT LICENSE LEVEL. CONTACT AVAYA OR AVAYA'S CHANNEL PARTNER FOR MORE INFORMATION ABOUT THE LICENSES FOR THE APPLICABLE HOSTED SERVICE, THE AVAILABILITY OF ANY FLEX LICENSES (IF APPLICABLE), PRICING AND BILLING INFORMATION, AND OTHER IMPORTANT INFORMATION REGARDING THE HOSTED SERVICE.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AVAYA CHANNEL PARTNER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA CHANNEL PARTNER; AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE ("AVAYA").

Avaya grants You a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. "Software" means computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products, pre-installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. "Designated Processor" means a single stand-alone computing device. "Server" means a Designated Processor that hosts a software application to be accessed by multiple users. "Instance" means a single

copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine ("VM") or similar deployment.

License type(s)

Designated System(s) License (DS). End User may install and use each copy or an Instance of the Software only on a number of Designated Processors up to the number indicated in the order. Avaya may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Avaya through electronic means established by Avaya specifically for this purpose.
Concurrent User License (CU). End User may install and use the Software on multiple Designated Processors or one or more Servers, so long as only the licensed number of Units are accessing and using the Software at any given time. A "Unit" means the unit on which Avaya, at its sole discretion, bases the pricing of its licenses and can be, without limitation, an agent, port or user, an e-mail or voice mail account in the name of a person or corporate function (e.g., webmaster or helpdesk), or a directory entry in the administrative database utilized by the Software that permits one user to interface with the Software. Units may be linked to a specific, identified Server or an Instance of the Software.
Database License (DL). End User may install and use each copy or an Instance of the Software on one Server or on multiple Servers provided that each of the Servers on which the Software is installed communicates with no more than one Instance of the same database.
CPU License (CP). End User may install and use each copy or Instance of the Software on a number of Servers up to the number indicated in the order provided that the performance capacity of the Server(s) does not exceed the performance capacity specified for the Software. End User may not re-install or operate the Software on Server(s) with a larger performance capacity without Avaya's prior consent and payment of an upgrade fee.
Named User License (NU). You may: (i) install and use the Software on a single Designated Processor or Server per authorized Named User (defined below); or (ii) install and use the Software on a Server so long as only authorized Named Users access and use the Software. "Named User", means a user or device that has been expressly authorized by Avaya to access and use the Software. At Avaya's sole discretion, a "Named User" may be, without limitation, designated by name, corporate function (e.g., webmaster or helpdesk), an e-mail or voice mail account in the name of a person or corporate function, or a directory entry in the administrative database utilized by the Software that permits one user to interface with the Software.
Shrinkwrap License (SR). You may install and use the Software in accordance with the terms and conditions of the applicable license agreements, such as "shrinkwrap" or "clickthrough" license accompanying or applicable to the Software ("Shrinkwrap License").

Copyright

Virtualization

Each product has its own ordering code and license types. Note that each Instance of a product must be separately licensed and ordered. For example, if the end user customer or Avaya Channel Partner would like to install two Instances of the same type of products, then two products of that type must be ordered.

Third Party Components

"Third Party Components" mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements ("Third Party Components"), which contain terms regarding the rights to use certain portions of the Software ("Third Party Terms"). As required, information regarding distributed Linux OS source code (for those products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the products, Documentation or on Avaya's website at: http://support.avaya.com/Copyright or such successor site as designated by Avaya. You agree to the Third Party Terms for any such Third Party Components.

THIS PRODUCT IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO (i) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD ("AVC VIDEO") AND/OR (ii) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE HTTP://WWW.MPEGLA.COM.

Note to Service Provider

The product or Hosted Service may use Third Party Components subject to Third Party Terms that do not allow hosting and require a Service Provider to be independently licensed for such purpose. It is your responsibility to obtain such licensing.

Preventing Toll Fraud

"Toll Fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

If You suspect that You are being victimized by Toll Fraud and You need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support website: http://support.avaya.com or such successor site as designated by Avaya. Suspected security vulnerabilities with Avaya products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation, Hosted Service(s), and product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation, Hosted Service(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party.
Avaya is a registered trademark of Avaya Inc.
All non-Avaya trademarks are the property of their respective owners. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: http://support.avaya.com, or such successor site as designated by Avaya.

Contact Avaya Support

See the Avaya Support website: http://support.avaya.com for product or Hosted Service notices and articles, or to report a problem with your Avaya product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: http://support.avaya.com (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Contents

Introduction

This white paper discusses a variety of topics pertaining to adding Media Servers (MS) to a Communication Manager (CM) configuration that includes Survivable Processors (SP). The intended audience is CM planners and administrators. The reader is assumed to already be familiar with how survivability is used with Port Networks (PN) and Media Gateways (MG). Each topic is independent—you can read them in any order.

Terminology

Term	Meaning
Active	A server, either a Main or a Survivable Processor (SP), is considered "active" if it has access to any media resources. These media resources might be in Port Networks (PN), Media Gateways (MG) and/or Media Servers (MS). An SP only gains access to media resources when the Main loses access to them. The Main regains access to media resources only when allowed to do so by the Recovery Rules.
Backup Servers	One or more Survivable Processors (SP), in preference order, assigned to back up a particular Network Region (NR) in the event of failure. The Backup Server lists for the various NR feed into the creation of Alternate Gatekeeper Lists (AGL) that are used by H.323 phones, and also determines which NR that the Split Registration Prevention (SRP) feature disables. The Backup Server list is administered on the "change ip-network-region" form.
Community	A set of Port Networks (PN) that share a common set of rules to follow in the event of a failover. The Community of a PN, as well as the settings for each Survivable Processor (SP), determines which server the PN will connect to in the event of a failover. The PN Community is administered on the "change system-parameters port-networks" form, and the Community of an SP is administered on the "change survivable-processor" form.
Enterprise Survivable Server (ESS)	A now obsolete term, replaced by "Survivable Core".
Failover	The act of transferring access to one or more media resources from a higher-priority server (usually but not necessarily the Main) to a Survivable Processor (SP) in the event of a failure that renders the higher-priority server unable to access those resources.
Fallback	The act of returning access to one or more media resources from a Survivable Processor (SP) to a higher-priority server (usually but not necessarily the Main) after the repair of the failure that caused the failover. Fallback is controlled by the Recover Rules. Fallback is also sometimes called "Recovery".
Fully Qualified Domain Name (FQDN)	Within the context of this paper, a FQDN is a symbolic name (e.g., pearl.avaya.com) associated with a SIP Entity. It defines, via Local Host Name Resolution (LHNR), an ordered set of IP addresses for that entity. The SIP Entity for purposes of this paper represents a Communication Manager (CM) and all of its Survivable Processors (SP).
Inactive	The opposite of "Active". See "Active".
Local Host Name Resolution (LHNR)	A mapping within Session Manager (SM) of a Fully Qualified Domain Name (FQDN) to an ordered set of IP addresses. LHNR overrides Domain Name System (DNS).
Local Survivable Processor (LSP)	A now obsolete term, replaced by "Survivable Remote".
Main	The Communication Manager (CM) system that manages all resources on a sunny day. The Main is backed up by one or more Survivable Processors (SP).

Term	Meaning
Media Gateway (MG)	Also sometimes called an H.248 Media Gateway. E.g., G350, G430, G450, G700. 
Media Gateway Controller (MGC)	The brains of a Media Gateway (MG).
Media Server (MS)	A software-based media processing server that provides media resources to Communication Manager (CM). CM can also get media resources from Port Networks (PN) and Media Gateways (MG). Unlike PN and MG, a Media Server has no line or trunk interfaces.
Media Server Element Manager	Within this context, a web interface that allows you to stop and start a Media Server (MS). It does other stuff too, but that's not relevant to this paper.
Media Server Reporting List (MSRL)	An administered list of Media Servers (MS) to which a particular Survivable Processor (SP) reports and receives status. The MSRL, in conjunction with the "Priority with respect to Media Servers", determines which one or more SP take over in the event of a failure. The MSRL is administered on the "change survivable-processor" form.
MGC List	A list of the IP addresses of servers to which a Media Gateway (MG) will go for service. The MGC List, administered on each MG individually. The MGC List has a "transition point", above which are the addresses for the Main and below which are the ordered addresses of the Survivable Processors (SP). The MGC List is administered on the Media Gateway CLI using the "set mgc list" command.
Network Region (NR)	A subset of the network, inside which there is considered to be unlimited bandwidth. Typically, a Network Region is a LAN, which is connected to other NR by a WAN. NR are administered on the "change ip-network-region" form, and are assigned on numerous other forms.
Port Network (PN)	G650, a 14-slot chassis that accommodates analog, digital, ISDN and IP interfaces. 
Priority	Priority is a value assigned to one particular IP address in Local Host Name Resolution (LHNR). Priority is used to control to which server the Session Manager (SM) will send new requests.
Priority Score	A number between 1 and 100 assigned to each Survivable Processor (SP) and used by Port Networks (PN), along with other factors, to order by preference which SP the PN will connect to in the event of a failover. Higher numerical values are higher preference. The Priority Score is administered on the "change survivable-processor" form.

Term	Meaning
Priority with respect to Media Servers	A number between 2 and 9999 which, combined with the Media Server Reporting List (MSRL), determines which Survivable Processor (SP) will begin using a Media Server (MS), and thus become active, in the event of a failure. Lower numerical values are higher priority. When multiple SP have the same MS on their MSRL, the highest priority functioning SP of the set will be the one using those resources. The Priority with respect to Media Servers is administered on the "change survivable-processor" form. If multiple SP have the same Priority with respect to Media Servers, their respective MSRL should be disjoint.
Processor Ethernet (PROCR)	The Ethernet port directly on the server. For a server with Port Networks (PN), there are two ways for a Media Gateway (MG) or H.323 endpoint to connect to it: either via a CLAN interface in a PN or directly via the PROCR. The Processor Ethernet is administered on the "change ip-interface procr" form.
Recovery	See "Fallback".
Recovery Rules	The rules that govern under what conditions recovery can take place. E.g., automatic recovery (perhaps after some time delay), scheduled recovery for a particular day and time, manual recovery. Recovery rules are administered separately for Port Networks (PN), Media Gateways (MG) and Media Servers (MS), on these forms: • change system-parameters port-networks • change system-parameters mg-recovery-rule • change system-parameters ms-recovery-rule
Session Manager (SM)	The SIP routing core of the Avaya Aura architecture. For purposes of this paper, SM is the agent that sends calls to Communication Manager (CM) servers. SM only sends calls to active servers, i.e., servers with media resources. OPTIONS messages are sent periodically from SM to each server, and the servers reply positively or negatively depending on whether they are active or inactive.
Signaling Group	The pipe by which Communication Manager (CM) talks to a Media Server (MS). Signaling groups are used for many other things too, but those uses are irrelevant to this paper. For purposes of this paper, we only care about the fact that if you busyout the signaling group ("busyout signaling-group") to the MS, we stop talking to it, and if you release the signaling group ("release signaling-group"), we start talking to the MS again.
SIP Entity	A SIP Entity is a Session Manager (SM) concept that groups the Main and all its Survivable Processors (SP) into one pool of servers to which it can send SIP requests. When a failover occurs, SM starts sending new SIP requests to whichever SP becomes active. If more than one SP become active, SM sends requests to the highest-priority SP.
Split Registration	Split registration is a potentially undesirable condition in which H.323 phones and the resources they need are connected to different servers.
Split Registration Prevention (SRP)	A Communication Manager (CM) feature that prevents fragmentation of the endpoints and media resources in certain situations. SRP only affects H.323 phones, Media Gateways (MG) and Media Servers (MS). It does not affect Port Networks (PN) or SIP phones. SRP is enabled with the "Force Phones and Gateways to Active Survivable Servers" field on the "change system-parameters ip-options" form.

Term	Meaning
Survivable Core	A type of Survivable Processor (SP) formerly known as ESS. Unlike a Survivable Remote, a Survivable Core has the ability to take over the entire enterprise or any portion of it.
Survivable Processor (SP)	A generic term that includes both Survivable Core and Survivable Remote. Use the "list survivable-processor" command to see the SP. On the Main, this command will show them all. On an SP, this command will only show you yourself.
Survivable Remote	A type of Survivable Processor (SP) formerly known as LSP. A Survivable Remote has more limited capabilities than a Survivable Core. Notably, a Survivable Remote cannot access Port Networks (PN).
System Access Terminal (SAT)	An interface into Communication Manager (CM) for administrative and maintenance functions.
System Manager (SMGR)	System Manager is the administrative interface for Avaya Aura. For purposes of this paper, it is used to administer the Communication Manager (CM) SIP Entity and the Local Host Name Resolution (LHNR) entries that define the IP addresses for the Main and all its Survivable Processors (SP).
Transition Point	A demarcation point in an MGC List. IP addresses above the transition point are considered to belong to the Main (either the PROCR or a CLAN), whereas IP addresses below the transition point are considered to belong to one or more Survivable Processors (SP). The transition point is expressed as an integer. E.g., if an MGC List has 6 entries and the transition point is 2, then the first two entries are above the transition point and the last four are below. The transition point is administered on the Media Gateway CLI using the "set reset-times transition-point" command.
Translations	The Communication Manager (CM) database that contains configuration parameters. In the case of a Main and the Survivable Processors (SP) that back it up, they all share the same translation file. Translations are only modified on the Main, and then are propagated to the SP when the "save translation all" command is entered on the Main SAT.
Weight	Weight is a value assigned to one particular IP address in Local Host Name Resolution (LHNR). When two LHNR entries have the same priority, weight is used to control to which requests are sent. For purposes of this paper, however, weight is irrelevant.

Reference: Administration Screens

This section provides sample administration screens.

Use “list survivable-processor” to check the status of your Survivable Processors (SP), e.g., whether they are registered, whether they are active and the status of their translations.

```
list survivable-processor
                              SURVIVABLE PROCESSORS
 Record Name/              Type         Reg Act         Translations      Net
 Number  IP Address                                     Updated           Rgn

  1     west-coast-ess    ESS S          y   y         8:33 2/13/2015    2
         123.45.6.78
        No V6 Entry
  2     east-coast-ess    ESS S          y   n         8:33 2/13/2015    3
         123.45.6.79
        No V6 Entry
  3     chicago-lsp       LSP            y   n         8:33 2/13/2015    103
         123.45.6.80
        No V6 Entry
  4     core-ess          ESS D          y   n         8:33 2/13/2015    1
         123.45.6.81
        No V6 Entry
```

Use “change survivable-processor” to configure your Survivable Processors (SP), set the Priority Score and Community for Port Network failover, set the Priority with respect to Media Servers, and establish their Media Server Reporting Lists.

```
change survivable-processor west-coast-ess                      Page   1 of   8
                              SURVIVABLE PROCESSOR
Type: simplex-ess  Cluster ID/MID: 2     Processor Ethernet Network Region: 2
                         Community: 2         Enable PE for H.323 Endpoints? y
                                               Enable PE for H.248 Gateways? y
SERVER A
    Server ID: 3
 V4 Node Name: west-coast-ess  Address: 123.45.6.78
 V6 Node Name:                 Address:
PORT NETWORK PARAMETERS
                   Community Size: all      System Preferred: n
                   Priority Score: 1         Local Preferred: n
                                                Local Only: n
```

```
change survivable-processor west-coast-ess                      Page   3 of   8
                              SURVIVABLE PROCESSOR
MEDIA SERVER PARAMETERS
  Priority with respect to Media Servers: 5
```

```
change survivable-processor west-coast-ess                          Page   4 of   8
                MEDIA SERVER REPORTING LIST FOR west-coast-ess
   Num NR  Node Name          Num NR  Node Name          Num NR  Node Name
   1   1   media-server-1
   2   2   media-server-2
   3   103 media-server-3
   4   4   media-server-4
   5   3   media-server-5
```

Use “change system-parameters ms-recovery rule” to set the Media Server Recovery Rule: automatic, scheduled or manual.

```
change system-parameters ms-recovery-rule                            Page   1 of   1
                          MEDIA SERVER RECOVERY RULES
FAILOVER PARAMETERS                      FALLBACK PARAMETERS
    Report Interval (sec): 60              Auto Return: no
  Report Expiration (sec): 180
```

Use “change system-parameters mg-recovery-rule” to set the Media Gateway Recovery Rules: automatic, scheduled or manual. This controls the timing of the fallback for one or more Media Gateways.

```
change system-parameters mg-recovery-rule 1                          Page   1 of   1
             SYSTEM PARAMETERS MEDIA GATEWAY AUTOMATIC RECOVERY RULE
 Recovery Rule Number: 1
 Rule Name: Scheduled
 Migrate H.248 MG to primary: time-day-window
 Minimum time of network stability: 3
                                  Time of Day
 Day of Week  00                       12                    23
 Sunday             X
 Monday             X
 Tuesday            X
 Wednesday          X
 Thursday           X
 Friday             X
 Saturday           X
```

Use “change system-parameters port-networks” to set the parameters for Port Network failover, including Community and Recovery Rules. These determine the failover and fallback strategies for Port Networks.

```
change system-parameters port-networks                               Page   1 of   2
                COMMUNITY ASSIGNMENTS FOR PORT NETWORKS
PN Community     PN Community     PN Community     PN Community     PN Community
 ------------     ------------     ------------     ------------     ------------
  1: 1             14: 1            27: 1            40: 1            53: 1
  2: 1             15: 1            28: 1            41: 1            54: 1
```

```
change system-parameters port-networks                               Page   2 of   2
                    PORT NETWORK RECOVERY RULES
FAILOVER PARAMETERS                      FALLBACK PARAMETERS
No Service Time Out Interval (min): 5          Auto Return: no
   PN Cold Reset Delay Timer (sec): 60
```

Use “change system-parameters ip-options” to control failover speed for H.323 phones, and also to turn on or off Split Registration Prevention (SRP).

```
change system-parameters ip-options                              Page   1 of   4
                          IP-OPTIONS SYSTEM PARAMETERS
…
                                      H.323 IP ENDPOINT
 H.248 MEDIA GATEWAY                   Link Loss Delay Timer (min): 5
  Link Loss Delay Timer (min): 5         Primary Search Time (sec): 75
                                  Periodic Registration Timer (min): 20
                                Short/Prefixed Registration Allowed? Y
```

```
change system-parameters ip-options                              Page   2 of   4

 Force Phones and Gateways to Active Survivable Servers? n
```

Use the Media Gateway CLI to set the server addresses and the transition point on a Media Gateway. In this example, on a G450 gateway, 123.45.0.67 is the Main and 123.45.0.68 is the Survivable Processor. This is how you set the failover strategy of a Media Gateway.

```
mygateway(super)# set mgc list 123.45.0.67, 123.45.0.68
Done!
```

```
mygateway(super)# set reset-times transition-point 1
Done!
```

```
mygateway(super)# set reset-times primary-search 20
Done!
```

```
mygateway(super)# show recovery

RECOVERY TIMES
-------------------------------
Primary Search  : 20
Total Search    : 30
Transition Point: 1
```

```
mygateway(super)# show mgc list

PRIMARY MGC HOST, Primary Search Time : 1 min(s)
IPv4 Address         IPv6 Address
-------------------- ----------------------------------------------
123.45.0.67          -- Not Available --

SECONDARY MGC HOST
IPv4 Address         IPv6 Address
-------------------- ----------------------------------------------
123.45.0.68          -- Not Available --
-- Not Available --  -- Not Available --
-- Not Available --  -- Not Available –
```

Use the Media Server Element Manager to stop and start the Media Server. This can be used to avoid accidental failover if you need to remove a Media Server from a MSRL, or if you need to busyout the Media Server signaling group for more than a brief time.

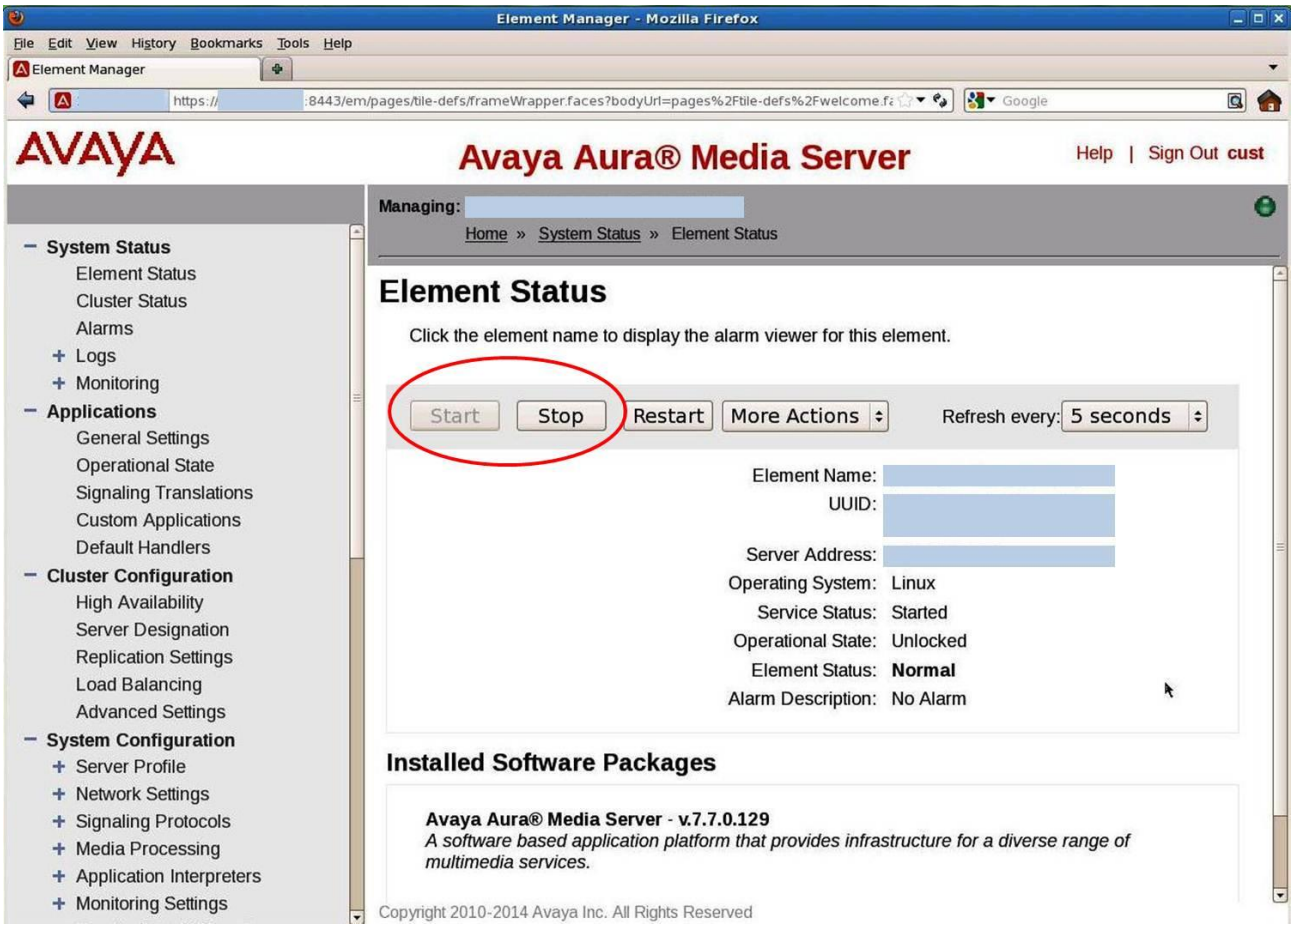

Use the CM SMI “Server Role” screen to set how a SP registers with the Main. If you want a Survivable Remote to be able to register with both the Main and a Survivable Core, you need to use a CLAN as a registration address. This is significant mainly if you are using Split Registration Prevention between the Survivable Core and the Survivable Remote.

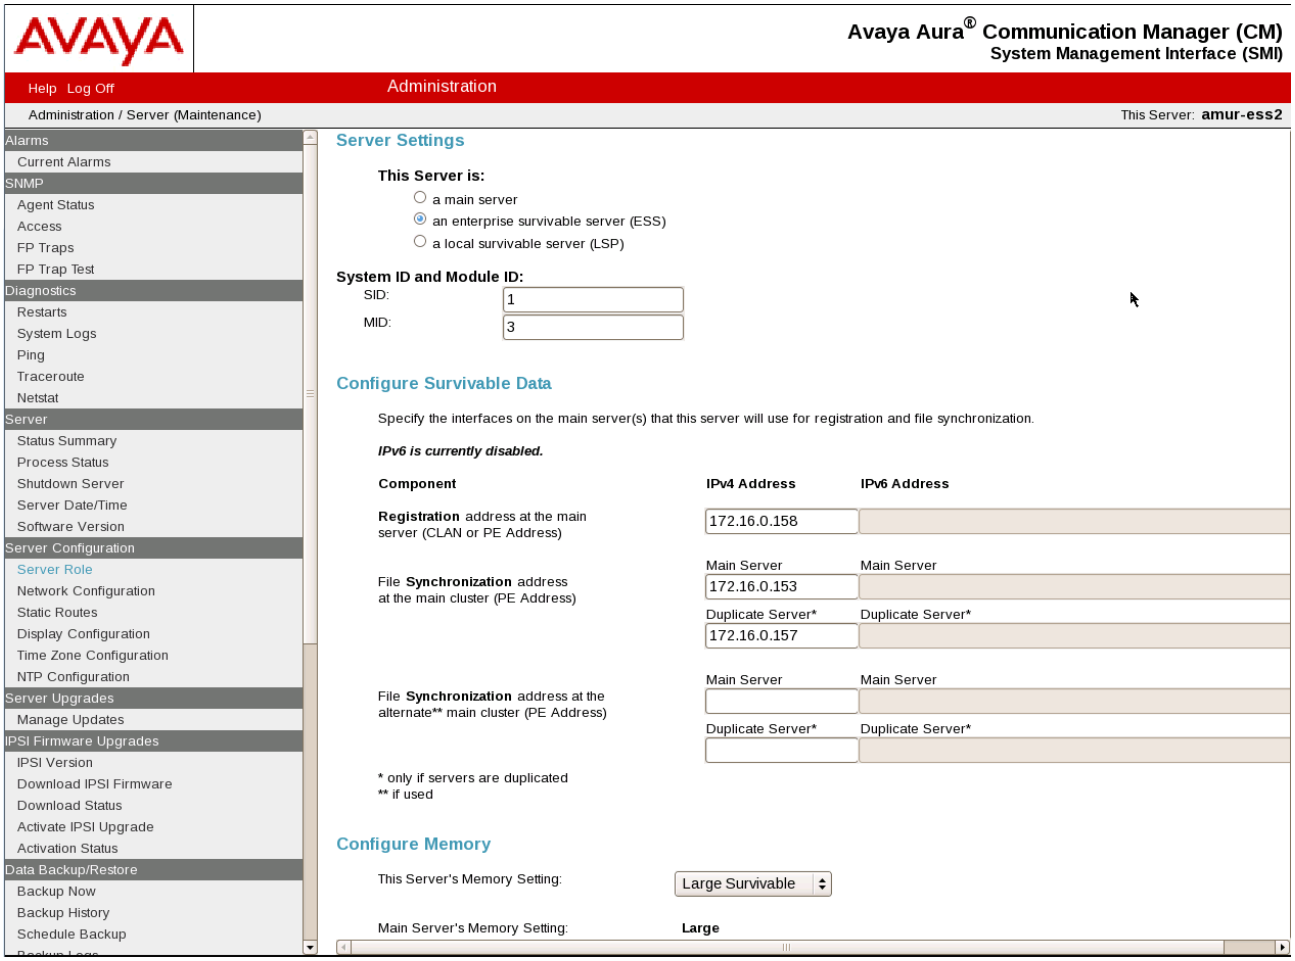

An FQDN to identify one set of servers (the Main and all its SP) is created in SMGR.

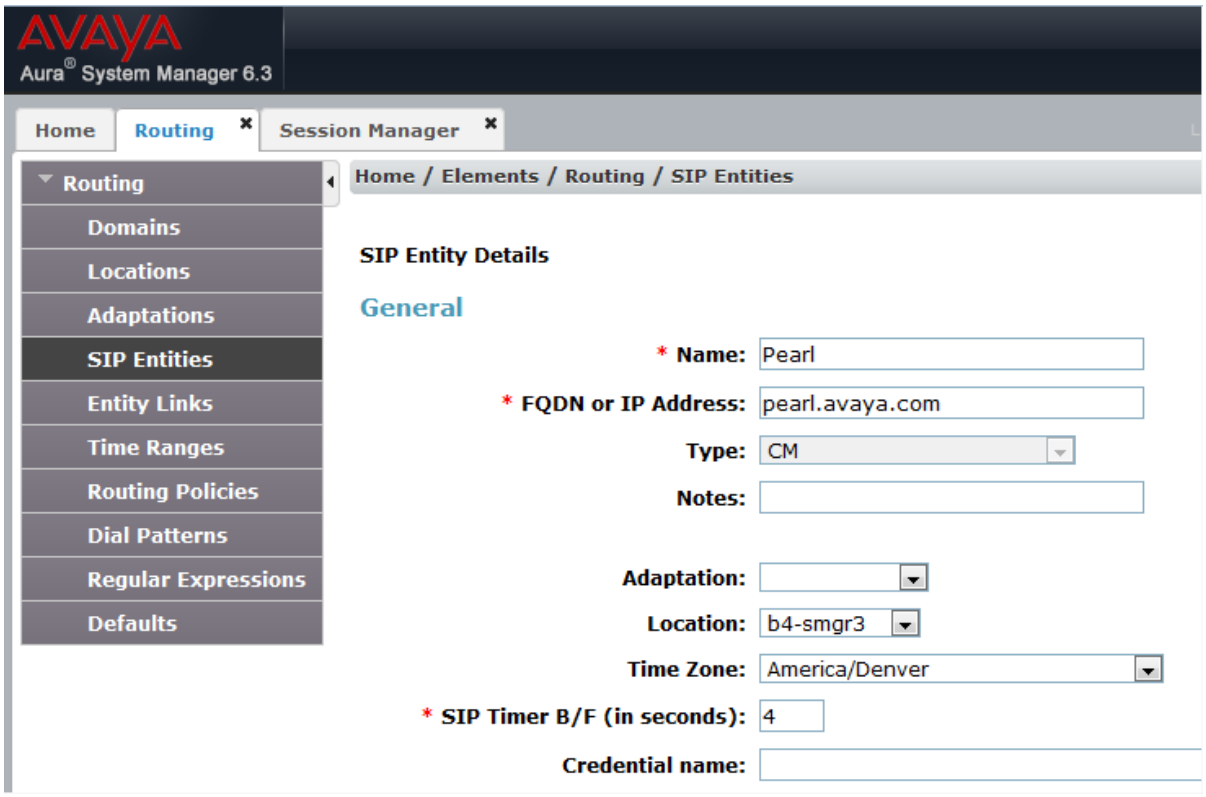

IP addresses and priorities are associated with the FQDN in Local Host Name Resolution.

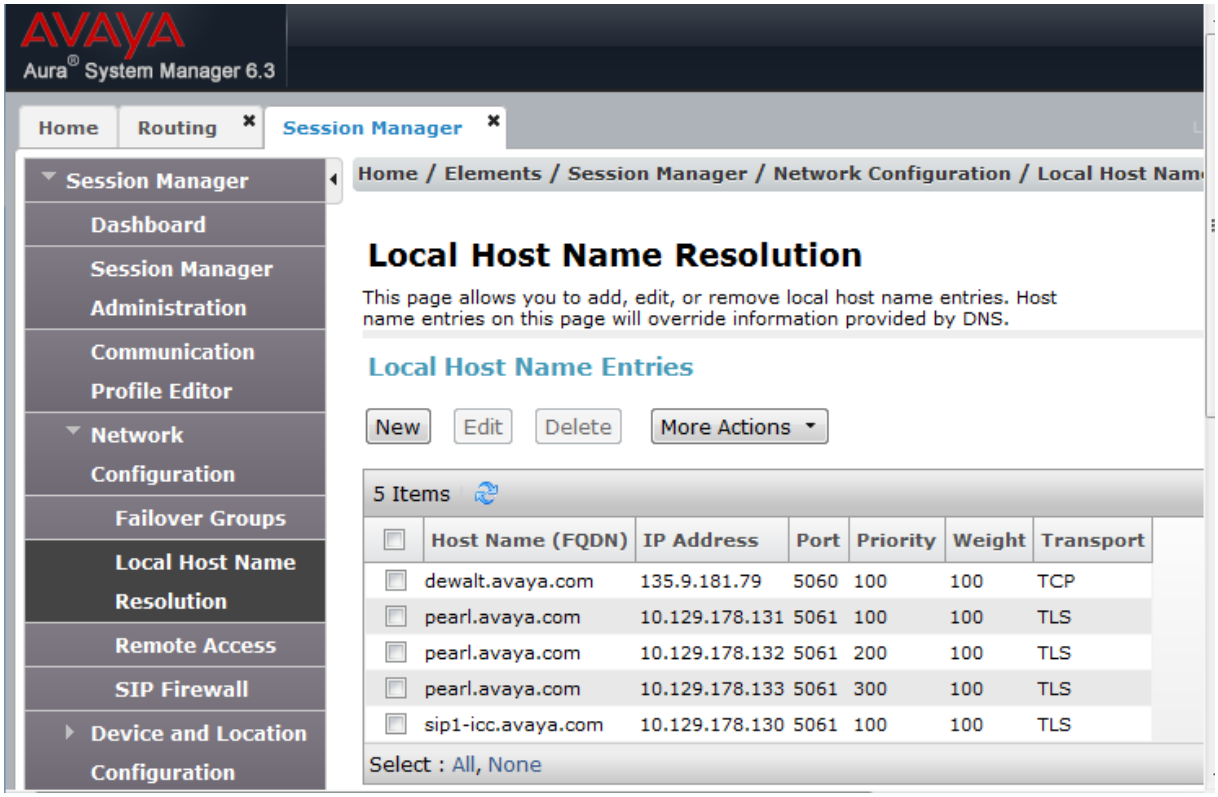

Guidelines: Consistent Failover Strategy for PN/MG/MS Environments

This section provides administration recommendations and guidelines.

When you have a configuration with a mixture of Port Networks (PN), Media Gateways (MG) and Media Servers (MS), you most likely want a failover strategy that causes all of these types of media resources to go to the same Survivable Processor(s) (SP) for the same type of failure. You probably also want your H.323 endpoint failover strategy to be consistent with your media resource failover strategy. This section contains suggestions for how you might achieve this behavior.

Where PN go on a failover is guided by:

- The “Community” of the PN (on “change system-parameters port-networks” form).
- The “Community” of the SP (on “change survivable-processor” form).
- The “PORT NETWORK PARAMETERS” of the SP (on “change survivable-processor” form).

Where MG go on a failover is guided by:

- The “mgc list” configured on the MG CLI.
- The “transition point” configured on the MG CLI.

Where MS go on a failover is guided by:

- The “Priority with respect to Media Servers” (on “change survivable-processor” form).
- The MEDIA SERVER REPORTING LIST (on “change survivable-processor” form).

Where H.323 endpoints go on a failover is guided by:

- The Network Regions (NR) of the endpoints (on “change ip-network-map” form).
- The “BACKUP SERVERS” for each NR (on “change ip-network-region” form).

So now let’s look at a few examples.

One Survivable Processor

Let's take the trivial case first. You only have one SP (SP-1).

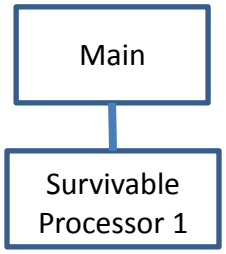

If the Main fails, you want all media resources and endpoints to go to SP-1. Also, if the network fragments such that the Main cannot reach certain NR, you want SP-1 to take over for those NR while the Main continues to serve the NR that it can still reach.

- H.323 endpoints[1]

	Backup Servers
All NR	SP-1

- PN[2]

	Community	***Priority Score***	***Local Only***	***System Preferred***
All PN	1			
SP-1	1	1	n	y

- MGC List[3]

For all MG
Main
SP-1

- MS[4]

	Priority wrt MS	***MSRL***
SP-1	2	At least one MS from each NR

[1] Put SP-1 as the backup server on the "change ip-network-region" form for all NR that have H.323 endpoints. That will add SP-1 to the AGL of all H.323 endpoints so that the endpoints will go look for SP-1 if they lose connectivity to the Main.

[2] You only need one community because you only have one SP. Put SP-1 in the same community as all your PN. "System Preferred" and "Local Only" are irrelevant because there is only one community. Priority Score is also irrelevant because you only have one SP.

[3] Put SP-1 after the transition point in the MGC List so that all MG will go to SP-1 if they lose connectivity to the Main.

[4] Since this is the only SP, any priority other than blank will do. You do not necessarily need to place a MS within each NR, but if you have any MS in a particular NR, it is best to include at least one MS from that NR in the MSRL so that each NR will continue to get service if it gets isolated from the Main. Including more than one MS from an NR in the MSRL is helpful in case one MS is down at the time of the failure.

Two Hierarchical Survivable Processors

Suppose you have two SP (SP-1 and SP-2).

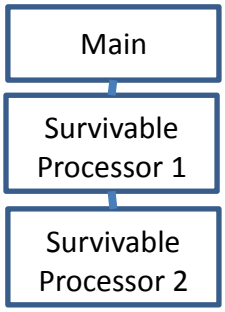

If the Main fails, you would like SP-1 to take over the whole network if it can. SP-2 becomes active only if at least one NR is servable by neither the Main nor SP-1. Note: This setup cannot be supported if the Split Registration Prevention Feature (SRP) is enabled because only one backup server is allowed per NR.

- H.323 endpoints[5]

	Backup Servers
All NR	SP-1 SP-2

- PN[6]

	Community	***Priority Score***	***Local Only***	***System Preferred***
All PN	1			
SP-1	1	100	n	y
SP-2	1	50		

- MGC List[7]

For all MG
Main
SP-1 SP-2

- MS[8]

	Priority wrt MS	***MSRL***
SP-1	2	At least one MS from each NR (both lists identical)
SP-2	3	

[5] Since you want all endpoints follow the same failover algorithm, the backup server lists for all NR can be identical. Put SP-1 first in the list to show that it is the preferred failover system. Put SP-2 second to show that it backs up SP-1.

[6] Since you want all PN to stay together if they can, you only need one community. SP-1 gets a higher "Priority Score" on the "change survivable-processor" form to show that it is preferred. The actual priority score values are irrelevant; only the relative values are significant. Since all PN can go to either SP if necessary, both SP can have the same community. "System Preferred" and "Local Only" are irrelevant because there is only one community.

[7] Put SP-1 first after the transition point in the MGC List so that all MG will try SP-1 first if they lose connectivity to the Main. Put SP-2 after SP-1 so that all MG will try SP-2 if they lose connectivity to both the Main and SP-1.

[8] SP-1 should get a higher priority (lower number) with respect to MS than SP-2 to indicate that it is preferred. The actual priority values are irrelevant; only the relative values are significant. Both MSRL should be identical so that each SP uses the same criteria to decide when to go active (using only the priority to differentiate their behavior).

Two Regional Survivable Processors

Suppose you have two survivable processors (SP-A and SP-B).

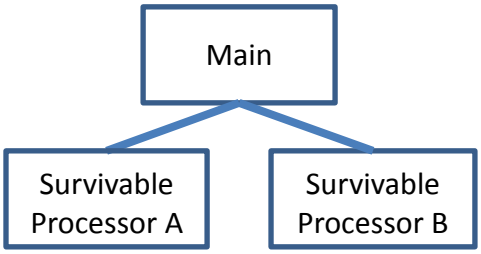

If the Main fails, you would like SP-A to be active for NR-A, and SP-B to be active for NR-B.

- H.323 endpoints[9]

	Backup Servers
All NR in NR-A	SP-A
All NR in NR-B	SP-B

- PN[10]

	Community	***Priority Score***	***Local Only***	***System Preferred***
All PN in NR-A	1			
All PN in NR-B	2			
SP-A	1	1	y	n
SP-B	2			

- MGC List[11]

For all MG in NR-A	***For all MG in NR-B***
Main	Main
SP-A	SP-B

- MS[12]

	Priority wrt MS	***MSRL***
SP-A	2	At least one MS from each NR in NR-A
SP-B	2	At least one MS from each NR in NR-B

[9] You need separate backup servers for the two sets of H.323 endpoints since you want different endpoints to go different SP for service in the event of a failure. All endpoints that you want to go to SP-A should be put in an NR that lists SP-A as the backup server. All endpoints that you want to go to SP-B should be put in an NR that lists SP-B as the backup server. This causes each set of endpoints to get the appropriate AGL for their region.

[10] Because we want PN to go to different SP in the event of a failover, we need two communities. The community numbers chosen are irrelevant as long as they are different from each other. Each PN should be in the same community as the SP that it should go to. Priority Score becomes irrelevant. Setting "Local Only" to "y" and "System Preferred" to "n" prevents a PN from going to the wrong SP.

[11] Different MG get different MGC Lists because we want different MG to go to different SP in the event of a failure. Below the transition point on each list, put only the SP to which you want that particular MG to go.

[12] Priority with respect to MS becomes irrelevant when you want each SP to make failover decisions based only on MS in its region. The MSRL should be the only deciding factor in these decisions. The MSRL for each SP should be disjoint, including only MS in NR that you want to be served by that SP. Note, however, that the MSRL does not restrict which MS a particular SP will use if it becomes active—it only restricts which MS a particular SP will use to decide to become active. Once active, an SP will use all the MS in the NR it backs up.

Hierarchical/Regional Mixture

Many environments will involve a mixture of hierarchical and regional servers. So let's look at a more complicated example.

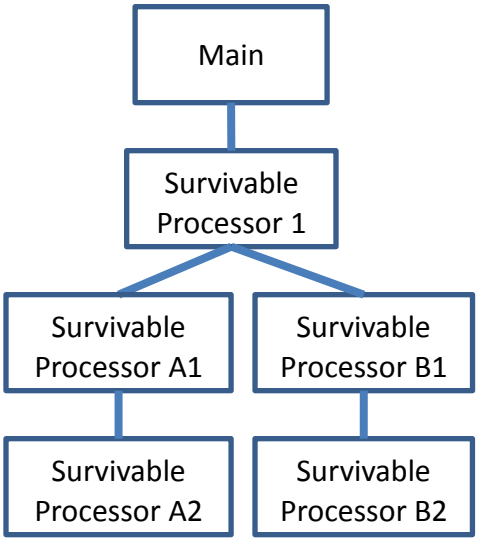

Accommodating such environments is merely a mixture of the two previous approaches. In this example, the Main and SP-1 are located where NR-1 is; SP-A1 and SP-A2 are located where NR-2 is; SP-B1 and SP-B2 are located where NR-3 is. You would like the whole system to converge on SP-1 if possible; otherwise, you would like the regional systems to converge on SP-A1 and SP-B1; SP-A2 and SP-B2 only come into play if the NR is isolated and SP-A1 or SP-B1 also fails. Note: This setup cannot be supported if the Split Registration Feature (SRP) is enabled because only one backup server is allowed per NR.

- H.323 endpoints[13]

	Backup Servers
All NR in NR-1	SP-1
All NR in NR-2	SP-1 SP-A1 SP-A2
All NR in NR-3	SP-1 SP-B1 SP-B2

[13] The list of backup servers for each NR should include where you want those H.323 endpoints to go for service in the event of a failure. Endpoints in NR-1 will go to SP-1. Endpoints in NR-2 will go to SP-1 if they can, or to SP-A1 if they cannot get to SP-1, or to SP-A2 if they can get to neither SP-1 nor SP-A1. Endpoints in NR-3 will go to SP-1 if they can, or to SP-B1 if they cannot get to SP-1, or to SP-B2 if they can get to neither SP-1 nor SP-B1.

- PN[14]

	Community	***Priority Score***	***Local Only***	***System Preferred***
All PN in NR-1	1			
All PN in NR-2	2			
All PN in NR-3	3			
SP-1	1	100	n	y
SP-A1	2	50	y	n
SP-A2	2	20		
SP-B1	3	50		
SP-B2	3	20		

- MGC List[15]

For all MG in NR-1	***For all MG in NR-2***	***For all MG in NR-3***
Main	Main	Main
SP-1	SP-1 SP-A1 SP-A2	SP-1 SP-B1 SP-B2

- MS[16]

	Priority wrt MS	***MSRL***
SP-1	2	At least one MS from each NR
SP-A1	3	Subset of the SP-1 MSRL that are in NR-2
SP-A2	4	
SP-B1	3	Subset of the SP-1 MSRL that are in NR-3
SP-B2	4	

[14] In example, we need three communities because we have three separate algorithms we want different PN to use to decide which SP to go in the event of a failure. Ideally, we'd prefer all PN go to SP-1 if they can. So we set SP-1 to the highest Priority Score and set it to "System Preferred" so that it will accept all PN regardless of community. Suppose, however, that the WAN between our three NR fails. In that case, we want the PN to go to their local SP. All the other SP are set to "Local Only" so that they accept only PN in the same community. SP-A1 and SP-B1 are given a higher Priority Score than SP-A2 and SP-B2 to indicate that they are to be preferred if both are available. A separate community is devoted to each region, and all PN and SP in that region are put in the same community.

[15] All MG want to have SP-1 in their MGC List immediately below the transition point because SP-1 is the preferred SP for all MG. In the event that an MG can get to neither the Main nor SP-1, however, they want the SP in their own region to be listed next, in priority order.

[16] SP-1 is given the highest priority (lowest number) with respect to MS to indicate that it is the first choice. Furthermore, the SP-1 MSRL is the union of the other MSRL so that all MS will be considered in SP-1's decision to go active. The other SP are given lower priority than SP-1 to indicate their place in the line of succession. Since SP-A1 and SP-B1 have disjoint MSRL, they can be given the same priority (or not—it doesn't really make any difference as long as each one is of lower priority than SP-1 and of higher priority than SP-A2 or SP-B2).

Review: Session Manager Administration

This section provides administration recommendations and guidelines.

Session Manager (SM) administration, via System Manager (SMGR), is unaffected by the introduction of MS. All previous considerations still apply. A brief review is included here.

- First, go to Home / Elements / Routing / SIP Entities. Here is where you will define a SIP Entity to represent the whole set of servers, i.e., the Main and all Survivable Processors (SP) that back it up. In addition to other settings, you will give this SIP Entity a Fully Qualified Domain Name (FQDN), such as "myCM.myCompany.com".
- Now go to Home / Elements / Session Manager / Network Configuration / Local Host Name Resolution (LHNR). The Main and each SP must be listed as separate IP addresses associated with the FQDN you administered for the SIP Entity. For each server, provide the IP Address, Port (of the signaling group—usually 5060 or 5061), Priority, Weight and Transport (TCP or TLS).
    - Typically, you would specify the PROCR address of each server as that server's IP address. Alternatively, you could use the CLAN addresses. This choice between these two approaches involves more concerns than survivability and is thus beyond the scope of this paper.
    - Start by giving the Main a Priority of 100. Lower numerical priority values indicate higher priority.
    - Then give each SP a unique priority greater than 100 (i.e., lower priority than the Main). If two SP back up different regions, then the relative order of the two priorities is immaterial, but they still should be unique in case SM can get to both. If one SP backs up another SP, then the higher priority SP should be given a lower priority number.
    - Because all servers have unique priorities, the Weight values are immaterial. You can set them all to 100.

Session Manager will send calls to the highest priority server in the FQDN that is responding positively to its OPTIONS messages. Servers that have no media resources (i.e., are inactive) will respond negatively to the Session Manager OPTIONS messages, which tells SM to send the calls elsewhere.

Consider the example below.

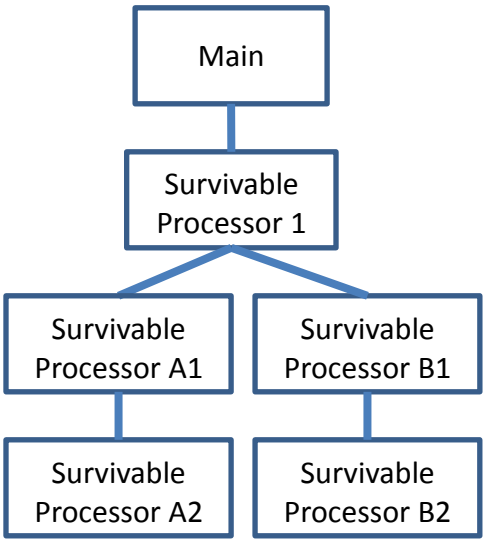

Here is one possible way (among many) you might configure the LHNR priorities and weights.

Host	IP Address	Priority	Weight
myCM.myCompany.com	Main	100	100
myCM.myCompany.com	Survivable Processor 1	200	100
myCM.myCompany.com	Survivable Processor A1	300	100
myCM.myCompany.com	Survivable Processor A2	500	100
myCM.myCompany.com	Survivable Processor B1	400	100
myCM.myCompany.com	Survivable Processor B2	600	100

Guidelines: Consistent Fallback Strategy for PN/MG/MS Environments

This section provides administration recommendations and guidelines.

In a system that includes mixtures of Port Networks (PN), Media Gateways (MG) and Media Servers (MS), separate recovery rules are used for each type of media resource. Most system managers will want to set the recovery rules to be consistent so that all media recovers at the same time. For all three types of media resources, the system manager has the option to set the recovery to be:

- automatic (i.e., as soon as possible and without manual intervention),
- scheduled (i.e., at a specific day and time, without manual intervention) or
- manual (i.e., not until administrator commands are entered).

The MG has two extra types of recovery that will not be discussed here: “0-active-calls” and “time-window-OR-0-active-calls”. Neither the PN nor MS recovery rule has an equivalent.

The recovery rules are set:

- On the “change system-parameters port-networks” form for PN. This one rule applies to all PN.
- On the “change system-parameters mg-recovery-rule” form for MG. You may create different rules for each MG if you desire. Split Registration Prevention (SRP), however, puts some limits on this.
- On the “change system-parameters ms-recovery-rule” form for MS. This one rule applies to all MS.

There is no administrable recovery rule for H.323 phones. Phones recover when the media resources recover.

Automatic Recovery

If you prefer automatic recovery:

- For PN, set “Auto Return” to “yes” on “change system-parameters port-networks”.
- For MG, set “Migrate H.248 MG to primary” to “immediately” on “change system-parameters mg-recovery-rule”. Also, set “Recovery Rule” to this recovery rule number on “change media-gateway” for the MG to which you want this rule to apply.
- For MS, set “Auto Return” to “yes” on “change system-parameters ms-recovery-rule”.

Scheduled Recovery

If you prefer scheduled recovery:

- For PN, this setting must be re-administered after each failover (i.e., each administration only allows one recovery). Set “Auto Return” to “scheduled” on “change system-parameters port-networks”. Select a day of the week and a time of day. This rule applies the next time this day and time come around. For example, if it is Wednesday at 13:00 when you submit this form, and you select Thursday at 02:00 for your scheduled recovery, the recovery will be initiated 13 hours after the time you type the command. The recovery rule is executed and then cancelled on Thursday at 02:00 (whether or not a recovery is needed at that time); the rule reverts to manual after that time.
- For MG, set “Migrate H.248 MG to primary” to “time-day-window” on “change system-parameters mg-recovery-rule”. Then put an “X” for each hour in which scheduled recovery is allowed. You need

at least one “X”. Unlike the PN rule, this rule is persistent. E.g., if you put an “X” for Thursday at 02:00, recovery will be allowed *every* Thursday at 02:00, not just the next one. Furthermore, you do not need to elect it again after the window passes. Also unlike the PN rule, you can select multiple recovery windows (e.g., you could choose every day at 02:00). Finally, set “Recovery Rule” to the recovery rule number on “change media-gateway” for the MG to which you want this rule to apply.

- For MS, set “Auto Return” to “scheduled” on “change system-parameters ms-recovery-rule”. Select a day of the week and a time of day. Like the MG rule, this is a one-hour window. Like the PN rule, you can only select one window. Unlike the PN rule, this rule is persistent. E.g., if you select Thursday at 02:00, recovery is allowed *every* Thursday between 02:00 and 02:59. You do not need to elect it again after the window passes.

Manual Recovery

If you prefer manual recovery:

- For PN, set “Auto Return” to “no” on “change system-parameters port-networks”.
- For MG, set “Recovery” to “none” on “change media-gateway” for the MG to which you want this rule to apply. Alternatively, set the “Migrate H.248 MG to primary” field of the recovery rule to blank on the “change system-parameters mg-recovery-rule” form.
- For MS, set “Auto Return” to “no” on “change system-parameters ms-recovery-rule”.

To initiate manual recovery, you will normally type all three of these commands back to back. It doesn’t matter in which order you type them.

- get forced-takeover ipserver-interface all
- enable mg-return all
- enable ms-return

Manual return can be used even if you have the recovery rules set for something else. This is a way of, for example, forcing the fallback without waiting for the scheduled time, or without waiting for the automatic time delay.

If you use Split Registration Prevention (SRP), you should enter “enable mg-return all” as a part of recovery, even if you do not have any MG on your system. This is necessary to re-enable the disabled network regions so that the H.323 phones can come back.

Information: Recovery from one SP to another SP

Although we generally talk about recovery from a Survivable Processor (SP) to the Main, it is also sometimes possible to recover from one SP to another SP. The considerations are different for Port Networks (PN), Media Gateways (MG) and Media Servers (MS). The introduction of MS expands the options available.

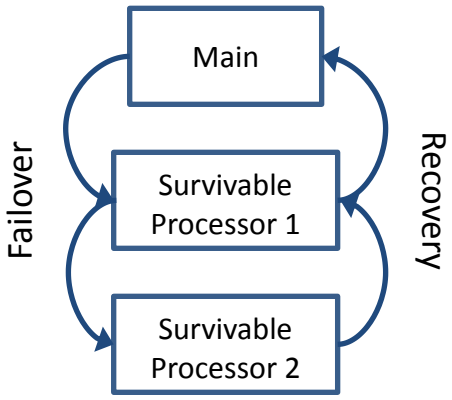

The only way MG can recover from one SP to another SP is if the configuration includes PN. Furthermore the MG must have been configured to register through a CLAN (rather than through a PROCR). Note: This is not necessarily the recommended approach. MG recovery is based on the MGC List. Once an MG fails over to below the transition point for service, it can only recover to something above the transition point (i.e., it can failover again to someplace lower, but it cannot recover to someplace higher on the list but still below the transition point). Generally the addresses above the transition point are the Main. If these are addresses of CLAN, however, and if the PN have failed over to an SP, then that SP is now above the transition point and can be recovered to. If the recovery rule is not automatic or scheduled, you would still need to manually enable the recovery on the SP. If the address above the transition point is the Main PROCR, however, then MG recovery can only be done to the Main.

A PN can only recover to the Main, never to another SP.

MS recovery is a simpler and more flexible concept. An MS can recover to a higher-priority SP just as easily as it can recover to the Main. There is no transition point, and no special consideration is given to the Main. Recovery to a higher-priority SP is based on the same rules as for recovery to the Main. Given a choice of servers to recover to, a MS will always recover to the highest priority server available. MS, however, can only recovery to a server that includes it on the MSRL (the Main MSRL implicitly includes all MS that are on any SP MSRL).

Failover Speed for MS: Report Interval and Report Expiration

This section provides administration recommendations and guidelines.

The Report Interval (RI) and Report Expiration (RE) values on the Media Server Recovery Rules ("change system-parameters ms-recovery-rule") can be used to establish failover speed. Fast failover reduces the length of a service disruption, but it does so at the risk of an increased number of disruptions caused by failovers and fallbacks.

Failover will occur between (RE – RI) and (RE + RI) seconds after the failure. Slow failover is accomplished by setting RE high. Fast failover is accomplished by setting RE low. Setting RI low will narrow the size of the failover window at the cost of increased network traffic and processing time. RE must always be at least one second larger than RI. A 3-to-1 RE-to-RI ratio is recommended to reduce risk of false failover.

Let's consider some examples.

- **Very rapid failover (RI=1, RE=2)**: Setting values this low is not recommended except in unusual situations. Failover will occur between 1 and 3 seconds after a failure. However, failover might and probably will also occur during Main restarts, Main interchanges, network congestion, Media Server congestion, etc. You may not want this.
- **Moderate failover (RI=60, RE=180)**: These values are default, and are the recommended values for most situations. Failover will occur between 120 and 240 seconds after a failure. This allows the system to accommodate most temporary conditions without a failover. Most system managers prefer to keep the media resources and H.323 endpoints at the Main in the event of a Main reboot or restart, or in the event of a short network outage.
- **Slow failover (RI=60, RE=400)**: These values can be used if you want to prevent a failover except in the case of persistent failures. Failover will occur between 340 and 460 seconds after a failure.

PN and MG also have parameters to control failover speed. These parameters are covered in the next section.

Failover Speed for PN, MG and Phones

This section provides administration recommendations and guidelines.

The previous section discussed how to control failover speed for Media Servers (MS). If your environment also includes Port Networks (PN), Media Gateways (MG) and/or H.323 phones, you probably want to set parameters for these to create approximately the same failover speed. How to do that is covered here.

PN failover speed is controlled by the "No Service Time Out Interval" on the "change system-parameters port-networks" form. It can be set from 2 to 15 minutes, with a default of 5 minutes. With the default value, the PN will failover after five minutes. Set it lower for faster failover and higher for slower failover. If you want it to approximately match the default MS failover speed, set it to 3 (minutes). This time is in addition to the link loss detection time.

MG failover speed is controlled by a timer administered on the MG itself. Use "set reset-times primary-search". The default is one minute. If you want it to approximately match the default MS failover speed, set it to 3 (minutes). This time is in addition to the link loss detection time (about 30 to 40 seconds).

Failover speed for H.323 phones is controlled by "Primary Search Time" on the "change system-parameters ip-options" form. It defaults to 75 seconds. If you want it to approximately match the MS failover speed, set it to 180 seconds. This time is in addition to the link-loss detection time (about 25 to 45 seconds if you use default values on the "change ip-network-region" form). You might, however, want to set the primary search time a little longer than 180 seconds (maybe 30 seconds longer) in order to give the media resources time to failover before you send the phones to the Survivable Processor (SP), since the SP is not going to accept the phone registrations until it has media.

Recovery Speed

This section provides administration recommendations and guidelines.

If you are using manual recovery, recovery begins soon after you type the recovery command. If you are using scheduled recovery, recovery begins as soon as the scheduled time arrives. So this discussion of recovery speed only applies to automatic recovery.

Automatic recovery speed for Media Servers (MS) depends on the "Time Delay" value on the "change system-parameters ms-recovery-rule" form. By default, it is zero minutes, which facilitates rapid recovery. If you are concerned, however, that your network may bounce during recovery, you can set this up as high as 120 minutes. Recovery won't take place until the server has continuous access to the MS for as long as specified. If you want it to approximately match Media Gateway (MG) and Port Network (PN) recovery, you'll need to set it to the same values you use for those (see below).

Automatic recovery speed for MG depends on the "Minimum time of network stability" value on the "change system-parameters mg-recovery-rule" form. By default, it is 3 minutes. You can set it anywhere from 3 to 15 minutes.

Automatic recovery speed for PN is controlled by the "IPSI Connection Up Time" on the "change system-parameters port-networks" form. There is no default. You can set it anywhere from 3 to 120 minutes.

Recovery of a H.323 phone does not happen until two other things both happen:

- The Survivable Processor (SP) must evict it, which it doesn't do until all the media resources have fallen back to some other server.
- The server that the phone is recovering to must allow it to register, which it doesn't do until it has at least some media resources to serve the Network Region (NR) of the phone.

Failover Strategy: Priority with respect to Media Servers and MSRL

This section provides administration recommendations and guidelines.

When planning your failover strategy, it is best to start by drawing a diagram of how you want the failover to occur. Below is an example of such a diagram. In this example, the Main and SP-1 are located where NR-1 is; SP-A1 and SP-A2 are located where NR-2 is; SP-B1 and SP-B2 are located where NR-3 is. You would like the whole system to converge on SP-1 if possible; otherwise, you would like the regional systems to converge on SP-A1 and SP-B1; SP-A2 and SP-B2 only come into play if the NR is isolated and SP-A1 or SP-B1 also fails.

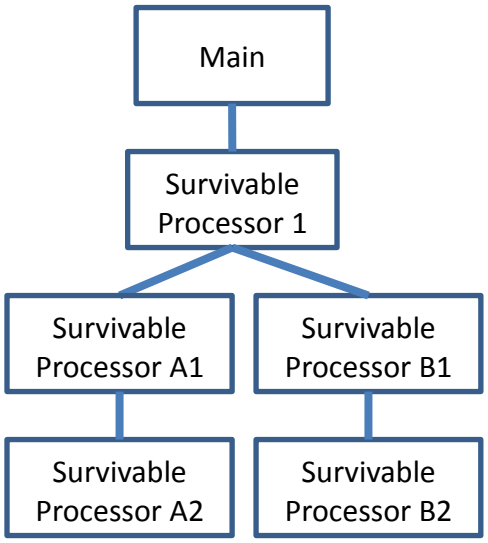

Each Survivable Processor (SP) is assigned a Priority with respect to Media Servers. This value can be from 2 to 9999. Lower numerical values are higher priority. Priority 1 (the highest priority) is reserved and implicitly assigned to the Main. The range is broad so that you can leave unused numbers between each SP to make it easier to add new SP in the future.

The priority can also be blank. Any SP with a blank priority will never become active due to Media Servers. That SP might still become active, however, if a PN or MG connects to it.

Each SP with a non-blank priority is also given a Media Server Reporting List (MSRL). The MSRL defines which MS should be used to determine when that SP goes active. Normally you would include one or more MS from each NR for which this SP is acting as a backup. At least two MS from each NR is recommended for redundancy (e.g., in case one MS were to fail). If desired, you can put up to all MS in the MSRL, at the cost of additional administration and additional network traffic. This is not recommended, however, if you have a lot of MS (because it increases server load and network traffic) or if you have a lot of SP (because it can create a large number of control sessions to the MS).

See “Hierarchical/Regional Mixture” on page 23 for suggested settings of Priority with respect to Media Servers and MSRL for the above example.

Here are some general recommendations:

- For an SP that you want to back up the entire network (e.g., Survivable Processor 1 in the diagram above), include in the MSRL at least two (in case one MS is down) MS from each NR in the enterprise.
- For an SP that you want to back up only part of the network (e.g., Survivable Processors A1, A2, B1 and B2 in the diagram above), include in the MSRL at least two MS from each NR in that part of the network.
- When two or more SP have the same or overlapping MSRL (e.g., Survivable Processor A1 and Survivable Processor A2 in the diagram above), give each SP a different Priority with respect to Media Servers. The higher-priority (smaller numerical value) SP will handle the failover in preference to the lower-priority (larger numerical value) SP.
- If multiple SP have the same Priority with respect to Media Servers (e.g., Survivable Processor A1 and Survivable Processor B1 in the diagram above), the MSRL for these SP should be disjoint, i.e., have no MS in common. If you violate this rule, the failover results may be unpredictable.
- To avoid ambiguity in case of misadministration of MSRL, it is safer to never give two SP the same Priority with respect to Media Servers.

Recommendation: Survivable Processor OOS for an Extended Time

This section provides procedural cautions.

This section gives recommendations that allow a booting Main to quickly discover whether it should be active or inactive. If an SP is already active when the Main boots, the Main should stay inactive until the recovery rules allow it to be active. But if no SP is active when the Main boots, the Main should quickly go active to minimize service disruption. The decision to go active can be delayed if the Main is unable to discover the state of one or more SP.

When the Main boots, it discovers the SP states by asking each MS on its Media Server Reporting List (MSRL) for status. Once the Main has discovered that all SP are inactive, it becomes active. If, however, one of the SP is down at the time the Main boots, the Main must wait a “reasonable time” for that SP to report before deciding that it is inactive. That “reasonable time” is determined to be the “Report Expiration” time on the Media Server Recovery Rule (default 180 seconds). If the SP has still not reported after that amount of time, the Main assumes that the SP is inactive and the Main becomes active.

What does this mean? In order to allow the Main to become active as soon as possible after booting, it is recommended that, if an SP is to be out of service for an extended period of time, you should set its “Priority with respect to Media Servers” (on the “change survivable-processor” form) to blank. This tells other servers not to expect reports from it. When you bring the SP back into service, you will need to re-administer the MSRL.

Information: Reboot of an Active SP

When an active Survivable Processor (SP) is rebooted while a higher-priority server (usually the Main) and the network are both healthy, the Media Servers (MS) being used by that SP will usually fall back to the Main without waiting for a recovery rule. This assumes Split Registration Prevention (SRP) has not auto-disabled the Network Region (NR) of the MS. This may cause an unexpected fallback.

Information: Media Server Service States

The Main will bring every Media Server (MS) into service that it can. The Survivable Processors (SP), however, won't. The reason for this is that bringing MS into service from many SP consumes resources that could otherwise be used for call handling by that MS. So each SP only brings into service those MS that it needs. The chart below shows which MS an SP will bring into service.

	SP Inactive	SP Active
MS on the SP Media Server Reporting List (MSRL)	Yes	Yes
MS in Network Regions (NR) backed up by this SP	No	Yes
Other MS	No	No

MS for which there is a "No" in the table above are shown with a state of "N/A" on "list media-server" and a state of "not-applicable" on "status media-server". "N/A" may also be shown on an SP in the first minute or two after the SP boots, while the SP is figuring things out.

```
list media-server

                              MEDIA SERVERS REPORT

                            License  Dedicated  Licenses
Num Sig-grp     Node-Name     Limit    Licenses   In-Use     NetRgn  Loc    State
1     301     ams1                                 0          1              N/A
2     302     ams2                                 0          2              N/A
3     303     ams3                                 0          2              INS
4     304     ams4                                 0          3              N/A
5     305     ams5                                 0          4              N/A
6     306     ams6                                 0          2              N/A
```

```
status media-server 2                                          Page   1 of   2
                              MEDIA SERVER STATUS


                    Media Server Number: 2
                                  State: not-applicable
                         Signaling-group: 302
                              Node Name: ams2
                             IP Address: 10.129.130.137
                          Network Region: 2
                             SW-Version:
              Voip Channel License Limit:
        Dedicated Voip Channel Licenses:
           Voip Channel Licenses in-use: 0
                           Load Factor: 0
             Estimated Channel Capacity: 0
                  Announcements Present: 0
```

Information: Split Registration Prevention

This section provides a general understanding of what to expect from this feature.

Overview

The Split Registration Prevention (SRP) feature ("Force Phones and Gateways to Active Survivable Servers" on the "change system-parameters ip-options" form) comes into play after a failure and the subsequent repair of that failure, but before the fallback has been initiated. SRP is designed to prevent situations in which some of the Media Gateways (MG) and H.323 phones are registered on the Main and others are registered on a Survivable Processor (SP). SRP also prevents some Media Servers (MS) from being used by the Main and some by the SP in the same Network Region (NR). SRP applies to neither Port Networks (PN) nor SIP phones.

SRP is normally used in conjunction with scheduled or manual recovery; it can, however, play a more limited role with automatic recovery. SRP causes the Main to reject registration attempts from phones and MG until the fallback has been initiated. SRP only functions when the network is healthy enough for the SP to be able to register with the Main.

SRP was designed primarily for use with MG. Its use is discouraged on systems with only Port Networks (PN), and has some limitations on systems that include PN along with MG and/or MS. With the introduction of Media Servers (MS), SRP functions in similar way as with MG, either when using MS alone or when using MS in combination with MG and/or PN.

SRP works between the Main and an active SP. It can also work between a higher-priority SP and a lower-priority active SP if the lower-priority SP has registered to the higher-priority SP via a CLAN (which requires that the PN containing the CLAN has connected to the higher-priority SP). You cannot, however, use SRP between the Main and SP-1 and between SP-1 and SP-2 in the same configuration.

SRP Example

Here's an example of how SRP works. Assume the following configuration, with manual fallback.

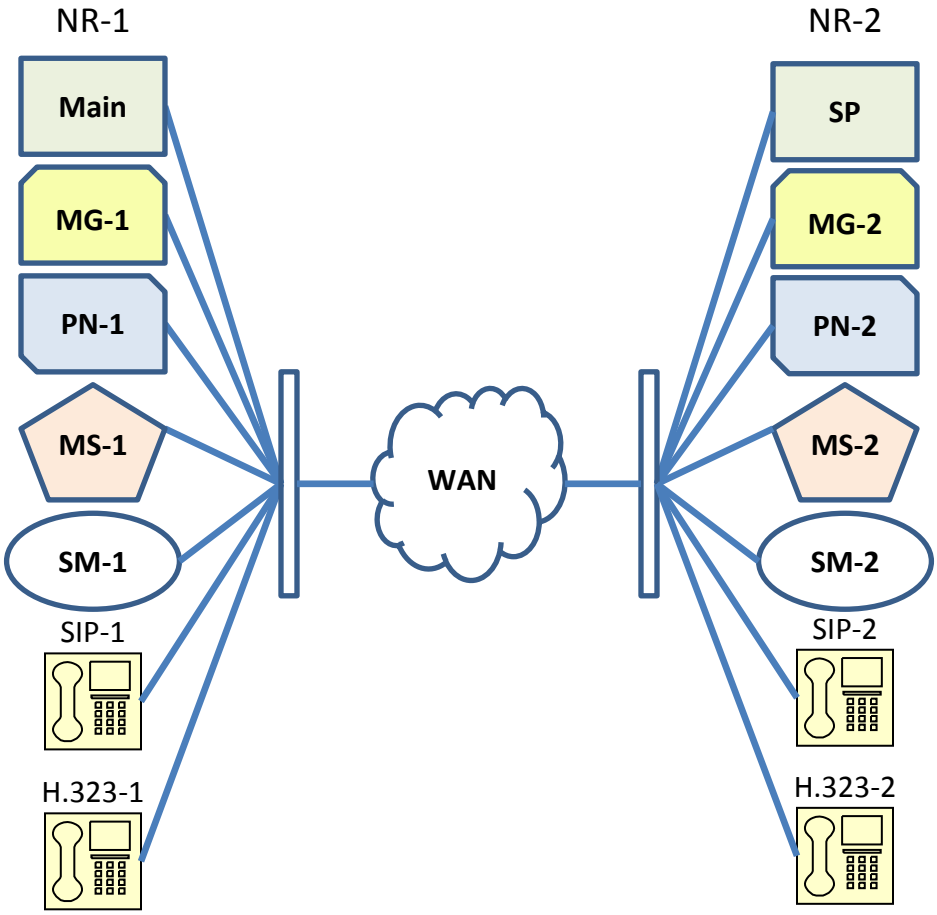

The following chart walks you through, from top to bottom, a failover, repair and fallback. It shows what happens to each resource at each step, with and without SRP. As you can see, SRP only comes into play between the repair and the fallback. Refer to the diagram above.

<table>
<tr><th></th><th>SRP</th><th>H.323 Phones</th><th>SIP Phones</th><th>PN</th><th>MG</th><th>MS</th></tr>
<tr><td>WAN fails, splitting the system by NR. All phones still have service within their NR.</td><td>N/A</td><td>H.323-2 reregisters with SP. H.323-1 stays registered with the Main.</td><td>SIP-1 stays registered with SM-1, which sends its calls to the Main. SIP-2 stays registered with SM-2, which sends its calls to the SP.</td><td>PN-2 connects to SP. PN-1 stays with the Main.</td><td>MG-2 reregisters with SP. MG-1 stays with the Main.</td><td>SP starts using MS-2. Main is still using MS-1.</td></tr>
<tr><td rowspan="3">WAN is repaired. SP registers with the Main, telling it that it is active.</td><td>Off</td><td colspan="5">Nothing changes from the above states.</td></tr>
<tr><td rowspan="2">On</td><td colspan="5">Main auto-disables NR 1, consolidating everything except PN-1 and SIP phones onto the SP</td></tr>
<tr><td>H.323-1 reregister with SP.</td><td>Both SM start sending all calls to the Main (because Main still has PN-1).</td><td>Nothing changes.</td><td>MG-1 reregisters with the SP.</td><td>Main stops using MS-1. SP starts using MS-1.</td></tr>
<tr><td>Fallback is manually initiated.</td><td>N/A</td><td>Both phones reregister with the Main.</td><td>Nothing changes.</td><td>PN-2 reconnects with the Main.</td><td>Both MG register with the Main.</td><td>SP stops using both MS. Main starts using both MS.</td></tr>
</table>

Precaution: Removing a MS from a SP MSRL

This section gives procedural cautions.

When you remove a Media Server (MS) from a Survivable Processor's (SP) Media Server Reporting List (MSRL), you run the risk of accidentally causing that SP to go active. This is because, between the time you make this change and the time you propagate the translations to the SP (using "save translation all"), the SP may see that the Main has stopped reporting and think that it has died. This can happen if you either (1) blank out one of the MS on the MSRL, or (2) blank out the "Priority with respect to Media Servers" (because that effectively clears all entries from the MSRL).

This is only a concern if the MS that you are removing from an MRSL is the last MSRL that that MS is on. E.g., if you remove MS-3 from the MSRL of SP-1, but MS-3 is still on the MSRL of SP-2, no precautions are needed. Use "status media-server" to see which MSRL a MS is on, but don't count the Main.

There are steps you can take to avoid this unintended failover, and steps you can take to recover from it if it happens.

Recovering from an SP that accidentally went active

If you don't care if the SP goes active temporarily (and you may not—most of the time it causes no harm[17]), then you can simply let it happen and recover from it. Here's how to recover:

1. Make your change from the Main SAT, using the "change survivable-processor" command.
2. As soon as possible, execute "save translation all" on the Main SAT.
3. Use "list survivable-processor" to monitor what happens over the next ten minutes or so.
    a. If you see that the "Translations Updated" has a new timestamp and the "Act" column stays "n", then you have avoided the problem and no recovery is necessary.
    b. If you see that the "Act" column has gone to "y", then you will need to recover.
        i. Log onto the SP.
        ii. Reboot the SP with a "reset system 4". The SP will reboot twice.
        iii. After it comes back up, go back to the Main SAT and execute "list survivable-processor".
        iv. Wait until that SP has registered (the "Reg" column is "y").
        v. Verify that it is not active (the "Act" column is "n") and that the "Translations Updated" has a new time stamp.
        vi. If the SP goes active again (the "Act" column is "y"), you may need to reboot the SP again (i.e., you rebooted it the first time before the translations were propagated).

Note: You cannot recover from this by executing "enable ms-return". Use the procedure above.

Avoiding having an SP accidentally go active

If you'd like to avoid the problem altogether, then you need to take steps to ensure that the SP cannot go active in the first place. There are two ways to do this.

[17] Caution: It may cause harm if you are using Split Registration Prevention because the Main may disable NR.

Avoidance Procedure A: Busyout the MS signaling group

One way to avoid this is as follows. Suppose you need to remove MS-3 from the last MSRL it is on.

1. Log onto the SP SAT.
2. Busyout the signaling group for MS-3.
3. Go back to the Main SAT.
4. Make your change from the Main SAT, using the "change survivable-processor" command.
5. As soon as possible, execute "save translation all" on the Main SAT.
6. The SP will automatically reboot, picking up the new translations and automatically releasing the signaling group.

Avoidance Procedure B: Stop the MS

Here's another possible avoidance:

1. Use the MS Element Manager to stop the MS.
2. Make your change from the Main SAT, using the "change survivable-processor" command.
3. As soon as possible, execute "save translation all" on the Main SAT.
4. Use "list survivable-processor" to wait until the translations get updated on all SP.
5. Use the MS Element Manager to start the MS.

Precaution: Busyout of a MS Signaling Group

This section gives procedural cautions.

When you busyout a Media Server (MS) signaling group, you run the risk of causing one or more Survivable Processors (SP) to go active.

Recovering from an SP that accidentally went active

If one or more SP goes active while a MS signaling group is busied out (verified with "list survivable-processor"), here's how to recover.

Recovery Procedure A: Release the signaling group

1. Release the signaling group. If you want to recover without doing this, then use "Recovery Procedure B: Remove the MS from all MSRL".
2. Enter "enable ms-return".

Recovery Procedure B: Remove the MS from all MSRL

1. Use "status media-server" to find out which SP have this Media Server on their MSRL.
2. Use "change survivable-processor" for all the SP identified.
3. Remove the MS for which the signaling group is busied out from the MSRL.
4. Enter "save translation all".
5. Follow the recovery procedure step 3 for "Recovering from an SP that accidentally went active" on page 41.

Avoiding having an SP accidentally go active

Avoidance Procedure A: Don't leave the MS busied out very long

If you only need to leave the signaling group busied out briefly:

1. Busyout the MS signaling group.
2. Quickly make whatever changes you need.
3. Release the MS signaling group less than two minutes (Report Expiration less Report Interval) after you busied it out.

If you need a little more than two minutes (up to six minutes), you can increase the time available between busyout and release by temporarily increasing the "Report Expiration" to the maximum value (400 seconds), and decreasing the "Report Interval" to 10 seconds. Then wait at least a minute (the previous "Report Interval") before busying out the signaling group. This will give you more than six minutes to release the signaling group after busying it out. You can then change the Report Interval and Report Expiration back to their previous values. You do not need to do a "save translation all" during this process.

Avoidance Procedure B: Stop the MS

If you need to leave the signaling group busied out for a long time:

1. Use the MS Element Manager to stop the MS.
2. Busyout the MS signaling group.
3. Make whatever changes you need.
4. Release the MS signaling group.
5. Use the MS Element Manager to start the MS.