Application Note

January 2022

Business Continuity Best Practices: Resiliency, High Availability and Disaster Planning with a Mitel IP-PBX

Description: This Application Note describes best practices and provides detailed information on configuring the fault tolerance, high-availability, resiliency, and redundancy features of the Mitel Unified Communications system.

Environment: Mitel IP-PBX versions 9 – 14.2.

Table of Contents

OVERVIEW **5**

VERSION **6**

TARGET AUDIENCE **7**

TESTING ENVIRONMENT **8**

DOCUMENT AND SOFTWARE COPYRIGHTS **9**

TRADEMARKS **10**

REVIEW OF THE MITEL COMPONENTS AND THE MITEL DISTRIBUTED ARCHITECTURE **11**

MITEL COMPONENTS 11
MITEL IP PHONES 11
SHOREGEAR VOICE APPLIANCES 12
MITEL VIRTUALIZED APPLIANCES 13
MITEL VIRTUAL IP PHONE SWITCH 13
MITEL VIRTUAL SIP TRUNK SWITCH 13
MITEL VIRTUAL APPLICATION SERVICE APPLIANCE 14
MITEL APPLICATION SERVICES END POINTS 14
MITEL COMMUNICATOR 14
MITEL'S DISTRIBUTED ARCHITECTURE 15

BUSINESS CONTINUITY DEFINITIONS **16**

DESIGN FEATURES, BENEFITS AND BEST PRACTICES **18**

SHOREGEAR VOICE APPLIANCES 18
FULL MESH CONNECTIVITY AND DISTRIBUTED INTELLIGENCE 19
LOCATION SERVICE PROTOCOL (LSP) 20
EXTENSION DIALING 20
TRUNK DIALING 22
LSP SUMMARY 22
SELECTION OF PRIMARY AND SECONDARY END POINTS FOR APPLICATION SERVICES 23
SELECTING THE PRIMARY AND SECONDARY APPLICATION SERVICES END POINTS 23
SELECTION EXAMPLES 25
BACKUP AUTO ATTENDANT 27
PEER-TO-PEER AUDIO PATH 28
DESIGN FEATURES, BENEFITS AND BEST PRACTICES: SUMMARY 31

SINGLE SITE REDUNDANCY **32**

SINGLE SITE – DISTRIBUTED CALL CONTROL 32
SINGLE SITE – REDUNDANT CALL CONTROL 34
ANALOG PHONE AND SIP PHONE FAILOVER 38
POWER FAIL TRANSFER PORTS 38
SINGLE SITE REDUNDANCY: SUMMARY 40

APPLICATION SERVICES **41**

EXPLICIT BACKUP EXTENSIONS 42
REPLICATED SERVICES: AUTO ATTENDANTS AND ACCOUNT CODE COLLECTION 44
DISTRIBUTED SERVICES: WORKGROUPS 45
DISTRIBUTED SERVICES: VOICE MAIL 45
VOICE MAIL FAILOVER: STORE AND FORWARD 46
SERVICES THAT FAILOVER: CALL DETAIL RECORD (CDR) COLLECTION 47
SERVICES THAT DO NOT FAILOVER 48
PAGING GROUPS 49
COMMUNICATOR CLIENT SERVICES 49

COMMUNICATOR CLIENT SERVICES: SUMMARY 55
ADMINISTRATION AND REPORTING 57
APPLICATION SERVICES: SUMMARY 57

MULTI-SITE REDUNDANCY **58**

MULTI-SITE CONFIGURATION 59
REMOTE SITE SURVIVABILITY – CALL CONTROL AND PSTN ACCESS 60
REMOTE SITE SURVIVABILITY – APPLICATION SERVICES 61
REMOTE SITE SURVIVABILITY – REDUNDANT CALL CONTROL 62
REMOTE SITE SURVIVABILITY – IP PHONE CONFIGURATION SWITCHES 64
MULTI-SITE REDUNDANCY: SUMMARY 64

DATABASE CHANGES AND SERVER REDUNDANCY **65**

'CALL FORWARD ALWAYS' OVERRIDE (*72/*73) 66
DISTRIBUTED DATABASE 68
CHANGING CALL HANDLING MODES WITH DISTRIBUTED DATABASE ENABLED 69
DISTRIBUTED DATABASE VS. DISTRIBUTED WORKGROUPS 69
WORKGROUP AGENT LOGIN STATE 69
SERVER REDUNDANCY 69
HQ SERVER REDUNDANCY WITH DOUBLE-TAKE 71
VMWARE HIGH AVAILABILITY (HA) 72
MULTIPLE SERVER REDUNDANCY WITH MICROSOFT HYPER-V 72
DATABASE CHANGES AND SERVER REDUNDANCY: SUMMARY 74

FAILURE SCENARIOS **75**

TOPOLOGY 75
FAILURE SCENARIOS 76
IF THE HQ SERVER GOES OFF-LINE, WHAT *STOPS* WORKING? 76
IF THE HQ SERVER GOES OFF-LINE, WHAT *CONTINUES* TO WORK? 76
IF THE WAN BETWEEN SITES FAILS, WHAT *STOPS* WORKING? 77
IF THE WAN BETWEEN SITES FAILS, WHAT *CONTINUES* TO WORK? 77
IF A DVS FAILS, WHAT *STOPS* WORKING? 78
IF A DVS FAILS, WHAT *CONTINUES* TO WORK? 79
WHAT IF A SHOREGEAR SWITCH FAILS? 79
WHAT IF A TRUNK CIRCUIT FAILS? 79
CAN I IMPROVE ON THIS DESIGN? 80

NETWORK CALL ROUTING **81**

OUTBOUND CALL ROUTING 81
CALL SCENARIOS 82
INBOUND CALL ROUTING 83
ANALOG TRUNKS 84
DIGITAL TRUNKS 84
PARENT AS PROXY 86
CALL COST PROMOTION 111
PSTN FAILOVER 112

SITE TREE HIERARCHY **115**

4, 5, AND 6-PARTY CONFERENCING 116
SPARE SWITCH SELECTION 118
SELECTION OF DRS SERVERS 118

SCHEDULES **119**

THE APPLICATION OF SCHEDULES 121
HUNT GROUP SUMMARY: 121
WORKGROUP SUMMARY: 121
AUTO ATTENDANT SUMMARY: 121
ACCOMMODATING DIFFERENT TIME ZONES 121

HUNT GROUPS: 122
WORKGROUPS: 122
AUTO ATTENDANT MENUS: 122
SCHEDULES APPLIED TO CALL HANDLING MODES 123
MITEL 11 AND HIGHER 123
MITEL 8.1, 9.X, 10.X 123

DISASTER RECOVERY PLANNING 125

ACCOMMODATING INACCESSIBLE SITES 125
ACCOMMODATING DOWN SITES 126
PREVENTING TOTAL LOSS OF SERVICES 126
EXTENSION REASSIGNMENT AND SOFTPHONE CONSIDERATIONS 127
DISASTER RECOVERY (DR) SITES 127
HOT-STANDBY SITES 127
ECC REDUNDANCY BEST PRACTICES 128

DISTRIBUTED ROUTING SERVICE 129

EXTENSION CALL ROUTING WITH DRS 129
PSTN CALL ROUTING WITH DRS 129
SELECTING DRS SERVERS 130
EXTENT OF DRS INVOLVEMENT IN CALL CONTROL 130
MAINTENANCE PAGES WHEN DRS IS ENABLED 130
DESIGN CONSIDERATIONS WITH DRS ENABLED 130

ADDITIONAL CONSIDERATIONS 133

EVENT FILTERS AND ALARM NOTIFICATION 133
ENTERPRISE CONTACT CENTER 133
SERVICE PPLIANCES 133
REDUNDANT POWER SUPPLIES 141

CONCLUSION 142

RESOURCES, REFERENCES AND FURTHER READING 143

APPENDIX A: HQ, DVS AND V-SWITCH FEATURES 145

APPENDIX B: VOICE MAIL PROMPT BEHAVIOR 148

VOICE MAIL PROMPTS 148

APPENDIX C: QUESTIONS AND ANSWERS 151

REBOOTING QUESTIONS 151
COLOCATION SITE QUESTIONS 152
COMMUNICATOR QUESTIONS 152
HUNT GROUP QUESTIONS 152

APPENDIX D: LSP TABLES 153

LSP COMMANDS 153
LSP EXAMPLES 153
LSPCONLIST 156
LSPTELLIST 156
LSPTELLIST 1 157
LSP STATUS DURING AN OUTAGE 160

OVERVIEW

Every enterprise, whether large or small, places a heavy reliance on their communication infrastructure to conduct and grow their business. It is imperative that your Unified Communication (UC) system provide extremely high levels of reliability, survivability and functionality, even when faced with outages, faults and unforeseen disasters.

The distributed architecture of the Mitel Unified Communication solution is uniquely tailored to exceed these expectations and is built with reliability, resiliency, survivability and fault tolerance in mind. Like any advanced system, proper planning, deployment strategies and best practice configurations are necessary keys to installing and maintaining a fully survivable and highly-available Mitel UC system.

This Application Note discusses the myriad Business Continuity-oriented features and functions of the Mitel UC system and describes when, and how, to properly deploy and configure these features.

VERSION

Version	Date	Contributor	Content
1.0	Might, 2012	J. Rowley	Original Published App Note
1.1	Might, 2012	J. Rowley	Minor grammatical updates
1.2	June, 2012	J. Rowley	Updates to Communicator protocols sections for Mitel 13
2.0.1	November, 2012	J. Rowley	Clarified and corrected Communicator for Windows language
2.0.2	December, 2012	J. Rowley	Typographical updates
3.0	Might, 2014	J.Rubio	Mitel 14.2 updates
3.1	November, 2014	J.Rubio	Incorporated feedback
3.2	February, 2015	J.Rubio	Additional edits, grammatical updates

TARGET AUDIENCE

This document is an advanced Application Note and is intended for the experienced Mitel deployment engineer.

Familiarity with the Mitel components, Mitel's distributed architecture, and deployment strategies for a multi-site Mitel system is presumed.

This Application Note is intended to be a supplement to:

- Mitel Maintenance Guide.

Refer to the Maintenance Guide for additional detail on many of these topics.

For more information on basic Mitel installation and administration tasks and procedures, see the complete set of manuals for your version, including:

- Mitel Planning and Installation Guide
- Mitel Administration Guide

Additional references and resources are located in the 'Resources, References and Further Reading' section at the end of this document.

TESTING ENVIRONMENT

Validation testing for the features and best practices in the document were performed predominantly with:

- Mitel version 12.3, 13.3 & 14.2

Not all features are present in all versions and not all features can be combined simultaneously. Where appropriate, version differences are cited, and mutually exclusive features are identified.

DOCUMENT AND SOFTWARE COPYRIGHTS

TRADEMARKS

REVIEW OF THE MITEL COMPONENTS AND THE MITEL DISTRIBUTED ARCHITECTURE

Before embarking on a discussion of the resiliency and fault-tolerant elements of the Mitel Unified Communications system, it is important to review and understand the core components and architectural designs of a multi-site Mitel deployment. This section serves as an in-depth review for the experienced Mitel installation engineer.

MITEL COMPONENTS

The Mitel Unified Communication System is made up of several components:

- Mitel IP Phones
- Mitel Voice Appliances (also known as ShoreGear voice switches, or simply 'switches', we use the term appliance and switch interchangeably throughout this document) Mitel Virtualized Appliances
    - Mitel Virtual IP Phone Switch
    - Mitel Virtual SIP Trunk Switch
    - Mitel Virtual Application Service Appliance
- Mitel Application Services end points (servers and V-switches)Mitel Communicator end user client software

Each of these is discussed below.

MITEL IP PHONES

The Mitel family of IP Phones includes a broad range of feature-rich, IP-based handsets designed to fit easily into your existing data network using standards-based TCP/IP protocols.

When connected to a data network, the Mitel IP Phones connect to a Mitel server, they might automatically download and install necessary firmware, or an update / download might be initiated by the system administrator. The Mitel IP Phones are assigned to a ShoreGear voice appliance (commonly referred to as a 'voice switch') to receive dial tone and call control and then become ready for operation.

The Mitel IP Phones provide audio processing for calls but leverage the intelligence built into the ShoreGear voice switches for call setup and call routing.

SHOREGEAR VOICE APPLIANCES

Each ShoreGear voice appliance, or voice switch, is a solid-state hardware device used to manage and control the PBX functionality of the Mitel system. When a ShoreGear voice switch is connected to a data network and added to a Mitel system, it automatically downloads and installs necessary firmware by a simple reset / reboot of the ShoreGear switch. It then establishes connectivity to all other ShoreGear switches and Mitel servers in your enterprise.

Each ShoreGear switch receives, from its managing Mitel server, a subset of the entire enterprise-wide configuration database. This database includes information on all trunk groups, other Mitel devices, sites, dial plans, user permissions, call control parameters and call routing options.

Each ShoreGear can be the host, or 'owner', of local end points (such as end user phones), local services (such as, Hunt Groups and Bridged Call Appearances) and local trunks. The complete list of all local end points is exchanged with all other ShoreGear switches (including the SoftSwitch on all Mitel servers) so that every ShoreGear switch becomes aware of the location of all end points and resources throughout the entire enterprise.

Combining the knowledge of the configuration database (received from the server and stored in non-volatile RAM, or NVRAM) and the dynamic exchange of all end points and all resources in the system (received from all other switches/SoftSwitches and stored in dynamic RAM, or DRAM) allows each ShoreGear switch to make independent and intelligent call routing decisions without the need to query a centralized call processor or centralized servers.

ShoreGear voice switches provide complete call control and call routing services to their locally owned end points, and act as gateways to the PSTN network via locally connected trunks.

ShoreGear voice switches support a variety of stations (analog phones, 3rd party SIP phones, Mitel IP Phones) and trunk types (analog loop start, analog DID, SIP trunks, digital trunking such as T1, E1, BRI, PRI). ShoreGear switches also provide advanced call routing functions like Hunt Groups and Bridged Call Appearances. The resources built into each ShoreGear can be dynamically assigned and reconfigured by an administrator in real-time, providing configuration flexibility in response to rapidly changing requirements.

When ShoreGear switches are added to your data network, they connect and communicate directly with each other, creating a fully meshed network of connectivity with all other Mitel voice switches (ShoreGear and SoftSwitches). When Mitel IP Phones are connected to the data network, the phones are automatically assigned to the IP Phone resources of individual ShoreGear switches. Trunk resources are created, configured and assigned (physically or logically) to specific ShoreGear switches throughout the organization by an administrator.

Inbound calls are managed and controlled by the ShoreGear switch that is connected to the active trunk (analog, digital or SIP). Outbound calls are managed and controlled by the ShoreGear switch that is managing the phone (or device) that is placing the call. Each ShoreGear switch has locally-stored knowledge of all call routing details (received from the configuration database) and is fully aware of the location of all end points and trunk resources (from all other switches) and is therefore able to make all call routing decisions locally.

Call routing decisions are made by the local switch and call setup is performed directly between the necessary switches. Each ShoreGear switch maintains direct connectivity to all other resources (ShoreGear voice switches and SoftSwitches). No intermediate gateway or proxy server is required.

> **Note:** Locally stored call routing knowledge becomes more limited when your organization grows beyond 100 ShoreGear switches. For more information, see the details on enabling Distributed Routing Services (DRS) on page 134 below.

MITEL VIRTUALIZED APPLIANCES

In response to changing industry trends, Mitel 14.2 introduced a fully virtual Mitel unified communications system.

Mitel 14.2 allows Mitel customers to deploy virtual software, continue using physical appliances, or mix and match virtual and physical appliances – all within the same unified communications system, managed through the same single web-browser interface, Mitel Connect Director.

Virtualization complements Mitel's distributed architecture and allows customers to benefit from virtualization features and Mitel's inherent N+1 redundancy to deliver high availability at the lowest cost, with increased scalability to meet the needs of enterprises of all sizes.

MITEL VIRTUAL IP PHONE SWITCH

The Mitel Virtual IP Phone Switch can support up to 250, 500 or 1,000 phones – depending on the resources allocated to the virtual machine running the Virtual IP Phone switch software. It has built-in capacity for 5000 SIP Proxy ports, regardless of the number of phones it can support. For every 100 phones / users it supports: 5 Backup Auto Attendant prompts, 6 Make-Me Conference ports, 4 Hunt Groups, 16 total Hunt Group users, 16 users per Hunt Group, 8 Pick-Up Groups, 40 total Pick-Up Group users, 24 users per Pick-Up Group, 500 Extension Monitor buttons, Extension Monitor Event Rate is 1 per second and 2500 Busy Hour Call Completions (BHCC). For additional details please refer to Mitel's Planning and Installation Guide.

The Mitel Virtual IP Phone switch can be used as a spare.

Note: The Mitel Virtual IP Phone switch supports both MGCP and SIP phones, analog phones are not supported. The Mitel Virtual IP Phone switch does not support any trunks, Nightbell or Overhead Paging features.

MITEL VIRTUAL SIP TRUNK SWITCH

The Mitel Virtual SIP Trunk Switch can support 100, 200 or 500 SIP trunks – depending on the resources allocated to the virtual machine running the Virtual SIP Trunk Switch software.

All SIP trunks on a Virtual SIP Trunk Switch support media proxy to provide parity with PRI trunks. Media proxy enables a number of features on a SIP trunk: call recording, silent monitor, silent coach, barge, whisper page, and external extension assignment (Office Anywhere). It also has the benefit of providing a single IP address to the ITSP for all media and SIP signaling traffic which simplifies SIP trunk deployments.

Note: The Mitel Virtual SIP Trunk Switch only supports SIP trunks. It does not support phones or any of the telephony features, such as make-me conferencing.

MITEL VIRTUAL APPLICATION SERVICE APPLIANCE

The Mitel Virtual Service Appliance can support 50 or 200 audio/web conferences – depending on the allocated resources to the virtual machine running the Virtual Service Appliance (collaboration) software:

Mitel Connect Director will detect the resources allocated to the virtual machine and calculate the maximum capacity of the Virtual Service Appliance (collaboration).

Note: For additional details on Mitel's Virtual Switches / Appliances please refer to the Mitel Planning and Installation, Administration and Maintenance Guides.

MITEL APPLICATION SERVICES END POINTS

The Mitel UC system requires one headquarters server, initially called the HQ server. The HQ server is used for administration and reporting and holds the master copy of the configuration database. It can also be used for other Application Services such as Voice Mail, Auto Attendant menus, Workgroups, Group Paging, Account Code Collection, and the control of end-user client software such as the Mitel Communicator client.

Every Mitel installation requires an HQ server.

Additional Application Services end points can be added to distribute these services, including additional Mitel servers (called Distributed Voice Servers, DVS, or Distributed Application Servers, DAS, the term DVS is used throughout this document) and voice switches that have embedded Application Services, called V-switches.

The HQ server and DVS' can all be virtualized, please refer to the Mitel release specific Build Notes for supported platforms and versions.

These additional Application Services end points (Mitel servers or ShoreGear V-switches) can be added to a single site to increase Application capacity or to provide Application redundancy. Further, they can be added to remote or branch office sites to provide localized site survivability and reduce WAN traffic for access to Applications Services.

MITEL COMMUNICATOR

The Mitel Communicator family of clients is a unified framework of end-user clients that provide feature-rich and multi-media tools to help manage and control communication functions from any supported device. The collaboration features and graphical call control capabilities of the Communicator client are offered for Microsoft Windows and Apple Mac desktops, via web browsers, and across mobile platforms including iPhones, iPads, BlackBerrys, Androids and Nokia devices.

Mitel Communicator provides integrated, advanced UC functionality on the desktop or mobile device, providing phone management, integration with CRM and ERP systems, click- to-dial features, instant messaging, desktop video, telephony and desktop/UC presence, desktop sharing and web-based collaboration. The Mitel Communicator client provides total control over every aspect of the Mitel user's Unified Communications experience.

MITEL'S DISTRIBUTED ARCHITECTURE

Unique to Mitel is the distributed nature of its call control and configuration database. Each ShoreGear

voice switch is a multi-purpose appliance, that:

- Acts as a gateway to the PSTN
- Provides dial tone and call control to a wide array of handsets
- Is fully knowledgeable of all call routing, user permissions, configuration settings and enterprise-wide telephony resources
- Has direct connectivity to all other devices across the entire multi-site Mitel Unified Communications system

By combining this fully-meshed topology, the ShoreGear's self-reliance, and the distributed awareness of all resources and devices, each component of the Mitel system is more intelligent and more capable than any other voice solution on the market today.

The benefits of this unique distributed architecture include:

- **Faster call routing**: No centralized call processor needs to be queried
- **More accurate route selection**: Each ShoreGear switch knows which resources are reachable and will always select an available and reachable resource
- **Rapid reaction to outages**: Local decisions are made using local information and dynamically update switch connectivity in real-time
- **Reduced impact of outages**: Every call routing decision is made locally, therefore loss of connectivity or loss of another UC component has little or no impact
- **Site survivability**: Each site is natively independent and survivable, so the loss of centralized services has little or no impact
- **More reliable configuration**: The fault-tolerant features of the Mitel system are enabled with little to no programming – often with just a single checkbox. This simplified configuration results in fewer human errors and fewer misconfigurations, leading to higher reliability
- **Simplified scalability:** Expanding individual sites, or adding sites, becomes a simple exercise in selecting the right-sized ShoreGear switch and connecting it to your data infrastructure. It will automatically mesh with the existing Mitel components and resources

In the following sections we will provide an in-depth review of each of these components, the unique architectural differentiators of the Mitel UC system, and look at exactly how each component and feature works. We will provide detailed examples and identify the proper usage and configuration for each area of the Mitel system to achieve the highest degree of functionality, fault tolerance and survivability.

BUSINESS CONTINUITY DEFINITIONS

The term ‘Business Continuity’ is defined in many different ways by different organizations, but is commonly used as a broad term encompassing multiple categories of outage preparation, disaster planning, and the mechanisms to accommodate or alleviate the effects of those situations and events.

Business Continuity topics often include: Disaster Preparation, Disaster Recovery, Survivability, Resiliency, Redundancy, High-Availability, Fault Avoidance, Fault Tolerance, Backup and Restore Procedures, Failover and Fail-back.

For the purposes of this Application Note we will use the following terms and definitions:

Term	Description
Resilient	The ability to continue functioning unimpeded despite unfavorable or adverse circumstances. For example: • Having a wide range of temperature tolerances allows a device to continue operating even when the air conditioning fails in a server closet. • Having fewer moving parts and/or fewer internal connectors means there are fewer things to break, and parts are unable to wiggle free. This enables a device to continue operating after sustained movement or rattling such as during a minor earthquake or site relocation. **Note:** Resiliency is a measure of how tolerant a device is to error conditions. It is a measure of how long a device will continue running before a device outage (or failure) occurs.
Redundant	Having a mechanism to recover following a failure or outage. For example: • Having a secondary Ethernet port to react to changes in LAN connectivity due to a cable failure or port failure. • Having spare, or extra, resources or components that are activated to replace failed component(s). **Note:** Redundant systems can recover fast or slow. They can recover fully or partially. They might recover in certain circumstances but not in others. Recovery can be automatic or manual. Redundancy can be expensive or affordable. Redundancy is a measure of what happens after a failure occurs.

Disaster Recovery	The procedures that are implemented to recover from a major disruption in service. For example: • Having equipment and/or locations available to accommodate the needs of a displaced work force due to the inability to occupy a facility (e.g. a gas leak or snow storm prevents workers from entering a building). • Locating critical systems and components in a data center or colocation facility so that the physical loss of an administrative building (e.g. due to a natural disaster such as a fire or tornado) does not result in loss of data and/or functionality of the majority of your workforce. • Deploying systems in multiple geographical locations so that a localized or regional event, such as power failure, natural disaster, or physical attack, does not disrupt more than a small portion of your business operations and systems. • Having a set of policies and procedures in place to continue at an acceptable level of business activity in an acceptable alternate environment. **Note:** Disaster Recovery is a measure of the capabilities to continue business operations after a major disruptive event.

From a practical and functional perspective, businesses require a Unified Communications system that is as resilient as possible (to accommodate error conditions in the first place, and to gracefully recover from failure conditions) and as redundant as can be afforded (to accommodate unexpected outages). Resilient and redundant elements must be combined with strategic placement of equipment and circuits, along with advanced planning, to create proper procedures to invoke during a business-affecting disaster.

Some customers will desire 100% redundancy on all parts and all components of their Mitel system. This document will help them achieve that.

Other organizations will choose to prepare for the most probable outages, the most enterprise-wide affecting failure conditions, and the costliest incidents that might interrupt their enterprise.

By using a Business Continuity assessment process, an organization can identify and assess the costs, and impact, of each type of outage or disaster. They can further determine which events are most likely to occur, which outages create the greatest impact, in terms of lost business or productivity, and which events are most affordable to prepare for and protect against.

An organization can then weigh the costs and benefits of preparation vs. the enterprise- disruption and potential loss of business revenue during each type of situation, and apply resources as necessary.

DESIGN FEATURES, BENEFITS AND BEST PRACTICES

The Mitel Unified Communications system is based on a unique, distributed architecture. This architecture is inherently tolerant to failures, outages and disaster scenarios as will be outlined below. Yet, proper planning, design and configuration must be performed whether you are deploying a single site or in a multi-site environment.

This section will describe many of the resiliency-oriented features and native redundancy capabilities of the Mitel system and its core components. It will explain the reasons for each design, describe how each feature works, and provide examples to illustrate typical use-case scenarios for the features.

Understanding these core architectural capabilities lays the foundation for designing a highly available, fault-tolerant and redundant Mitel system.

SHOREGEAR VOICE APPLIANCES

Hardware appliances are inherently lower maintenance and more reliable than servers. All core PBX functionality on a Mitel system is performed by the hardware-based ShoreGear voice switches or the Mitel Virtual IP Phone / Trunk switches.

Functions performed by each ShoreGear switch include PBX functionality such as call set up, call tear down, class of service and permission decisions, trunk selection, call routing, and least cost routing.

The ShoreGear hardware-based appliances are purpose-built devices using a highly reliable, embedded operating system. They are solid-state devices having no hard-drives with spinning platters, are built without daughter cards and have all components soldered directly onto the motherboard, including RAM, FLASH memory, Digital Signal Processing resources (DSPs) and CPU processors. These units are easier to install (no cards or modules to insert), easier to replace (no pre-configuration is necessary) and are less susceptible to environmental issues (such as temperature variations and vibrations). Combined, these design elements result in an industry-leading Mean-Time Between Failures (MTBF) of 12-19 years for each unit.

ShoreGear switches are self-leveling (upgrading or downgrading to match the version of code running on the other components of the Mitel system), self-monitoring (providing real-time feedback and alerts to administrators on elements such as internal temperature and fan speed) and provide built-in redundancy capabilities (e.g. dual, redundant physical- layer Ethernet ports and an automatic Power Fail Transfer relay for emergency dial tone during power outages).

The Mitel Virtual IP Phone Switch and Virtual Trunk Switch can leverage VMware High Availability (HA) capabilities to provide additional reliability for a server-based VMware guest.

In total, these design elements make each individual voice appliance in a Mitel system more resilient and more redundant than components of other manufacturer’s UC solutions; especially when compared to server-based systems.

Combining the expected defect rates (based on individual components) and actual defect rates (based on returns) along with the average time to replace a defective unit, ShoreGear appliances each exceed 99.999% (‘Five 9’s’) of reliability and uptime. Each unit that Mitel manufactures is built to provide fault-free service well into its second decade of deployment.

ShoreGear Model	Predicted MTBF (hours)	Demonstrated MTBF (hours)	Demonstrated MTBF (years)	Availability (1 hour MTTR)
120/24	84,500	530,000	60.5	99.9998%

60/12	91,000	528,000	60.2	99.9998%
40/8	132,300	520,000	59.3	99.9998%
T1 & E1	154,200	601,000	68.5	99.9998%
30 & 50	190,600	252,800	28.8	99.9996%
90	171,400	215,000	24.5	99.9995%
50V	175,800	N/A	N/A	99.9994%
90V	159,400	N/A	N/A	99.9994%
30BRI & 90BRI	172,600	N/A	N/A	99.9994%
90BRI-V	162,900	N/A	N/A	99.9994%
220T1A	163,500	N/A	N/A	99.9994%
220T1 & 220E1	189,300	N/A	N/A	99.9995%

1. N/A = Not enough data (to date). MTBF = Mean Time Between Failures.

2. MTTR = Mean Time to Repair.

Source: ‘Building Reliable IP Telephony Systems: How Architecture and Design Differentiate Mitel from the Competition’

By Ed Basart, Chief Technology Officer, Mitel

FULL MESH CONNECTIVITY AND DISTRIBUTED INTELLIGENCE

Two architectural design aspects of the ShoreGear appliances should be highlighted.

First, each voice switch maintains a full mesh of connectivity with all other voice switches and Mitel servers (SoftSwitches) in your organization. This means that each ShoreGear switch maintains a dynamic table of ’reachable’ destinations, listing all the telephony end points and Application Services that are distributed throughout the enterprise.

Second, each voice switch receives and maintains a subset of the entire enterprise-wide configuration database including the complete set of knowledge required to make all call routing decisions. Every user group setting, class of service, user permission, call path and dial plan detail is stored internally, and written to NVRAM, by each ShoreGear switch. As changes occur throughout the organization, updates are incrementally shared and distributed to each ShoreGear switch in near real-time.

These two unique design aspects are of *huge*architectural significance when considering the fault-tolerant nature of a UC system. Being fully aware of all configuration and dial plan settings; being fully knowledgeable of all end points and all system-wide resources; and being able to communicate directly to every other destination, means that every call routing decision is made instantly and locally by the ShoreGear switch that ‘owns’ the call.

Note: These architecturally designed aspects change when your organization grows beyond 100 ShoreGear switches. For more information, see the details on enabling Distributed Routing Services (DRS) on page 134 below.

LOCATION SERVICE PROTOCOL (LSP)

ShoreGear switches share their switch-specific end points with other switches in the system using the Mitel proprietary Location Service Protocol (LSP). Switches keep current by propagating their changes and receiving updates from other switches.

The data sent and received for each *extension end point* includes three pieces of information:

1. The dialable extension number of the end point
2. A ‘forwarding destination’ number (or extension)
3. The IP address of the switch that owns the end point

The data sent and received for each *trunk group* includes three pieces of information:

1. The Trunk Group ID
2. The Individual Trunk’s port number (in that trunk group)
3. The IP address of the switch that owns the trunk

For example, when you add a new user within Mitel Connect Director, the web-based administration interface, and assign them to a phone serviced by ShoreGear switch ‘A’, the HQ server (owner of the master copy of the configuration database) communicates the new addition to ShoreGear ‘A’. ShoreGear ‘A’ then uses the Location Service Protocol (LSP) to inform all other switches of the new extension it now owns.

If a new user (or trunk) is added to ShoreGear switch ‘B’, which is managed by a Distributed Voice Server (DVS), the addition of the new user is communicated from the HQ server to the DVS server, and from the DVS server to ShoreGear switch ‘B’. ShoreGear ‘B’ then uses LSP packets to inform all other switches of the new resource that has been assigned to it.

Each ShoreGear switch learns of its own assigned resources (users, phones, trunks) from its managing Mitel server. This information, being part of the local copy of the configuration database received from the server, is stored locally in non-volatile FLASH memory. If a switch is rebooted, this stored information is re-loaded from local FLASH memory and is updated, as needed, once connectivity to its managing Mitel server is restored.

Each ShoreGear switch learns of the resources owned by other switches by exchanging LSP packets with all other switches (includingphysical ShoreGear switches, Virtual switches and the SoftSwitch running on each Mitel server). The resource information (users, phones, trunks) discovered from other switches is stored in DRAM. If a switch is rebooted, it loses this dynamic content and must re-contact and re-learn this information from all other switches.

When a ShoreGear switch needs to contact another user, or resource, it looks in its internal LSP table to determine the IP address of the resource’s managing switch. The ‘Source switch’ contacts the ‘Destination switch’ to initiate the call. If the destination extension (user, Workgroup, Hunt Group, and so on.) has re-routed the call to another destination (such as ‘Forward to the Voice Mail extension’), the Destination switch will instruct the Source switch to redirect to the alternative end point/extension/number.

EXTENSION DIALING

Consider the following *extension dialing* example:

- A user in the Seattle site dials the extension of a user at the Boston site. The Source ShoreGear switch in Seattle knows which switch owns the destination extension (by looking in its local LSP table), and initiates communication with the Destination ShoreGear switch in Boston. The Boston switch will accept the call, and ring the Boston user's phone.
- If the Boston user has forwarded all calls to Voice Mail (that is, a 'Do Not Disturb' configured Call Handling Mode), the Boston switch will instruct the Seattle switch to contact the Mitel Voice Mail extension. The Seattle switch will reach out to its nearest Voice Mail end point asking for the Voice Mail box of the user in Boston. If the Boston user's voice mail box is assigned to a different DVS or V-switch, the call will be redirected to the appropriate Voice Mail end point.

  Note: See 'Selection of End Points for Application Services' on page 26 below for more details on the selection of the 'nearest' Voice Mail server.

The LSP table in the Seattle ShoreGear switch has an entry for the Boston user which includes:

- The Boston user's extension
- The IP address of the Boston ShoreGear switch managing that end point
- The Voice Mail extension as the current 'forwarding destination' for the Boston user

  Note:

- The 'forwarding destination' stays the same whether the Boston user is in 'Standard' Call Handling Mode ('5 rings, no answer, forward to Voice Mail') or 'In a Meeting' Call Handling Mode ('Forward always to Voice Mail'). The ultimate forwarding destination for the user is still the Voice Mail extension.
- If the Boston user were to change their Call Handling Mode setting to 'Forward always to the Receptionist at x12345', then updated LSP information will be propagated from the Boston switch to all other voice switches in the Mitel system. The 'forwarding destination' for the Boston user will be changed to x12345.

When a call is placed to an extension, the Source ShoreGear switch reaches out to the Destination ShoreGear switch to establish the call. Based on the settings of the destination end point, the Destination switch might redirect the call to another location.

If the Source switch cannot reach the Destination switch (e.g. due to a WAN failure) the Source switch will route the call to the previously learned 'forwarding destination' for that extension. If that forwarding destination is also unreachable, the call will be directed to the Source switch's internal Backup Auto Attendant (see page 31).

Some end points have explicitly assigned 'Backup Extensions'. These include Work Groups and Hunt Groups. Other end points, like Users, have forwarding destinations that change based on user settings and Call Handling Modes. The LSP table maintained by each ShoreGear switch lists the 'currently assigned' forwarding destination (for users) or Backup Extension (from Workgroups, Hunt Groups, and so on.) for each entity, and will use that number (or extension) if the original destination's managing switch is unreachable.

It is important to note that forwarding destinations are not iterative. Consider the following:

- User 'A' lists Workgroup 'B' as their Call Handling Mode forwarding destination
- Workgroup 'B' lists Hunt Group 'C' as its explicit Backup Destination

- If a call is routed to User 'A', and the ShoreGear switch managing User 'A' is unreachable, the Source ShoreGear switch will route the call to Workgroup 'B'. If the server managing Workgroup 'B' is *also* unreachable the call will be directed to the internal Backup Auto Attendant of the Source switch.

- The call will *not* be routed to Hunt Group 'C'

Note: For additional details see the Explicit Backup Extensions section on page 46 below.

TRUNK DIALING

Consider the following *trunk dialing* example:

- A user in the Seattle site dials a Boston PSTN number. The local Source ShoreGear switch in Seattle knows that there is a PRI trunk and several analog trunks in Boston that can be used to place the call as a local, non-toll call. The Source switch in Seattle reaches out to the Destination switch in Boston that owns the PRI circuit to initiate the call. The Destination ShoreGear switch accepts the connection, and places the call out an available channel on the PRI.
- If the PRI circuit had been out of service, or completely in use, the Destination switch in Boston would have refused the connection. The Source switch in Seattle would then initiate a new connection to the Boston switch owning the analog trunks. If those trunks were also busy, or out of service, the Seattle switch would place the call as a long distance call from a local trunk based in Seattle.

When a call is placed to an external number, the Source ShoreGear switch communicates with the Destination switch that owns the 'best' trunk to use for placing the call. If those trunks are busy, the Source switch will continue to reach out to Destination switches, in order of trunk group preference, until the call is placed.

Note: See 'Network Call Routing' on page 99 below for additional details on the selection and prioritization of trunksgroups

LSP SUMMARY

Call routing decisions are made locally by the ShoreGear voice appliance that 'owns' the call. No external database lookup or server-based query is required. All call routing is handled by the local ShoreGear switch because it has knowledge of the relevant components of the enterprise-wide database and has direct communication with all other devices throughout the distributed Mitel architecture.

This distributed intelligence makes the Mitel Unified Communications system inherently more reliable than other UC systems. An outage of one component, or section, of the Mitel system does not interfere with the proper operation of the other components in the system. There is no centralized call processing server, or master control unit: each ShoreGear switch is independently fully aware and fully capable, and will continue to work and interoperate with all other Mitel components that are up, running and reachable.

Note: Locally stored call routing knowledge becomes more limited when your organization grows beyond 100 ShoreGear switches. For more information, see the details on enabling Distributed Routing Services (DRS) on page 134 below.

Note: See Appendix D: LSP Tables for more information on LSP table contents.

SELECTION OF PRIMARY AND SECONDARY END POINTS FOR APPLICATION SERVICES

Another task that each ShoreGear switch performs is to select the two ‘nearest’ Application Services end points. These end points can be V-switches or Mitel servers (HQ or DVS), and are used whenever a generic Application Service (like Voice Mail, an Auto Attendant menu, the Account Code Collection service, and so on.) is needed.

Most entities in the Mitel system have a single, specific, dialable extension and are managed by one specific Mitel device. These entities include Users, Workgroups, Hunt Groups, and Route Points. By exchanging LSP packets with all other switches (both ShoreGear switches and the ‘SoftSwitch’ running on every Mitel server) each ShoreGear appliance learns the whereabouts of every extension (entity) in the system.

Some services, like Auto Attendant menus and the Account Code Collection service, are fully distributed and run on every Application Services end point including V-switches, the HQ server and all DVS servers.

Other services, like the Voice Mail service, are fully distributed to all Application Service end points but each end point only has control over specific parts of the whole system. For Voice Mail, every user is assigned by the administrator to a single Voice Mail storage end point in Mitel Connect Director. Voice Mail end points manage their locally assigned users and work together with all other Voice Mail end points to provide backup Voice Mail services for each other.

Every Voice Mail instance (running on the HQ server and all V-switches and all DVS servers) knows which Voice Mail end point every user is assigned to. Standard ShoreGear switches do not. A standard ShoreGear switch routes all calls destined for an Auto Attendant menu, the Voice Mail system or the Account Code Collection service to its primary Application Services end point. If that end point is unreachable (and has therefore been marked as ‘inactive’ in the ShoreGear’s local LSP table) the call is routed to the ShoreGear switch’s secondary Application Services end point. If that end point is also unreachable (and marked as ‘inactive’ in the local LSP table) the call is routed to the local ShoreGear switch’s Backup Auto Attendant (see page 31).

SELECTING THE PRIMARY AND SECONDARY APPLICATION SERVICES END POINTS

When a ShoreGear switch connects to the Mitel system it receives (from the Telephony Management Service, or TMS, running on its managing server) an ordered list of all Mitel servers and V-switches that are ‘at or above’ the newly connected switch’s site.

This information is derived from the configuration of the Site Hierarchy Tree as displayed in Mitel Connect Director.

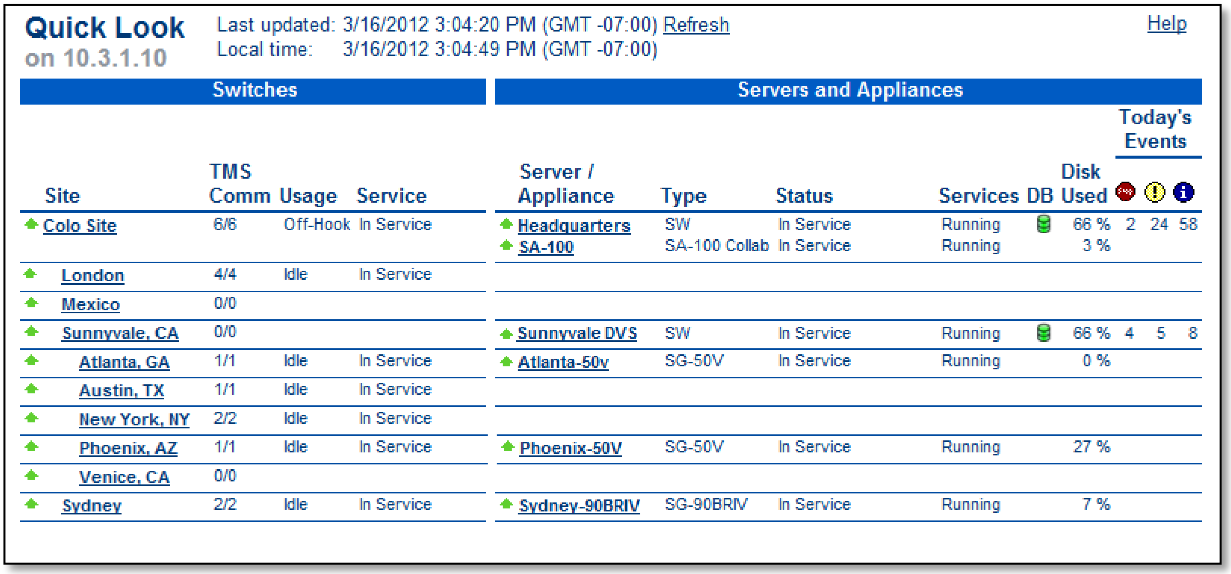

Quick Look on 10.3.1.10

Last updated: 3/16/2012 3:04:20 PM (GMT -07:00) Refresh
Local time: 3/16/2012 3:04:49 PM (GMT -07:00)

Help

Switches				Servers and Appliances								
										Today's Events		
Site	TMS Comm	Usage	Service	Server / Appliance	Type	Status	Services	DB	Disk Used	🛑	⚠	ℹ
▲ Colo Site	6/6	Off-Hook	In Service	▲ Headquarters ▲ SA-100	SW SA-100 Collab	In Service In Service	Running Running	⛁	66 % 3 %	2	24	58
▲ London	4/4	Idle	In Service									
▲ Mexico	0/0											
▲ Sunnyvale, CA	0/0			▲ Sunnyvale DVS	SW	In Service	Running	⛁	66 %	4	5	8
▲ Atlanta, GA	1/1	Idle	In Service	▲ Atlanta-50v	SG-50V	In Service	Running		0 %			
▲ Austin, TX	1/1	Idle	In Service									
▲ New York, NY	2/2	Idle	In Service									
▲ Phoenix, AZ	1/1	Idle	In Service	▲ Phoenix-50V	SG-50V	In Service	Running		27 %			
▲ Venice, CA	0/0											
▲ Sydney	2/2	Idle	In Service	▲ Sydney-90BRIV	SG-90BRIV	In Service	Running		7 %			

The Site Tree Hierarchy in Mitel Director

Each ShoreGear switch uses this ordered list to select a primary (or 'closest') Application Services end point, and a secondary (or 'next closest')

Application Services end point. Closest, or nearest, is an indication of relative proximity based on the Mitel Site Hierarchical Tree. These two Application Services end points can be V-switches, Distributed Voice servers (DVS) or the HQ server.

The following selection rules are followed by each ShoreGear switch when selecting its primary and secondary Application Services end points:

1. Local site resources are preferred over parent site resources

IMPORTANT

Physical location and geographical distance between sites is immaterial to the Mitel system.

Only the parent/child relationship of the Mitel Tree Hierarchy is used for selecting Application Services End Points.

- Parent site resources are preferred over grandparent site resources, and so on.

1. Servers are preferred over V-switches

2. At least one of the two Application Services end points *must* be a server

Since there will always be an HQ server at the top of the tree, every switch will always have at least one Application Services end point: the HQ server. As your topology grows and you add additional Mitel servers and/or V-switches, each ShoreGear switch will re-select its two 'closest' Application Services end points.

SELECTION EXAMPLES

Given the tree hierarchy displayed above, the ShoreGear switches at each site will select the following as their primary and secondary Application Services end point:

Site	Primary	Secondary
Colo Site	Headquarters	<none>
London	Headquarters	<none>
Mexico	Headquarters	<none>
Sunnyvale, CA	Sunnyvale DVS	Headquarters
Atlanta, GA	Atlanta-50v	Sunnyvale
Austin, TX	Sunnyvale DVS	Headquarters
New York, NY	Sunnyvale DVS	Headquarters
Phoenix, AZ	Phoenix-50v	Sunnyvale
Venice, CA	Sunnyvale DVS	Headquarters
Sydney	Sydney-90BRIv	Headquarters

Switches will always select the two 'closest' end points. Only *two* will be selected (unless only one is found in the tree hierarchy). If the topology changes, the selection process is repeated.

For example, if a new DVS is added to the Colo Site, the selection process is re-initiated by all ShoreGear switches at and below the Colo Site. The ShoreGear switches at the Colo Site, as well as in London and Mexico, would be able to select the new DVS to compliment the original selection of the HQ server, populating their list with two end points instead of their previous single end point.

If there are more Application Services end points than needed, each with the same relative 'nearness', then each ShoreGear switch will randomly select from the list. This means, through the law of averages, that of the many ShoreGear switches at the Colo Site, London and Mexico, half will pick the HQ server as their primary and the DVS as their secondary.

And the other half will pick the DVS as their primary and the HQ server as their secondary.

Similarly, ShoreGear switches at sites such as Sunnyvale, Atlanta, Phoenix, and Sydney will always select the local V-switch or DVS server as their 'closest' end point, but will randomly pick between the HQ server and the DVS server at the Colo Site as their secondary end point selection.

> **Note:** If there is both a Mitel server and a V-switch of equal 'nearness' the server will be selected before the V-switch.

V-switches can never be selected as *both* the primary and the secondary end points. At least one end point will *always* be a Mitel server; even if that means the HQ server (at the top/root of the tree hierarchy) is selected as the secondary end point for a ShoreGear switch.

> **Note:** Servers are preferred because they have more storage capacity and more simultaneous ports of access than V-switches.

DVS servers will always select themselves as their primary Application Services end point. They will choose the next nearest server or V-switch as their secondary end point.

V-switches will always select themselves as their primary Application Services end point and will always select a server (HQ or a DVS) as their secondary end point.

> **Note:** Deploying two V-switches at a remote site will not cause them to use each other as backup. Each V-switch will choose itself as its primary Application Services end point and an 'up-the-tree' server as its secondary Application Services end point.

Only the two 'closest' Application Services end points are stored. If a ShoreGear cannot reach (and therefore has marked as 'inactive') its primary Application Services end point, the ShoreGear switch will route the call to its secondary Application Services end point. If that end point is also unreachable, the call will be routed to the internal Backup Auto Attendant (BAA) within the ShoreGear switch itself (see page 31).

For example, consider the following:

- An inbound call is received on a Seattle PRI trunk and is transferred to a user in Boston. The Source ShoreGear-T1k switch in Seattle knows what the destination user's controlling switch is in Boston and will communicate directly with the Destination switch in Boston. The Destination ShoreGear switch will ring the Boston user's phone. After 5 rings, if the call is not answered, the Destination switch instructs the Source switch to send the caller to the Mitel Voice Mail system extension.
- The ShoreGear-T1k in Seattle looks in its LSP table, finds that the Voice mail system extension is mapped to its primary Application Services end point, a V-switch in Seattle and sends the call there. The V-switch knows that the Boston user is assigned to a Voice Mail server in Boston and instructs the ShoreGear-T1k switch to send the call to the Boston Voice Mail server extension. The Source switch communicates with the Boston server's SoftSwitch (owner of the Boston servers' Voice Mail extension), the Boston server's SoftSwitch answers the call and plays the Boston user's greeting and the caller then leaves their message.
- If the Boston server had been unreachable, the local Mitel V-switch would have accepted the call and retained the voice mail message itself, and queued the message for later delivery once the Boston server connection was restored.

As a best practice it is recommended to add a DVS server to the root site. This ensures that all switches, including those at the root site have two Application Services end points to use for Voice Mail and Auto Attendant services.

Further benefits and uses of these Application Services end points will be illustrated throughout the remainder of this document.

BACKUP AUTO ATTENDANT

Each ShoreGear voice appliance intelligently routes calls to the most appropriate destination using its internal knowledge of the enterprise's configuration and resources. As described previously, this is possible because each switch maintains a dynamic list of all resources and their locations (it knows where the call should go), stores a copy of the relevant configuration database elements (it knows what permissions and restrictions apply) and it can communicate directly with every other component in the Mitel system (it knows how to get there).

If a call needs to be routed to an end point that is assigned to a single switch, such as a ShoreGear-based end point (e.g. an end user's phone, a Bridged Call Appearance, a Hunt Group) or a server-based Application endpoint (e.g. a Workgroup, a Group Paging extension) the ShoreGear switch that 'owns' the call already knows if it can, or cannot, reach that particular device by checking in its local table of reachable Mitel components (the LSP table).

If the intended destination is *not* reachable, the ShoreGear switch will route the call to the most appropriate alternative destination, which might be an explicit Backup Extension, a learned 'forwarding destination', or the primary (or secondary) Application Services end point for that switch.

Note: See the 'Explicit Backup Extensions' section on page 46 below for complete details on explicit backup extensions.

For Application Services such as the Account Code Collection service, Auto Attendant Menus or Voice Mail, the call is routed directly to the switch's primary Application Services end point.

Note: See the 'Selection of End Points for Application Services' section on page 26 above for complete details on the selection of primary and secondary Application Services end points.

In the unusual case when a call cannot be properly serviced by the originally intended destination, *and* the appropriate alternative destination(s) also cannot be reached, each ShoreGear switch has a built-in 'Auto Attendant of last resort'. This final 'safety net' is called the Backup Auto Attendant (BAA).

The Backup Auto Attendant is a set of pre-recorded prompts, stored in non-volatile FLASH memory, that present a set of courteous and appropriate (and always reachable) options to the caller. The Backup Auto Attendant prompts are used as a last resort, offering local and reachable options.

Some typical Backup Auto Attendant prompts are:

- "That extension is not valid."
- "I'm sorry, that extension cannot be reached at this time. If you want to reach someone else, and know their extension, you might dial it now. For assistance press zero."

The Backup Auto Attendant prompts are Mitel-created, non-editable, pre-recorded prompts and are stored in FLASH memory on each switch. This means that even if a ShoreGear switch becomes entirely isolated (WAN failure, Ethernet failure, and so on.) the switch is still able to play appropriate prompts to callers. Some (or all) destinations might be unreachable but the BAA 'safety net' offers a final 'best effort' at presenting suitable and functional options in a courteous manner to the caller.

Additionally, each site can be configured with a site-specific backup operator. This backup site operator is presented as an option to callers by the Backup Auto Attendant.

If ever an IP Phone displays 'B Auto Attendant' you have reached a Backup Auto Attendant prompt being played from the local ShoreGear switch.

PEER-TO-PEER AUDIO PATH

The Mitel UC system is designed to scale without compromising features or capabilities. To achieve this, all voice streams are established as peer-to-peer communication flows. In other words, there is no central server, 'proxy' device, or voice gateway which an audio stream must communicate with or relay through. All voice packets are sent directly from the source device to the destination device.

In most cases, each new voice stream (for example, an audio call) will both originate and terminate on:

- An IP phone (for example, Mitel IP Phone, SIP Phone, SoftPhone, and so on.)
- A ShoreGear voice switch (trunk, voice switch or conference appliance; for example, analog phone, conference bridge port)
- A Mitel server (for example, Voice Mail, Auto Attendant Menu, Contact Center IVR, and so on.).

Call setup (coordinating the end points, negotiating codec selection, and so on.) is conducted by the ShoreGear switch. However, the actual voice packets (the Real Time Protocol, or RTP, media stream) flows from the IP address of the source device directly to the IP address of the destination device.

For example, in the following diagram you will see a direct, peer-to-peer exchange of voice packets between:

1. The Headquarters' ShoreGear 220-T1A and an IP Phone (for example, an inbound PSTN call to a user's direct inward dial, or DID/DDI, number)
2. The Remote Office's ShoreGear 220-T1A and an IP Phone (for example, an outbound PSTN call placed by the Mitel user)
3. An IP Phone and another IP Phone (for example, a site-to-site, extension-to-extension call)
4. An IP Phone and a Mitel server (for example, a Voice Mail or Auto Attendant call)
5. The SIP Session Border Controller (InGate) and an IP Phone (for example, a SIP trunk-based call)

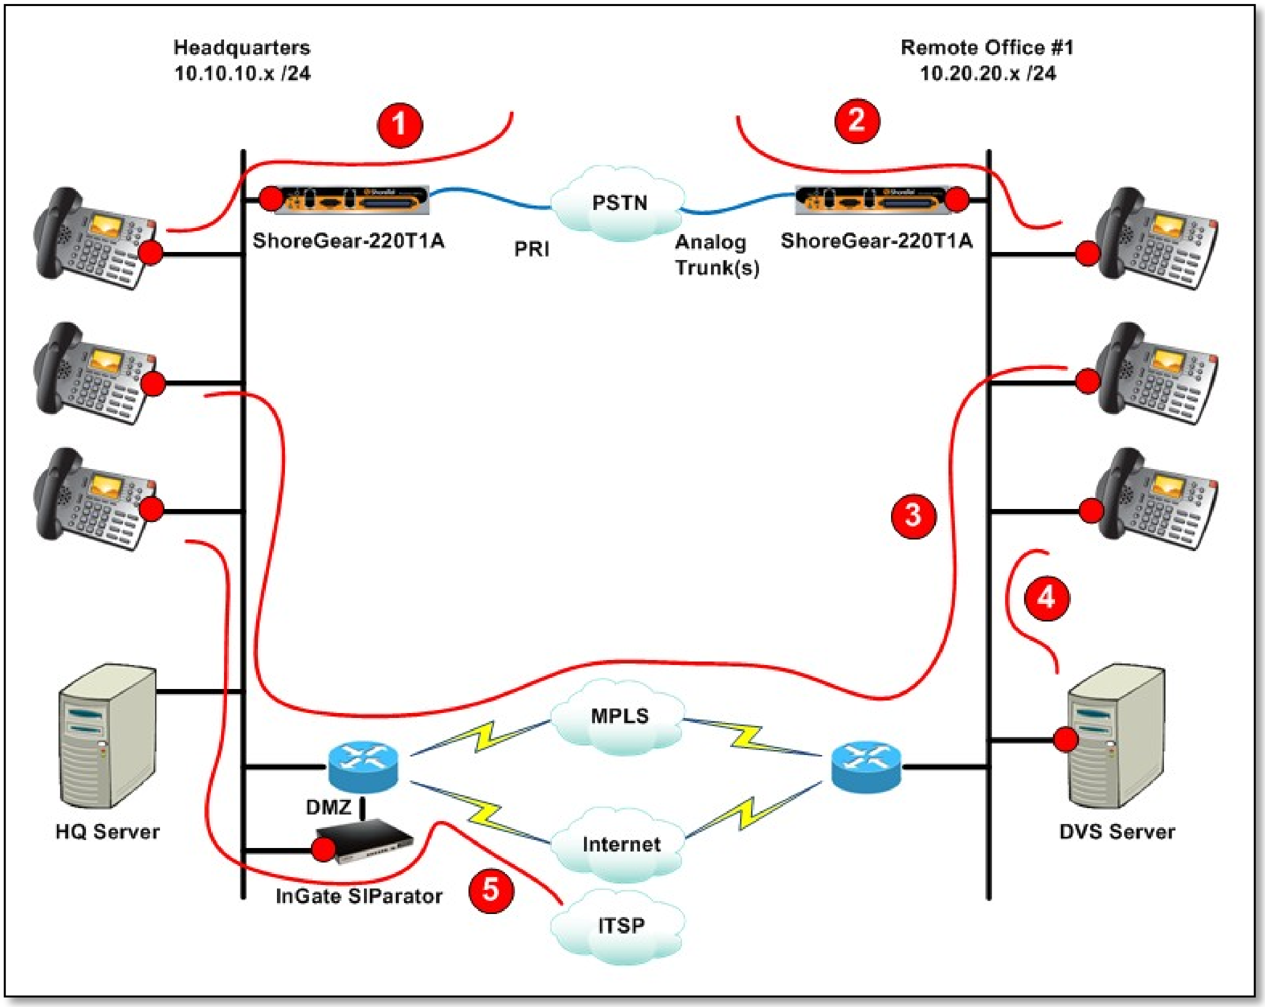

Peer-to-peer audio paths.

In each of these packet flows, there are two unidirectional streams of RTP packets between the IP addresses of the two end points (represented by the red 'dots' on each device) involved with each call.

This peer-to-peer audio path architecture is important to understand, especially as you begin to consider call flows and behaviors when components fail or become unreachable.

Once a call is set up and in progress, only the two RTP end points and the intermediary network devices (such as IP routers, Ethernet switches, and so on.) are involved. A failure in any other component of the Mitel UC system (ShoreGear switch, Mitel server, and so on.) will not affect or interfere with the active call.

Note: A failure with the ShoreGear trunk switch (PRI or analog) could result in a dropped call.

DESIGN FEATURES, BENEFITS AND BEST PRACTICES: SUMMARY

The resiliency and fault tolerance-oriented features of the Mitel hardware components include:

1. Reliable, virtual or hardware-based call control using ShoreGear voice appliances.
2. Fully-meshed communication between all ShoreGear switches and Mitel servers.
3. Distributed Intelligence, whereby all ShoreGear switches store a copy of the relevant portions of the configuration database, enabling them to independently make full call routing decisions.
4. A dynamic exchange of destination information shared between all Mitel switches.
5. The selection of two Application Services end points used when an Application service is needed, with automatic failover between them.
6. An embedded Backup Auto Attendant acting as a final safety net when all other call route options have been exhausted.
7. Peer-to-peer audio path that avoids the need for any intermediary device and stays functional even during ShoreGear switch outages.

SINGLE SITE REDUNDANCY

SINGLE SITE – DISTRIBUTED CALL CONTROL

While each individual ShoreGear voice appliance has been engineered to provide trouble free service for many years, it doesn't mean that problems won't arise. Problems can be caused by external connection issues, failures in other devices, circuit outages or accidents.

We will now look at some of these fault conditions and discuss some best practices for single site design and deployment. We will use an example site of 100 users.

With a single site of 100 IP Phones you could easily deploy a single ShoreGear voice switch to provide all call control and PSTN connectivity. Plus we add a single Mitel server for system administration and application services.

This creates a fully functional and feature rich UC environment with a highly reliable, solid- state device at the core of your PBX functions. However, several 'single points of failure' still remain.

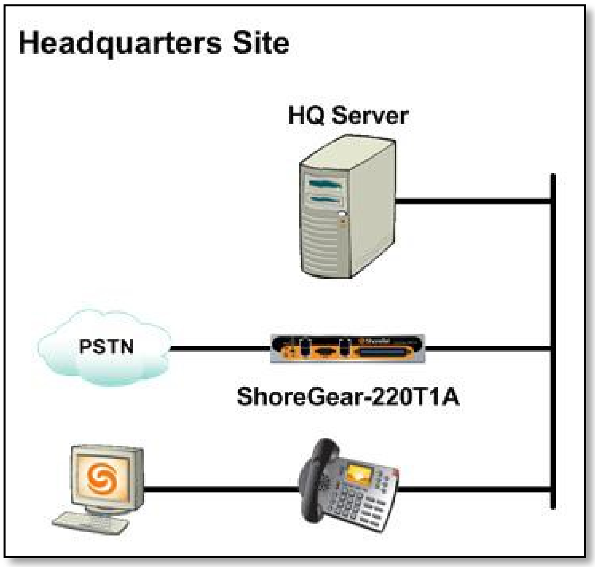

Simple single site for 100 users

Specifically, there are three areas of concern:

1. If there is a problem with the single trunk circuit, you have lost *all* ability to take and place external calls. (Consider a wiring worker who mistakenly disconnects your circuit while working in a Telco closet.)
2. If there is a problem with the single ShoreGear appliance, you have lost call control for *all* 100 users. (Consider a janitor whose vacuum trips a circuit breaker.)
3. If there is a problem with the single Ethernet cable (or with the port on your Ethernet switch), you have lost call control for *all* 100 users. (Consider an Ethernet switch that stops forwarding traffic due to an unexpected software lockup.)

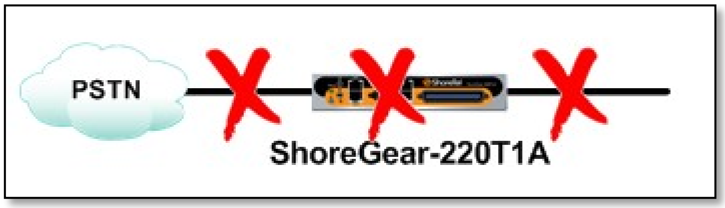

Three separate ‘single points of failure’

To address these issues, it is recommended that you consider the following design changes:

1. Split your trunking into multiple circuits (where feasible) and distribute the separate trunks amongst multiple ShoreGear switches.
2. Instead of deploying one large ShoreGear switch, deploy a few medium sized or multiple smaller-sized ShoreGear switches.
3. Each ShoreGear switch has two auto-selecting LAN ports. Connect one port to your primary Ethernet switch and connect the other port to either a different Ethernet switch, a different blade (in a chassis-based Ethernet system), or even to the same Ethernet switch.

By distributing your trunks and changing from one trunk circuit to multiple trunk circuits, you lessen the impact of a single cable cut or circuit disconnect. The loss of one trunk only reduces your call capacity rather than stopping all call capacity. (For even better survivability, ensure that separate trunks enter your facility using different physical entry points and cable runs.)

By distributing your call control (changing from one ShoreGear switch to several) you lessen the impact of losing a single ShoreGear switch. The loss of one switch disables some phones in your organization but not all phones.

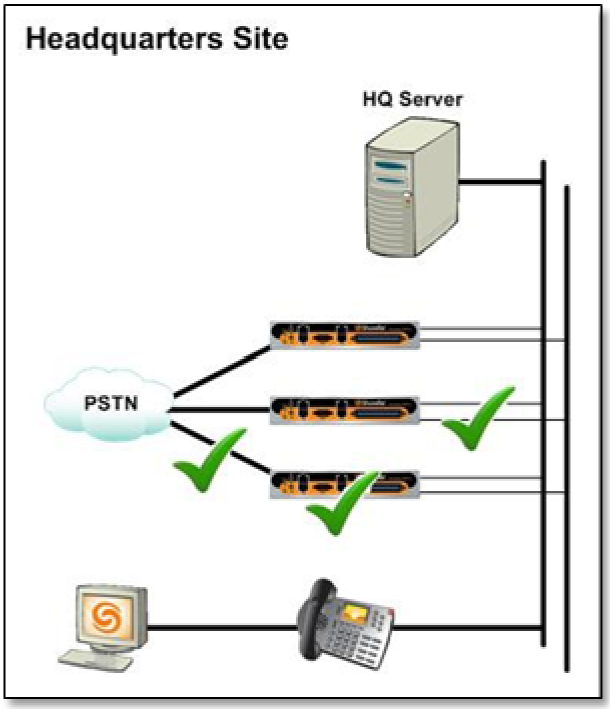

Single site with no single point of failure

By connecting the redundant LAN ports on each voice switch, you have avoided an outage due to the sudden failure of an Ethernet switch, port or cable. The ShoreGear switch will immediately change from one Ethernet port to the other.

All three single points of failure have now been addressed. Any outage or fault has been changed from a total loss of functionality to merely a reduction in capacity or reduced service.

Yet this reduction in service, especially the loss of a ShoreGear switch, still creates a very serious service outage. To remedy this, we take advantage of the automatic load balancing features and built-in IP Phone failover capabilities of the Mitel system.

SINGLE SITE – REDUNDANT CALL CONTROL

In the section above we started with very resilient, hardware-based PBX components to avoid faults in the first place. We then removed any single point of failure through the deployment of multiple intelligent devices instead of one single device. But now, instead of lessening the impact of a failure, we want the system to automatically work around the failure. We want the system to be redundant.

To illustrate call control redundancy, let’s change our example to a single site with 150 IP Phones. For the sake of simplicity, we will ignore trunking for the moment.

To supply resources for 150 IP Phones we could use a single ShoreGear-220 appliance. But based on the best practices described above, we have wisely chosen to deploy three ShoreGear-50 appliances instead. This removes any single point of failure in our design.

With all three ShoreGear-50 appliances up and running, we then proceed to plug in our 150 IP Phones. The Mitel system will accept those phones and evenly distribute them to the three ShoreGear switches, assigning 50 phones to each switch.

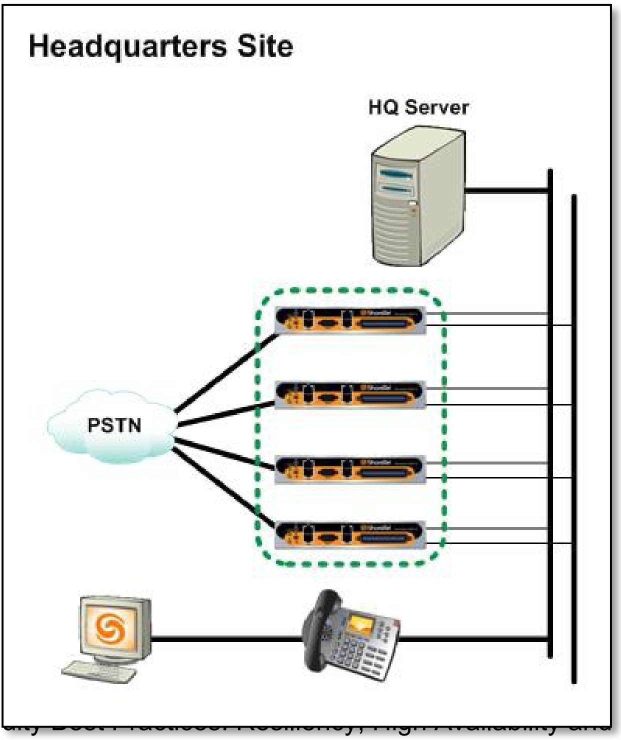

Single site with fully redundant call

But the loss of one of those switches means a severe disruption: 50 IP Phones will completely lose service. To remedy this, we need to add only one more ShoreGear appliance of equal or greater capacity to the site. This will provide redundancy for all devices at the entire site. We need three ShoreGear- 50's; and we add one more ShoreGear-50: we use the term N+1 to describe the fact that no matter how many voice appliances you need (N) you only have to add one more (+1) to achieve complete redundancy for all devices at that site.

Now, with four ShoreGear switches up and running instead of just three, we plug in our 150 IP Phones. The Mitel system will accept those phones and load balance the phones across all four active switches, assigning 37 (or 38) phones to each voice switch.

If any one ShoreGear appliance becomes unavailable for any reason, the 37+ phones serviced by the failed switch will be automatically re-assigned to the extra capacity available on the three remaining voice switches.

Note that the added switch (commonly, but incorrectly, described as the '+1 switch') is not held in reserve. It is actively used and actively load balanced along with all other active switches at that site, more commonly referred to as N + 1 Load Balancing. It is the fact that there is spare (extra) capacity spread across all the remaining appliances at that site that allows the remaining switches to accommodate the reassignment of the phones that are 'orphaned' by a switch going off-line. The addition of the 4th switch brings the total number of IP Phone resources at the site to 200. If any of the four ShoreGear switches becomes unavailable, there are still 150 remaining IP Phone resources; enough to supply dial tone and call control to all 150 phones at that site.

Any switch can fail – there is extra capacity for those 'orphaned' phones to be reassigned to available IP Phone resources on the other active switches.

The failover process takes anywhere from a few seconds to a maximum of about 4-6 minutes. Active calls on an affected phone will continue until the call is ended by the user.

Enabling this IP Phone failover capability is done with a single checkbox. In Mitel Connect Director, navigate to '**Administration** > **IP Phones** > **Options**' and select the **Enable IP Phone Failover** option.

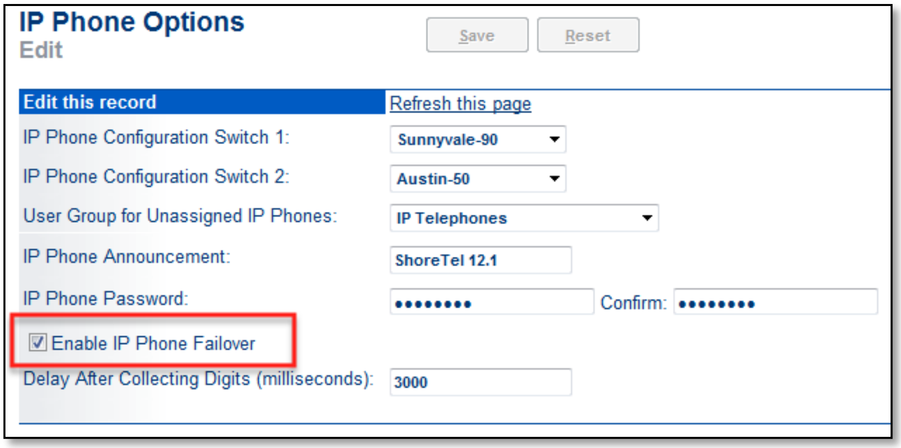

IP Phone Options
Edit

Edit this record

IP Phone Configuration Switch 1: Sunnyvale-90

IP Phone Configuration Switch 2: Austin-50

User Group for Unassigned IP Phones: IP Telephones

IP Phone Announcement: ShoreTel 12.1

IP Phone Password: •••••••• Confirm: ••••••••

☑ Enable IP Phone Failover

Delay After Collecting Digits (milliseconds): 3000

Enabling IP Phone Failover features for all sites with one checkbox

From a practical perspective, moving from one large voice switch to several small switches is of nominal, if any, price difference. And the addition of one more switch to provide complete redundancy is, in this example, an approximate 33% uplift in price. This is far more affordable than other manufacturers that force you to duplicate *all* components (a 100% uplift in price) to achieve the same results.

By making redundancy simpler and more affordable, the Mitel system enables even small organizations to deploy a fully redundant UC system

ANALOG PHONE AND SIP PHONE FAILOVER

Be aware that analog phones (modems, fax machines, lobby phones, and so on.) are physically connected to a port on a single ShoreGear switch. Failure of that switch will disable all directly connected analog devices, as well as all directly connected trunks circuits.

For improved survivability, it is advised to spread your analog devices amongst multiple ShoreGear switches. Then, if one switch fails, it only affects *some* of your analog devices instead of *all* of them.

SIP phones, such as the Mitel IP-8000 and other 3rd party SIP phones, use a similar registration process as Mitel IP Phones. Whereas Mitel IP Phones use IP Phone Configuration Switches to be assigned to specific ShoreGear IP Phone resources, SIP phones use SIP Proxy resources to be assigned to ShoreGear IP Phone resources.

Please note that the Mitel IP 400 series phones utilize the SIP protocol, but do not utilize SIP Proxy resources and instead use the IP Phone Configuration Switches.

For more information and additional detail on IP Phone Configuration Switches, SIP Proxy switches and SIP devices, see '*Application Note 10298: Mitel IP Phone Failover Features'*

POWER FAIL TRANSFER PORTS

Like any critical business tool, your phone system should be connected to quality Uninterruptable Power Supplies (UPS) to avoid disruption in service due to short power outages. A UPS system provides line conditioning and might be sized to supply backup power anywhere from a few minutes, to several hours or more.

However, during an extended power outage, UPS systems will drain completely and the Mitel components will inevitably be shut off.

To accommodate this situation, each ShoreGear switch that supports analog connections (analog phones and analog trunks) has a built-in power fail transfer relay that enables 'pass-through dial tone' in case of a complete power outage. By examining the cabling diagram of the RJ-21X Telco connector on each ShoreGear switch, you will notice that two ports, one analog trunk port and one analog station port, are identified as the designated Power Fail Transfer (PFT) ports.

It is recommended that you connect a PSTN Analog Loop Start trunk to the PFT trunk port, and connect a standard analog phone to the PFT station port. During normal operation, when the ShoreGear switch has power, the analog phone is simply a standard extension on the Mitel system and the analog trunk is simply a standard trunk controlled by the Mitel system.

During normal operation, lifting the handset of the analog phone will provide internal dial tone generated by the ShoreGear switch. Dialing a Trunk Access Code (such as '9') followed by a PSTN number will cause the ShoreGear switch to collect your dialed digits, evaluate your permissions and then select the most suitable trunk in the enterprise to place the call.

If the ShoreGear switch loses power, the power fail relay will close, physically connecting the wiring from the PFT trunk port to the PFT station port. All other functions of the ShoreGear switch will stop. All other physically connected analog phones will stop functioning. All other physically connected analog trunks will stop functioning.

The only capability still functional is the one analog phone connected to the PFT port. Lifting the handset of that one phone will provide direct dial tone from the telephone company's Central Office (CO) analog trunk.

Note: Since the dial tone is coming directly from the telephone company's Central Office, no Mitel Trunk Access Code is required. Dialing an emergency number such as '911' should be dialed directly, with no preceding Trunk Access Code.

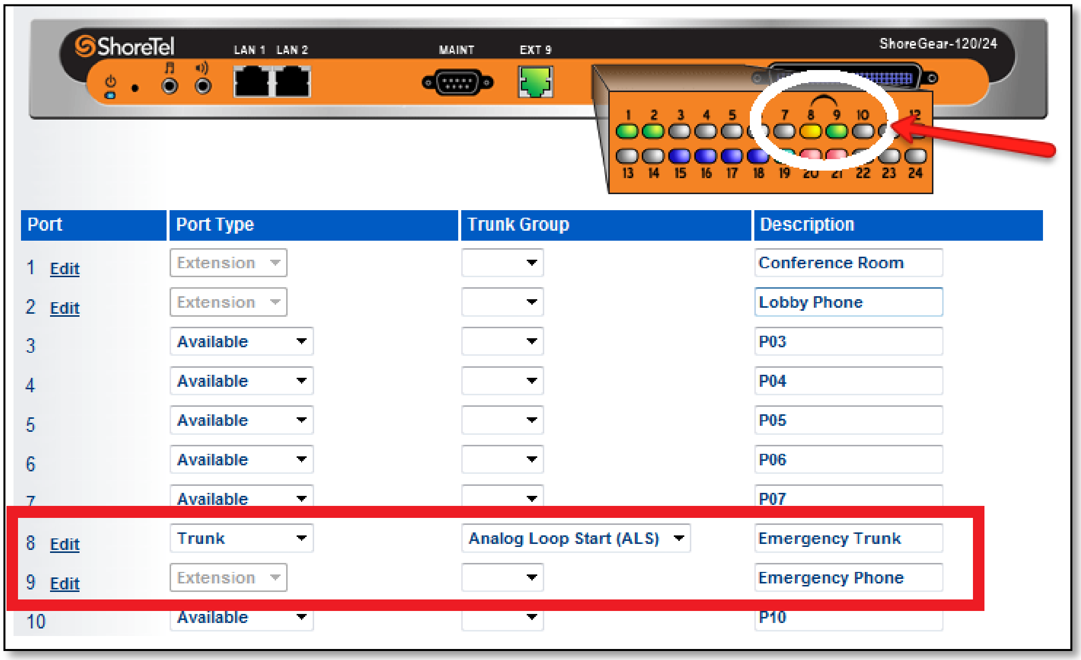

Port	Port Type	Trunk Group	Description
1 Edit	Extension		Conference Room
2 Edit	Extension		Lobby Phone
3	Available		P03
4	Available		P04
5	Available		P05
6	Available		P06
7	Available		P07
8 Edit	Trunk	Analog Loop Start (ALS)	Emergency Trunk
9 Edit	Extension		Emergency Phone
10	Available		P10

Notice the 'umbrella' mark over ports 8 and 9 on a ShoreGear 120/24.

Port 8 is the PFT Trunk Port. Port 9 is the PFT StationPort

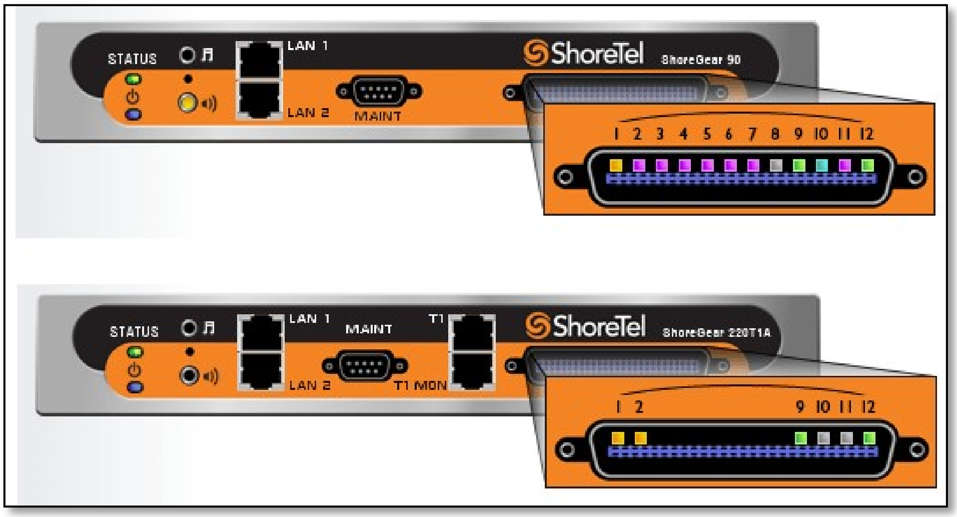

The ShoreGear 90 and ShoreGear 220T1A both use ports 1 and 12 for Power Fail Transfer

You can have only one trunk and only one analog phone per ShoreGear switch assigned as a Power Fail Transfer ‘pair’. It is recommended that at least one (preferably two) analog trunk(s) and analog phone(s) be configured at every site for emergency calling purposes during an extended power outage.

If you have a site with only digital trunks (such as a T1/E1 PRI) or only SIP trunks, it is strongly recommended that you add at least one (preferably two) analog trunk(s) and analog phone(s) for emergency calling purposes during an extended power failure.

SINGLE SITE REDUNDANCY: SUMMARY

Best practices for designing a single site for high availability with no single point of failure include:

1. Distributing call control among multiple ShoreGear switches. This reduces the impact of any device failure.
2. Distributing trunk circuits to multiple ShoreGear switches. This reduces the impact of any single cable cut or circuit failure.
3. Connecting each ShoreGear switch using both the primary and the secondary LAN ports to your Ethernet infrastructure – preferably to different Ethernet switches or blades. This promotes survivability in case of a cable, port or Ethernet switch failure.

APPLICATION SERVICES

The Mitel architecture separates call control from Application Services. The hardware- based ShoreGear appliances perform all of the call control while Application Services are provided by the HQ server, by Distributed Voice Servers (DVS) and by V-switches.

Note: See Appendix A: HQ, DVS and V-switch Features for a complete discussion of all Application Services performed by HQ servers, Distributed Application Servers and V-switches.

Now, let's return to our single site example.

So far, we have made the call control *resilient* (by starting with well-designed, virtual or hardware-based appliances), *fault-tolerant* (by distributing trunks, distributing call control and connecting both LAN connections on all ShoreGear switches) and *redundant* (by adding extra capacity to accommodate the failure of any one appliance, aka 'N+1').

But the Application Services are still running on a single point of failure: the Mitel HQ server.

Servers can go offline for a variety of reasons: hardware faults, a software crash, or for scheduled or unscheduled maintenance activity (such as an off-line backup, hardware upgrade or an operating system patch). In fact, the Mitel architecture presumes that server-based services will be off-line with some degree of regularity and is designed to make those occurrences less-impacting or even non-impacting.

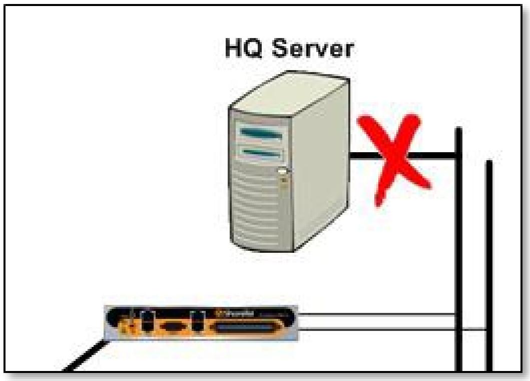

Loss of server connectivity disrupts many important services.

Remember, since all core PBX functions are performed by the ShoreGear voice appliances, the loss of a server, even an outage of the primary HQ server has no impact on the taking and placing of calls. All inbound and outbound calls continue to function normally. All permissions and classes of service are still enforced. All call setup, call teardown, call routing, trunk selection and least cost routing is still fully functional. Hunt Groups continue to function. Extension-to-Extension calls work, PSTN calls work and intercom calls work.

Why? Because all of these services are performed entirely by ShoreGear appliances. No dependence on the Mitel server(s) is required.

Yet, even though the core PBX features remain fully functional, other important capabilities would be disabled during the server outage including: Administration, Reporting, Voice Mail, Workgroups and Communicator client services. None of these lost services prevent the phones from ringing, or stop a user from dialing from their phone – but many users rely so heavily on these UC features that they have essentially become mission critical features in their own right.

The Mitel system offers many options to improve the availability and redundancy of these Application Services as will be discussed below.

EXPLICIT BACKUP EXTENSIONS

Some Application Services, like Workgroups, can be configured with an explicit ‘Backup Extension.’ This backup extension is used when the Workgroup’s managing server is unreachable. It is recommended to create a Hunt Group that has similar membership, ring patterns and overflow destinations as each Workgroup. Hunt Groups are serviced by a virtual or hardware-based ShoreGear appliance and remain functional even if the server is off-line.

Alternatively, you can leverage the Distributed Workgroup features of Mitel 10 and above and create a second Workgroup that is assigned to run on a different server and select that Workgroup as the explicit Backup extension of the first Workgroup. In this fashion one Workgroup server can be backed up by another Workgroup server.

Note: Distributed Workgroups (page 50) and Distributed Databases (page 82) are mutually exclusive.

By selecting a Hunt Group as the ‘Backup Extension’ of a Workgroup, any calls that cannot be routed to the Workgroup (e.g. during a server outage or WAN outage) will instead be handled by the Hunt Group.

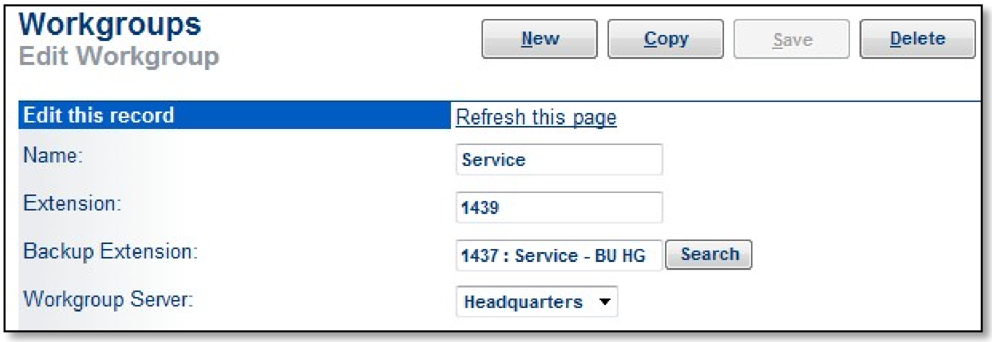

Name: Service

Extension: 1439

Backup Extension: 1437 : Service - BU HG

Workgroup Server: Headquarters

The ‘Service’ Workgroup is assigned to run on the server called ‘Headquarters’ and uses the Hunt Group ‘Service – BU HG’ as its explicit Backup Extension

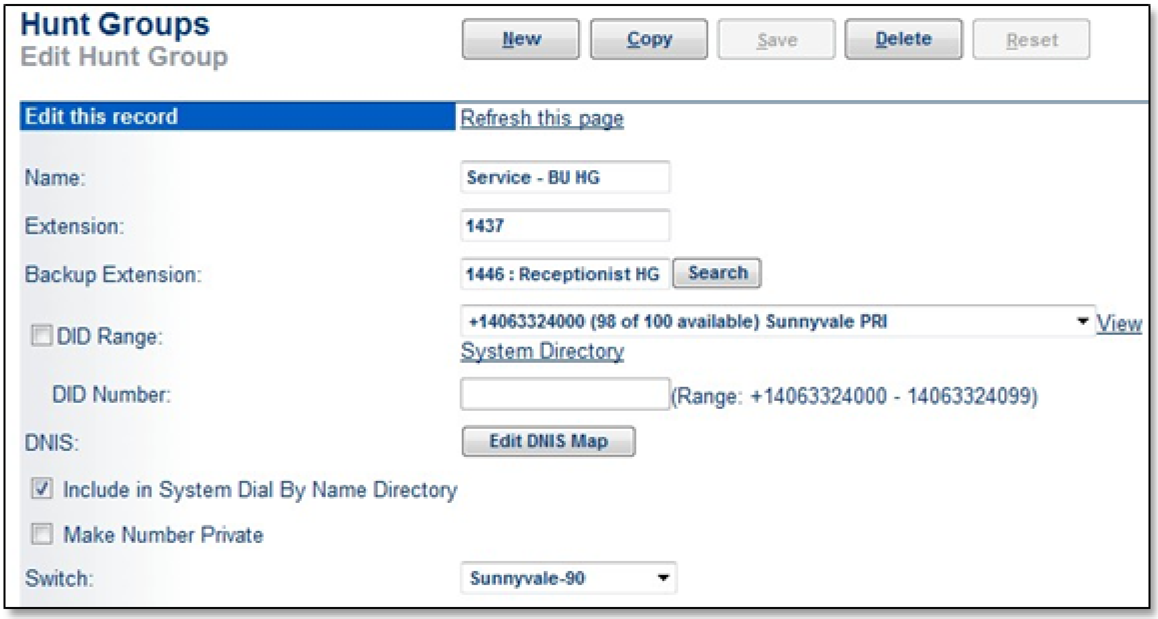

Hunt Groups
Edit Hunt Group

Edit this record | Refresh this page

Name: Service - BU HG

Extension: 1437

Backup Extension: 1446 : Receptionist HG

□ DID Range: +14063324000 (98 of 100 available) Sunnyvale PRI

DID Number: (Range: +14063324000 - 14063324099)

DNIS:

☑ Include in System Dial By Name Directory

□ Make Number Private

Switch: Sunnyvale-90

The Hunt Group named ‘Service – BU HG’

The Hunt Group ‘Service – BU HG’ was created and programmed as a backup for the ‘Service’ Workgroup. This Hunt Group runs on the ‘Sunnyvale-90’ ShoreGear switch and can route calls in a similar way to the original ‘Service’ Workgroup until the HQ server is once again reachable.

A Workgroup can use nearly any dialable extension as its explicit Backup Extension, including other Workgroups, Hunt Groups, Auto Attendants or a user’s extension.

Similarly, Hunt Groups can use almost any dialable extension as their Backup Extension, including other Hunt Groups, Workgroups, and so on.

Note:

- It is important to choose a backup destination that is running on a *different* entity than the primary entity. In our example, we would not want to back up the ‘Service’ Workgroup with another Workgroup running on the same ‘Headquarters’ server. If the Headquarters server was unreachable, *both* Workgroups would be unreachable and the backup extension would be of no benefit.
- Similarly, we would not want to assign the ‘Service – BU HG’ Hunt Group to run on the same ShoreGear switch that the ‘Receptionist HG’ Hunt Group is running on. If that ShoreGear appliance is unreachable, both Hunt Groups would be unreachable and the call would be routed to the Backup Auto Attendant.

The Mitel entities that have explicit Backup Extensions include:

- Bridged Call Appearances
- Hunt Groups
- Workgroups

You should consider your specific call flows and site and trunk configurations when selecting effective backup extensions. For instance, if calls enter on trunks located at a remote site and are serviced by a Workgroup on a server at the HQ site, selecting a ShoreGear switch at the HQ site to host the Workgroup’s backup Hunt Group would work during a server failure but not during a WAN outage as both the primary destination (the Workgroup) and the Backup Extension (the Hunt Group) are inaccessible due to the WAN failure.

Selecting a remote-site ShoreGear switch to host the backup Hunt Group is a better alternative. If the WAN is down the HQ-site calls still route to the Workgroup and remote- site calls route to the Hunt Group. If the server is down all calls route to the remote-site Hunt Group.

Explicit Backup Extensions are only used ‘one layer deep’. If the original destination is unavailable, the ShoreGear switch will route the call to the explicit Backup Extension of the unavailable entity. If the backup destination is *also* unreachable, the call will be sent to the local switch’s Backup Auto Attendant.

Multiple iterations of backup extensions are not followed.

For example, consider the ‘Service’ Workgroup above. The Workgroup (running on the Headquarters server) lists the Hunt Group ‘Service – BU HG’ as its Backup Extension. The Hunt Group (running on the Sunnyvale-90 ShoreGear switch) lists the Hunt Group ‘Receptionist HG’ as its Backup Extension.

If a ShoreGear switch needs to route a call to the ‘Service’ Workgroup but the server that manages that Workgroup is unreachable, the switch will route the call to the Workgroup’s Backup Extension: the ‘Service – BU HG’. If the ShoreGear switch that manages that Hunt Group is *also* unreachable the call will be routed to the Backup Auto Attendant of the Source ShoreGear switch. The call will *not* be routed to the ‘Receptionist HG’.

Doing so would require that a ShoreGear switch follow ‘the trail of backup extensions’ recursively. They do not. A ShoreGear switch will only follow one Backup Extension and then route the call to its internal Backup Auto Attendant.

REPLICATED SERVICES: AUTO ATTENDANTS AND ACCOUNT CODE COLLECTION

Some services, like Auto Attendant Menus and the Account Code Collection service, are fully replicated on all Application Services end points including the HQ server, all DVS servers and all V-switches.

When you save an Auto Attendant menu in Mitel Connect Director, all recorded prompts, schedule assignments and call flow patterns of the AA Menu will immediately be replicated, in full, to all servers and all V-switches.

This means that any call that is directed to an Auto Attendant can be successfully serviced by the ‘nearest’ server or V-switch. Any server or V-switch within the network might be used; and the ‘closest’ will always be chosen.

Services that are fully redundant across all Mitel servers and V-switches include:

- Auto Attendant (AA) Menus
- Account Code Collection (ACC) Service

If a call is intended to be routed to an Auto Attendant Menu and neither the primary nor the secondary Application Services end points are available, the call will be sent to the ShoreGear switch’s internal Backup Auto Attendant.

bIf a call is intended to be routed to the Account Code Collection service and neither the primary nor the secondary Application Services end points are available, the call will be forwarded according to the following rules:

- If the user’s User Group has Account Code Collection as ‘optional’, the call will be placed.
- If the user’s User Group has Account Code Collection as ‘forced’, the call will be denied.

Note: See ‘Selection of End Points for Application Services’ on page 26 above for more details on the selection of the primary and secondary Application Services end points.

DISTRIBUTED SERVICES: WORKGROUPS

Some services are distributed to multiple Application Services end points in the Mitel environment but are *not* fully replicated. This means that the service will be fully functional when the primary end point is available, but are limited, or non-functional, if the primary end point is unreachable. Workgroups is a distributed, but not replicated, service.

When you create a new Workgroup in Mitel Connect Director, you define which Mitel server (HQ or a particular DVS) that should ‘own’ and manage the Workgroup. If a call is directed to a Workgroup, the controlling ShoreGear switch that owns the call will check its LSP table. If the destination server that owns the Workgroup is reachable, the call will be forwarded to the SoftSwitch running on the server. If the destination server that owns the Workgroup is *not* reachable, the ShoreGear switch will direct the call to the Workgroup’s explicit Backup Extension.

Workgroups are required to have an explicit Backup Extension. If the owner of the backup extension (server, V-switch or ShoreGear switch (virtual or hardware-based)) is not reachable, the call will be directed to the local ShoreGear switch’s Backup Auto Attendant.

Note:

- Only the HQ server and DVS servers support Workgroups; V-switches do not.
- Distributed Workgroups (page 50) and Distributed Databases (page 82) are mutually exclusive. Once a Workgroup has been assigned to a DVS server you cannot enable a distributed database on any DVS servers. Conversely, if a distributed database has been assigned to any DVS server you cannot assign a Workgroup to any DVS server.

DISTRIBUTED SERVICES: VOICE MAIL

Voice Mail is another service that is distributed but not replicated. The distributed Voice Mail service differs from the distributed Workgroup service in how it fails over during an outage.

The distributed Voice Mail service uses a ‘Store and Forward’ failover mechanism and is described below.

The HQ server, all DVS servers and all V-switches run an instance of the distributed Voice Mail service.

Each user (or other Mitel entity) with a voice mail box is assigned to *one* specific Voice Mail storage end point. The selected Voice Mail server or end point will be the sole storage location for that user, and holds the user’s:

- Recorded name
- Recorded greetings (up to five, one for each of the user’s five Call Handling Modes)
- All messages left by inbound callers

Every Mitel server and all V-switches know the assigned Voice Mail storage end point for each and every user at all sites. ShoreGear voice switches do not.

If a user, controlled by one switch, places a call to another user, controlled by another switch, the Source switch will contact the Destination switch to connect the call. If the destination user has forwarded their calls to another destination (e.g. 'Forward always to Voice Mail' or 'Forward after 3 rings to an external PSTN number') the Destination switch will instruct the Source switch to redirect the call to the new destination.

If the redirection is to the Mitel Voice Mail system extension, the Source switch will contact its primary (or secondary) Application Services end point. That Application Services end point, being fully knowledgeable of every user's assigned voice mail storage location, will instruct the Source switch to redirect the call to the end point that has been assigned as the destination user Voice Mail end point. If the destination's assigned voice mail end point is unreachable, the local Application Services end point will answer the call itself, play an appropriate alternative prompt, and record and retain the voice mail message for later delivery once connection to the original destination end point is restored.

Note:

- See the section 'Selection of End Points for Application Services' on page 26 for more details.
- See Appendix A: HQ, DVS and V-switch Features for additional information on services performed by each type of Application Services end point.

If both the primary and secondary Application Services end points are not available, the Source ShoreGear switch will send the caller to the switch's internal Backup Auto Attendant.

Note: See Appendix B: Voice Mail Prompt Behavior for details on the exact behavior of the Voice Mail Prompts that are played when an alternative Application Services end point answer a call.

VOICE MAIL FAILOVER: STORE AND FORWARD

As stated earlier, the Mitel Voice Mail service is a distributed service, but is not replicated. If a caller is able to reach a user's assigned voice mail end point (HQ server, DVS server or V-switch) the caller will hear an appropriate prompt and will leave a message. The completed message will immediately be placed into the destination user's voice mail box.

The destination user will see the message in their Mitel Communicator client (and optionally, in their e-mail inbox) and the Message Waiting Indicator (MWI) on their phone will begin blinking.

If the destination user's assigned voice mail end point is *not* reachable, the call will be answered by the local ShoreGear's primary (or secondary) Application Services end point. The caller will hear an appropriate prompt and will leave a message. The message is then kept, and stored, on the alternative end point. It has not yet reached the original destination's voice mail box – it is still waiting to be delivered.

The answering Application Services end point, whether a Mitel server or V-switch, relies on its fully-meshed table of connectivity to all other servers and switches to be notified when the original destination user's assigned voice mail end point is once again reachable. At that point, the answering Voice Mail end point transmits (forwards) the stored message to the assigned Voice Mail end point using the SMTP protocol. The message will be deposited in the destination user's voice mail box and the destination user will see the message in their

Communicator client, optionally in their e-mail inbox, and the Message Waiting Indicator (MWI) on their phone will begin blinking.

This delivery mechanism is analogous to the Postal Service: If the mail box at your house is unreachable because the roads are too icy and impassable (akin to a WAN circuit being down), your postal mail will not be delivered that day. The mail is not lost – it is safely stored in the mail truck waiting for an opportunity to be delivered. Once the path is clear (akin to the WAN circuit being restored), the postal truck can once again deliver your mail. All the stored messages will flood into your mail box in a burst.

Server-based Voice Mail features are only available on Mitel servers which have ‘Allow Voice Mailboxes’ checked in Mitel Connect Director.

☑ Allow Voice Mailboxes

Unchecking ‘Allow Voice Mailboxes’ on any server will prevent that server from being able to host any user’s voice mail and prevents the server from acting as a ‘Store and Forward’ backup for other voice mail servers. This will also prevent the server from being selected as either a primary or secondary Application Services end point for any ShoreGear switch.

Note: It is possible to uncheck ‘**Allow Voice Mailboxes**’ on the HQ server. However, it is recommended to leave this box checked, even if you choose not to assign any user’s voice mail to the HQ server, so that the HQ server can be used as an Application Services end point for ShoreGear switches. That is unless you are instructed to disable this feature by third party vendors that are integrating with the Mitel system.

SERVICES THAT FAILOVER: CALL DETAIL RECORD (CDR) COLLECTION

Similar to the Mitel Voice Mail service, the Mitel Call Detail Record (CDR) collection service is fully distributed, but not replicated.

Every IP Phone is managed by a ShoreGear switch. Every ShoreGear switch is managed by a Mitel server. All ShoreGear switches send completed call detail information to their managing servers and in turn, all Mitel servers collect and forward the CDR information to the HQ server for consolidation and the running of reports.

A standard ShoreGear switch transmits call statistics to its managing Mitel server at the conclusion of each call (actually, at the end of each call stage, or leg). If a ShoreGear switch cannot communicate with its managing server at the end of a call, the CDR information for that call is discarded and lost.

A ShoreGear V-switch also transmits call statistics to its managing Mitel server at the conclusion of each call. If a V-switch cannot communicate with its managing server at the end of a call, the CDR information for that call is discarded and lost.

An exception to this is when the call is answered by an Application Service on the V-switch itself (that is, Auto Attendant menus, the Account Code Collection service, or Voice Mail). In that case, the V-switch will cache the CDR information for the locally answered calls, and forward the data to the HQ server (just like a DVS server would do). If the HQ server is unreachable, the CDR information will be temporarily held in DRAM on the V-switch. The V-switch will hold up to 120 minutes of stored CDR data (by default). When connectivity to the

HQ server is restored, the most recent two hours of CDR information will be pushed to the HQ server.

If the loss of connectivity extends beyond 120 minutes, the temporarily stored CDR data will begin to 'time out'. Any CDR data older than 2 hours will be discarded. Only the most recent two hours of CDR data will be cached and pushed to the HQ server when connectivity to the HQ server is restored.

Note: For V-switches, only locally-answered Application Services calls (such as Auto Attendant Menus, Voice Mail services and Account Code Collection services) are cached and forwarded to the HQ server. All other call activity (for managed IP Phones, and so on.) is immediately pushed to the V-switch's managing server, just like a standard ShoreGear switch.

Mitel Distributed Voice Mail servers (DVS) collect all CDR data from the ShoreGear switches it manages and calls to its own Application Services (Auto Attendant menu, Workgroup, Route Points, Account Code Collection service, and so on.) and forwards that CDR data to the HQ server. If the DVS cannot communicate to the HQ server it will temporarily store the CDR information in DRAM.

Like the V-switch, the DVS will retain up to 120 minutes (2 hours) of stored CDR data by default.

If connectivity to the HQ server is restored before that time, all CDR information will be successfully pushed to the HQ server. If the loss of connectivity extends beyond 120 minutes, the temporarily stored data will begin to age out. Any CDR data older than 2 hours will be discarded. Only the most recent two hours of CDR data will be cached and pushed to the HQ server when connectivity to the HQ server is restored.

Note: If required, DVS and V-switch retention of cached CDR information can be increased to longer than 2 hours using a server registry key setting.

SERVICES THAT DO NOT FAILOVER

As stated above, some services, such as Auto Attendant Menus and Account Code Collection, are fully replicated. For these fully replicated services, any Application Services end point (server or V-switch) can be used.

Other services are distributed, but not replicated. Some (such as Workgroups) use a Backup Extension mechanism for failover. Others (Voice Mail and CDR collection) implement a 'Store and Forward' technique.

Finally, some Application Services are completely dependent on one particular Application Services end point. If that end point is unavailable, the service stops functioning for the users assigned to that end point.

Application Services that *do not* failover include:

- Paging Groups
- Communicator Client Services
- Administration
- Running Reports

PAGING GROUPS

When creating a Paging Group you define the Mitel server (HQ or a particular DVS) that will 'own' and manage the Paging Group. Paging Groups *do not* have an explicit Backup Extension. If the destination server that owns the Paging Group is not reachable, the ShoreGear switch will direct the call to the local ShoreGear switch's Backup Auto Attendant.

If group paging is a mission critical feature, consider creating a second Paging Group assigned to a different server using the same Extension List. If the first Paging Group extension becomes unusable due to a server failure, inform your employees to use the second Paging Group.

COMMUNICATOR CLIENT SERVICES

The Communicator client provides many valuable Unified Communication features to the Mitel end user, including access to:

- Visual Voice Mail
- Call History
- Quick Dialer
- Graphical Call Control
- Directory Access
- Settings and Options

The Mitel Communicator client runs on Microsoft Windows, the Apple Mac, and in supported Web browsers.

COMMUNICATOR FOR WINDOWS

There is a direct correlation between the Mitel user, their assigned voice mail storage location and the ShoreGear switch and Mitel server that is driving the call control on the user's phone.

COMMUNICATOR FOR WINDOWS IN MITEL 8 – 12

In Mitel through version 12, the Communicator for Windows client uses a protocol called CSIS (Client Server Internet Service) to communicate with the Mitel server that hosts the user's voice mail box. CSIS is used for all voice mail activity and for changing settings such as within Options/Preferences and selecting Call Handling Modes.

A Mitel administrator uses Mitel Connect Director to assign a user to a specific Mitel server or V-switch for storing (and controlling) their voice mail box.

Additionally, the Communicator for Windows client uses a protocol called TAPI (the Telephony API) to communicate with the Mitel server that manages the ShoreGear switch that controls the user's phone (IP Phone, analog phone or SoftPhone).

Additionally, the Communicator for Windows client uses a protocol called TAPI (the Telephony API) to communicate with the Mitel server that manages the ShoreGear switch that controls the user's phone (IP Phone, analog phone or SoftPhone).

More specifically:

- A user is assigned to an IP Phone

- The IP Phone is assigned to a ShoreGear switch (as dictated by an IP Phone Address Map)
- The ShoreGear switch is managed by a Mitel server (either the HQ server or a DVS server, as configured on the switch edit page in Mitel Connect Director)

  This relationship dictates what Mitel server is used for call control (via TAPI) of the user's Mitel phone.

  Based on this relationship, the user's phone (physical or SoftPhone) gets all of its PBX call control from the ShoreGear switch it is assigned to, and can be managed by a Communicator for Windows client via the Mitel server that the ShoreGear switch is controlled by.

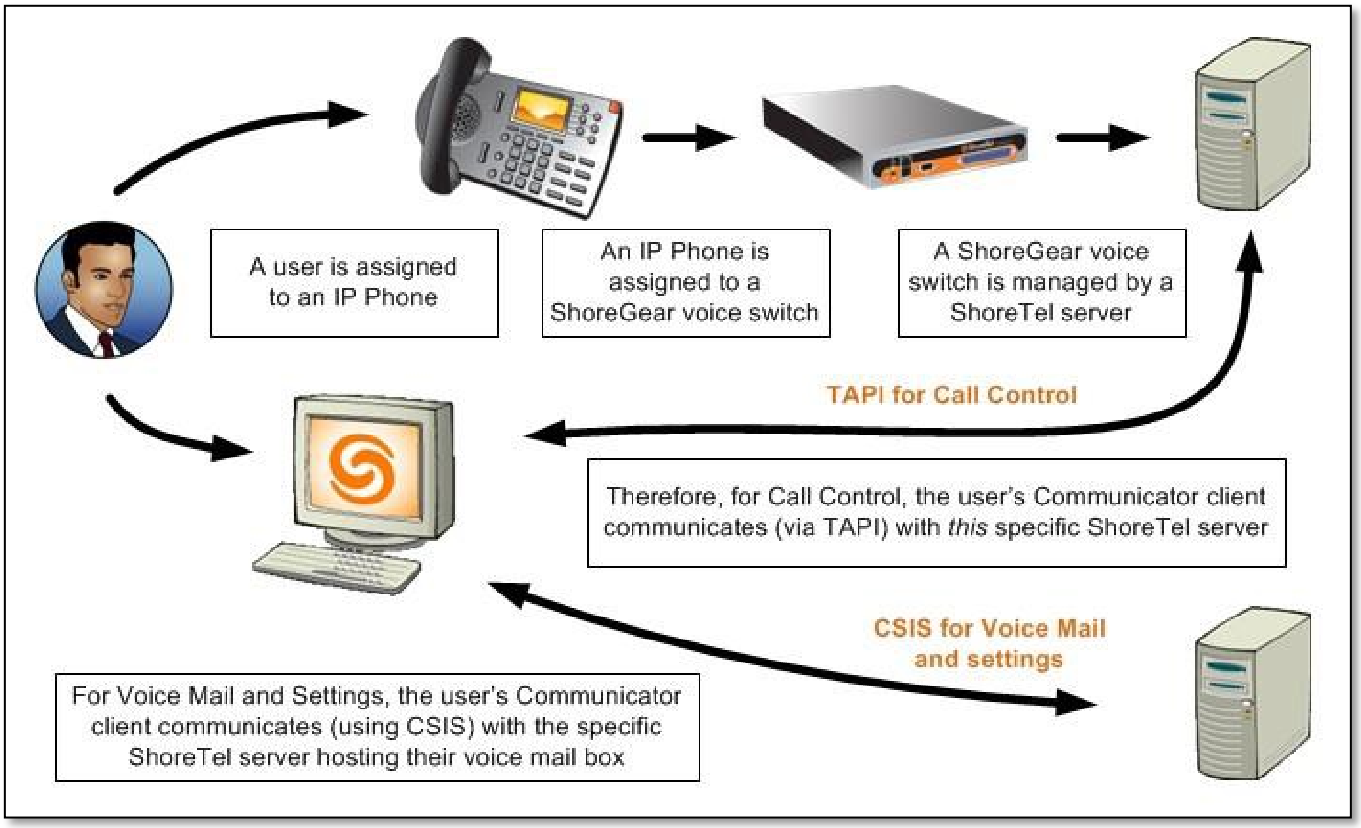
A user is assigned to an IP Phone
An IP Phone is assigned to a ShoreGear voice switch
A ShoreGear voice switch is managed by a ShoreTel server
TAPI for Call Control
Therefore, for Call Control, the user's Communicator client communicates (via TAPI) with *this* specific ShoreTel server
CSIS for Voice Mail and settings
For Voice Mail and Settings, the user's Communicator client communicates (using CSIS) with the specific ShoreTel server hosting their voice mail box

Communicator for Windows (in Mitel 12 and prior) establishes connectivity to both the user's voice mail server (using CSIS) and the server that manages the user's phone switch (using TAPI).

Communicator for Windows, when first launched, will prompt the user for a Mitel server. The user can enter any Mitel server (HQ or a DVS) but it is recommended that the user enter the HQ server (as explained below). The Communicator for Windows client will contact the user-entered Mitel server to authenticate. That server might happen to be the server that is configured as the user's assigned Voice Mail Server and might (or might not) also be the managing server for the user's current ShoreGear switch.

If another server (or servers) provide these services, the Communicator for Windows client will be instructed (redirected) to contact those servers.

If Communicator for Windows contacts the server as initially configured by the end-user and is redirected to a different Mitel server (or servers), the Communicator for Windows client will cache the new CSIS and/or TAPI server's IP address(es). The next time the Communicator for Windows client is started, it will initially contact the cached CSIS server and the cached TAPI server. If those servers are not available, the client will contact the original, user-entered Mitel server.

Once up and running, if the Communicator for Windows client loses connectivity to its assigned CSIS server it will lose the ability to interact with voice mail (play, record, and so on.) and will be unable to change settings such as changing Options or selecting Call Handling Modes. Features like Quick Dialer and the System directory tab will continue to function, but only because the data they display has already been retrieved via CSIS and is cached in RAM by the Communicator for Windows client. If the client is restarted it will lose that cached information, rendering Quick Dialer and the Directory tab non- functional, until CSIS is restored.

If the Communicator for Windows client cannot reach its assigned TAPI server it will lose all call control features.

In either of these cases the Communicator client will display a red 'X' icon in the top right corner.

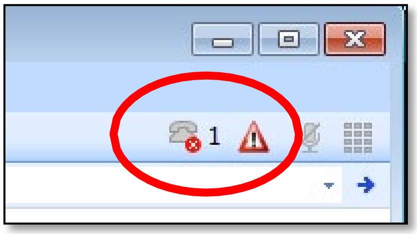

Note: Additional features within Communicator for Windows (such as IM presence, IM chat, desktop conferencing and web sharing, point-to-point video, and so on.) communicate with other services and components (such as a Mitel Service Appliance-100/400) using other protocols and might or might not continue to function depending on other considerations.

As stated earlier, if a Communicator for Windows client cannot reach its assigned (cached) CSIS server it will reach out to the originally entered Mitel server. If this originally entered server is the HQ server, it is possible for the 'orphaned' Communicator for Windows client to be successfully reassigned to another server for CSIS and TAPI control.

For example, if you have both an HQ server and a DVS server you could instruct all your users to enter the HQ server as their initially entered Mitel server when installing the Communicator client. If the DVS ever fails, you could use Mitel Connect Director to edit the users with the DVS as their voice mail server and reassign them to use the HQ server as their voice mail server. This results in the users' CSIS server being dynamically reassigned to the HQ server.

In addition you could use Mitel Connect Director to manually reassign the ShoreGear switches previously serviced by the DVS to be serviced by the HQ server. This results in the users' TAPI server being dynamically reassigned to be the HQ server.

In this example, each Communicator client, after failing to contact the cached DVS server, will revert back to contacting the HQ server (the originally entered server) and will learn that their

CSIS and TAPI servers have been reassigned to be the HQ server. In this fashion, voice mail services and call control services are restored to users even while the DVS is inaccessible.

Note: Be aware that forcing a user's voice mail box to a new server while their original voice mail server is unreachable can result in unpredictable behavior for previously stored voice mail messages. When the DVS server is returned to service, there is no guarantee that previously stored messages will be properly merged with new messages in the new storage location. Use this technique only as a last resort.

COMMUNICATOR FOR WINDOWS IN MITEL 13

Starting in Mitel 13, the Mitel Communicator for Windows client replaces both the CSIS and TAPI protocols with an updated Mitel proprietary communications protocol called CAS – the Client Application Services protocol. CAS simplifies the exchange of packets between Mitel devices and servers and allows the Communicator client to communicate with only the *one* server that is managing the ShoreGear switch controlling their phone. All call control, voice mail activity and settings changes are handled through the CAS communication with this one server. This one CAS server provides any necessary coordination with other Mitel servers for services like voice mail playback and with the HQ server for database changes.

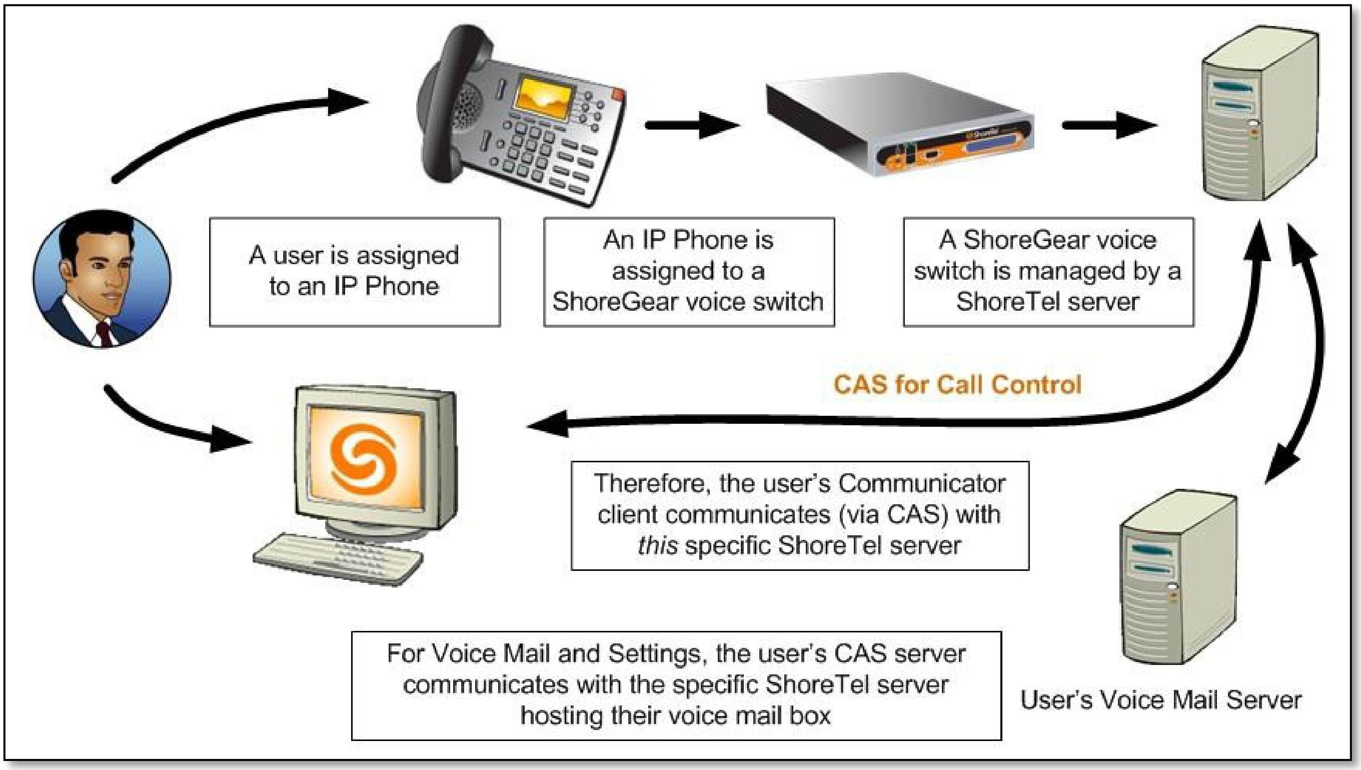
A user is assigned to an IP Phone
An IP Phone is assigned to a ShoreGear voice switch
A ShoreGear voice switch is managed by a ShoreTel server
CAS for Call Control
Therefore, the user's Communicator client communicates (via CAS) with *this* specific ShoreTel server
For Voice Mail and Settings, the user's CAS server communicates with the specific ShoreTel server hosting their voice mail box
User's Voice Mail Server

CommunicatorforWindows(inMitel13andabove)establishesconnectivity only to the server that manages the user's phone switch (using CAS).

During initial setup, as with prior versions, a user can enter the IP address or Fully Qualified Domain Name (FQDN) of *any* Mitel server and the Communicator clientwill be redirected to the server that is currently managing the ShoreGear switch controlling their phone. This assigned server information is cached by the client and the next time Communicator for Windows is launched it will reach out to the cached server first.

The client also uses CAS to learn the IP address(es) of the HQ server. If, upon re-launch, the client is not able to reach the cached server, it will reach out to the HQ server in an attempt to learn what server it should be assigned to. If a user has been reassigned to a phone controlled by a ShoreGear switch that is managed by a different server, the CAS service on the HQ server will redirect the Communicator client to the proper CAS server (which will in turn be cached for use during the next launch).

If the user's CAS server is unreachable but the HQ server is still reachable, the Communicator for Windows client will connect to the CAS service on the HQ server until its home CAS server is either reassigned or becomes reachable again. The CAS service on the HQ server allows only for the display of basic information (Directory content, Quick Dialer lookups) but all call control features, voice mail features, settings and Call Handling Mode changes are still unavailable until the home CAS server is either reassigned or becomes reachable again.

Note: Non-voice services (such as IM & IM Presence) would be unaffected by loss of CAS services and will continue to function.

COMMUNICATOR FOR WINDOWS (ALL VERSIONS)

Communicator for Windows, when launched, must be able to authenticate with either the HQ server or with a DVS enabled with a Distributed Database (page 82). As stated above, when a Communicator for Windows client is started (or restarted) it will contact its assigned DVS server. That DVS server must either be able to contact the HQ server to authenticate the user or that DVS must have a copy of the Distributed Database to authenticate the user.

If the HQ server is unreachable or the DVS server has not been enabled with a Distributed Database the Communicator for Windows client will not be able to authenticate and will not provide any services.

COMMUNICATOR FOR MAC

Communicator for Mac, introduced in Mitel 12, uses CAS (like Communicator for Windows in Mitel 13) (see page 60).

Communicator for Mac, when first launched, will prompt the user for a Mitel server. The user can enter *any* Mitel server (HQ or a DVS). The Communicator for Mac client will contact that Mitel server to authenticate. If that server is the managing server for the ShoreGear switch controlling the user’s phone, the client will authenticate and become operational. If not, the client will be redirected to the proper server and will authenticate with that new Mitel server.

Note: If you attempt to communicate with a DVS, that doesn’t have Distributed Database enabled, and the main HQ server is off-line, you will not be able to authenticate, until the HQ server is back online.

If Communicator for Mac contacts the server as configured by the end-user, and is redirected to a different Mitel server (the one that manages their phone’s ShoreGear switch), the Communicator for Mac client will cache that new server. The next time the Communicator for Mac client is started, it will again contact the original, user-entered server and be redirected. If that server is *not* available it will contact the cachedserver.

Note: This behavior is different than the Communicator for Windows client in Mitel 12 and below which will contact the cached server first.

COMMUNICATOR FOR WEB

Communicator for Web, introduced in Mitel 11, uses the CAS protocol like Communicator for Mac (and Communicator for Windows in Mitel 13). Unlike Communicator for Windows and Mac, Communicator for Web requires that the Mitel server name (or IP address) be entered directly in the URL address of the browser.

The user can enter any Mitel server (HQ or a DVS) in the URL when using Communicator for Web and will be redirected to the proper CAS server.

If the user has been configured in Mitel Connect Director to be authenticated by the local Mitel user database, the URL is:

http://<Server_IP_Address>/login

If the user has been configured in Mitel Connect Director for Active Directory authentication, the URL is:

http://<Server_IP_Address>

The Communicator for Web client will contact that Mitel server to authenticate. If that server is the managing server for the ShoreGear switch controlling the user’s phone, the client will authenticate and become operational. If not, the client will be redirected to the proper server and will authenticate with that new Mitel server.

Communicator for Web does *not* store or cache information about the redirection.

If a user notices they are being redirected to a different server they can change the URL to point towards the new server instead of the original server. This will avoid the redirection and authentication will take place slightly faster.

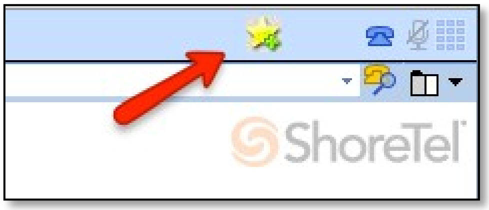

If Communicator for Web is redirected from the user-entered IP address to a different CAS server, you will see a 'Star' icon that provides a direct URL for easy bookmarking.

COMMUNICATOR CLIENT SERVICES: SUMMARY

All desktop Communicator clients (Windows, Mac and Web) can be pointed at *any* Mitel server for their initial, first-time launch, and for any and all subsequent launches. Whichever Mitel server they connect to will redirect the client to the proper server(s).

If the server that manages the ShoreGear switch that controls the user's phone is not reachable, all call control features (enabled via TAPI or CAS, depending on the client) of the Communicator client will stop functioning. There is no failover to another server.

Note: In the case of Communicator for Windows in Mitel 12 and prior, be aware of the different protocols and destination servers that are used. It is possible that connectivity to one server could be down but connectivity to another server remains functional. This can create situations where notification of new voice mail messages, playback of voice mail over the PC speakers and changing Call Handling Modes is functional (CSIS server is up) but call control is not (TAPI server is down); or vice-versa.

Note: If a user is a member of a Workgroup, their Communicator for Windows client (in Mitel 12 and prior) might need to contact one server (for call control), another server (for personal voice mail) and yet another server or servers (for Workgroup voice mail).

If a Communicator client is up and running (driven by the HQ or a DVS) and the server managing the user's Voice Mail box is unreachable, call control activities will still function (via TAPI or CAS) but Call Handling Mode/Options/Preferences settings and voice mail features will stop (since the CSIS server is unreachable).

If a Communicator client is up and running via a DVS and connectivity between the DVS and the HQ server fails, the user will still be able to perform call control and voice mail actions but will not be able to make database changes such as selecting a new Call Handling Mode or altering settings in the Options/Preferences dialog box. This is true for CAS clients as well as TAPI/CSIS clients.

Note: See the section on Distributed Databases on page 82 for more information allowing CHM changes even when the HQ server is unreachable.

Note: For additional details on the failover features of SoftPhone, part of the Communicator forWindows Professionalaccess-level,seeApplicationNote 10338: 'Mitel SoftPhone – Features, Functions and Details'.

If a user reassigns their extension to a different location, the server they are connected to for call control services (TAPI or CAS) will switch as needed. For example: a Headquarters user who travels to a remote office and reassigns their extension to an IP phone at the remote office will have their Communicator client's TAPI (or CAS) connection redirected to the DVS at the remote office controlling their new phone. If they return to their hotel and launch the Mitel SoftPhone the TAPI/CAS services will be reassigned to the server that controls the ShoreGear switch that the SoftPhone is associated with.

ADMINISTRATION AND REPORTING

The HQ server is the sole point of administration for the Mitel system and is the sole point used for running Call Detail Record (CDR) reports.

If the HQ server is up and running at the Headquarters site but unreachable by a remote- site administrator (for example, due to a WAN outage), the remote administrator would be unable to access the Administration portal to make changes or run reports but no data or services would be lost at the HQ site.

If the HQ server was down (e.g. due to a hardware failure, disconnected Ethernet connection, reboot, or scheduled or unscheduled maintenance such as an OS upgrade, backup, patch, and so on.) all administrators would be unable to make configuration changes or run reports.

As described in the sections above, most other services that the HQ server performs are either replicated (AA, ACC) or distributed (Workgroups, Group Paging) or offer a robust failover/fallback mechanism (Voice Mail, CDR collection). But some services are still dependent on a single server and making the Mitel server(s) fully redundant should be considered for any large scale deployment

See the section 'Database Changes and Server Redundancy' on page 78 below for further details on HQ database replication and server redundancy options.

APPLICATION SERVICES: SUMMARY

In summary, the resiliency and high-availability options for Application Services include:

1. Explicit Backup Extensions provide an alternative call handling entity for Workgroups, Hunt Groups and Bridged Call Appearances when a ShoreGear switch or Mitel server is unreachable.
2. Services like Auto Attendants and the Account Code Collection process are fully replicated to all servers and all V-switches. This reduces WAN traffic by handling requests by the local, or 'nearest,' Application Services end point.
3. Services like Workgroups and Group Paging are distributed, allowing local servers to provide local services. This increases the capacity of the Mitel system and, in the case of Workgroups, allows one server to back up another server, retaining full features and functionality even during a server outage.
4. Services like Voice Mail are distributed to multiple end points, with all end points acting as a 'single-image voice mail system.' All voice mail end points work together to provide a 'Store and Forward' safety net and failover mechanism.
5. A robust voice mail prompt mechanism is able to provide a friendly and functional caller experience, even during periods of outage.
6. An intelligent 'Store and Forward' mechanism for delayed delivery of Call Detail Records is leveraged during communication outages.
7. The Communicator client is empowered by one or more specific servers (for TAPI, CSIS and CAS). If those specific servers are not reachable the Communicator client does not automatically fail over. You can manually reassign a user (or ShoreGear) to induce a failover by altering server assignments for users and ShoreGear switches in Mitel Connect Director.

MULTI-SITE REDUNDANCY

So far, we have discussed the following:

- Resiliency features of the virtual or hardware-based ShoreGear switches.
- Distributing call control to multiple ShoreGear switches at a site to reduce the impact of an outage.
- Distributing trunks to multiple ShoreGear switches at a site to reduce the impact of an outage.
- Connecting both, redundant LAN connectors on each hardware-based ShoreGear switch to avert trouble during an Ethernet outage.
- Deploying Uninterruptable Power Supplies for power protection and power backup.
- Connecting analog PSTN trunks and emergency analog phones to Power Fail Transfer ports for emergency dialing during prolonged power outages.
- Configuring explicit Backup Extensions to provide alternative call routing for some services.
- Configuring distributed Workgroups to survive server outages.
- Understanding the native 'Store and Forward' failover mechanisms of the Voice Mail service.
- Understanding the native 'Store and Forward' failover mechanisms of the Call Detail Record collection process.
- Understanding the fully distributed nature of services like Auto Attendant Menus and the Account Code Collection process.
- Defining the features and functionality of all Application Services end points.

Let's continue the discussion by focusing on the resiliency and redundancy features specifically tailored to multi-site deployments.

MULTI-SITE CONFIGURATION

As your organization expands from a single site with one Mitel Application Server to multiple sites and multiple Application Services end points (DVS and V-switches) we should start by reviewing the important design strategies and configuration settings for a multi-site and multi-server Mitel system.

- Each user must be assigned to a specific Voice Mail server. In general, you should assign users to a voice mail storage end point that is as close to their location as possible. This often will be a DVS or V-switch at their local site. Alternatively you might choose to locate the voice mail servers close to the inbound trunks. For example, if most of your trunks enter at a Colo Site, consider hosting the user's voice mail boxes on servers at the Colo Site.
- Voice Mail server assignment is configured on the General tab in the Individual User's edit page.

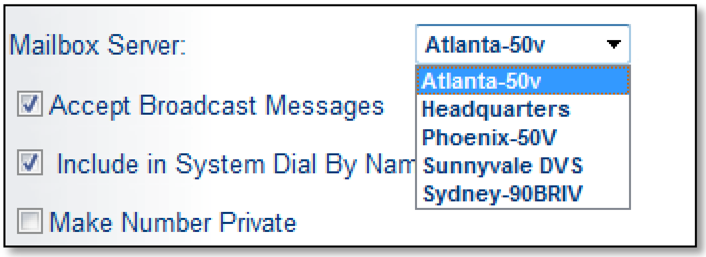

Mailbox Server: Atlanta-50v ▾

☑ Accept Broadcast Messages

☑ Include in System Dial By Nam

☐ Make Number Private

- Each ShoreGear switch must be assigned a managing Mitel server. In general, you should assign ShoreGear switches to be managed by the closest Mitel server. This will often be a DVS server at the same site as the ShoreGear switch. Managing server assignment is configured in the switch's edit page.

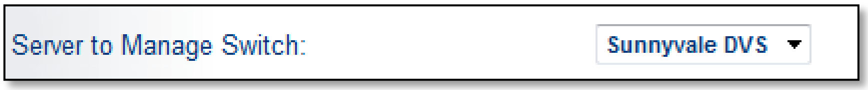

Server to Manage Switch: Sunnyvale DVS ▾

- IP Phones are assigned to ShoreGear switch resources via IP Phone Address Maps. IP Phone Address Maps associates IP address ranges with Mitel sites. For each location, be sure to add the *entire*range (or ranges) of IP addresses that a site's phones might be assigned (either via static IP assignment or DHCP assignment). Ranges should also be large enough to cover any computers that will use SoftPhone as well. IP Phone Address Maps are configured in Mitel Connect Director under 'IP Phones > IP Phone Address Map'

Note: If the Mitel system fails to find an appropriate IP Phone Address Map, for an IP Phone IP address, it will assign a ShoreGear switch from the Headquarters site.

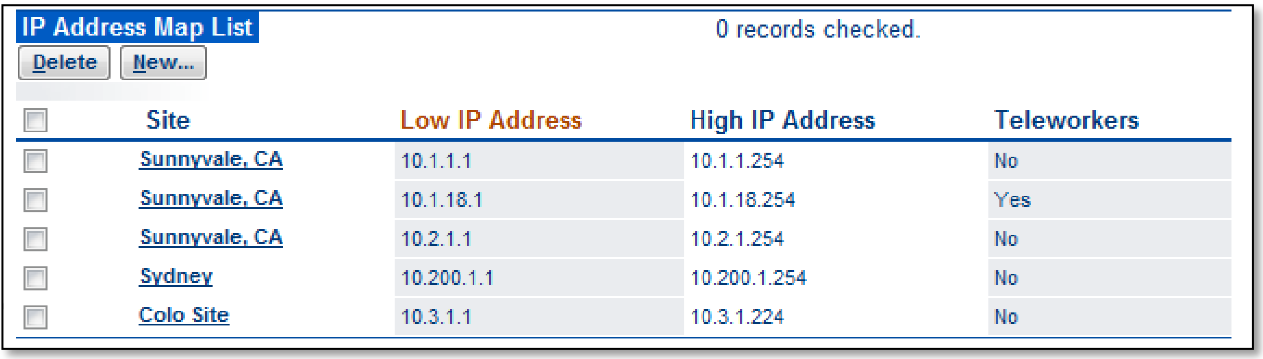

IP Address Map List

0 records checked.

☐	Site	Low IP Address	High IP Address	Teleworkers
☐	Sunnyvale, CA	10.1.1.1	10.1.1.254	No
☐	Sunnyvale, CA	10.1.18.1	10.1.18.254	Yes
☐	Sunnyvale, CA	10.2.1.1	10.2.1.254	No
☐	Sydney	10.200.1.1	10.200.1.254	No
☐	Colo Site	10.3.1.1	10.3.1.224	No

REMOTE SITE SURVIVABILITY – CALL CONTROL AND PSTN ACCESS

It is completely acceptable to deploy multiple handsets at a remote location yet centrally locate the ShoreGear call control appliances and Application Services at a centralized location, such as a server room at the Headquarters site.

This is common for campus environments that have high speed, reliable and redundant links between buildings, *and* that have a stated business policy of centralizing storage, centralizing management and centralizing trunking.

A centralized solution is perfectly acceptable … unless an unexpected cable cut disconnects the remote site(s) from the centralized resources. Such an event would disable *all* phones and *all* trunk services to the affected sites. Application Services like Voice Mail, Auto Attendant menus and Communicator services would also fail.

A centralized design is also less desirable when the data links to the remote sites are of low bandwidth or are unreliable, such as global deployments with unmanaged Internet-based VPN connections between sites.

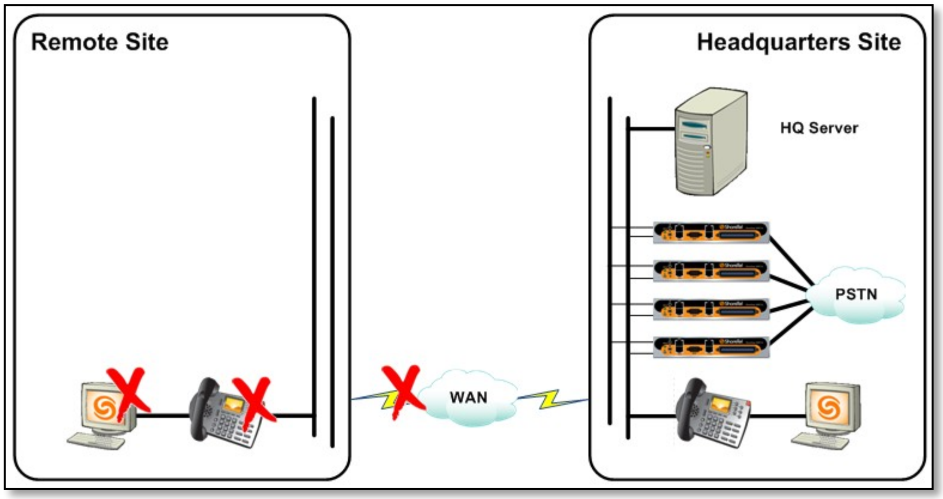

A remote site with a WAN outage disables all telephony and UC features for that site

Distributing your ShoreGear call control switches to the remote site will avoid this dilemma.

Consider a Headquarters site of 150 users and a remote site of 50 users. Instead of having all 200 IP Phone resources at the central location (in the form of four ShoreGear-50 switches) relocate one of the ShoreGear-50 switches to the remote location.

Additionally, deploying (or moving) local PSTN trunk(s) to the remote site will enable local site-survivable telephony features for the remote site.

Also, consider deploying trunking services from *different* PSTN carriers (carrier diversity) to protect against failures in a specific carriers network.

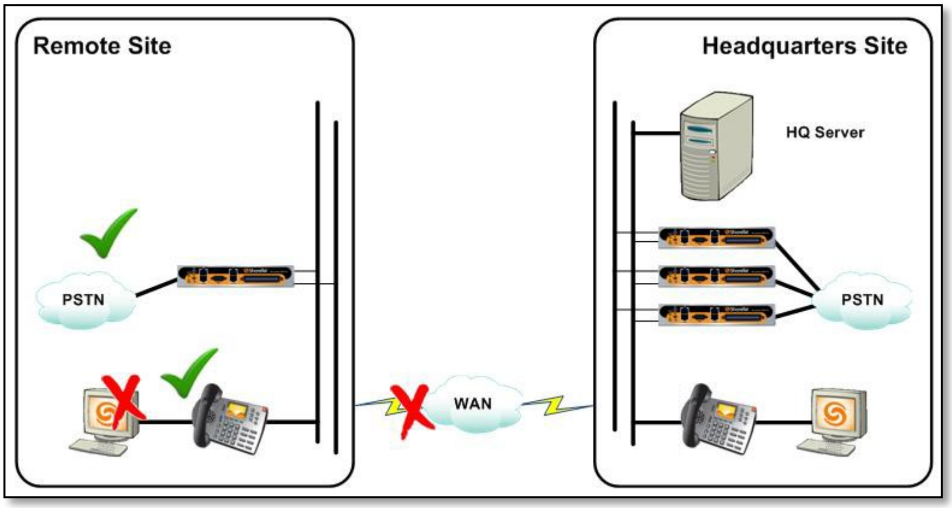

Moving one ShoreGear switch and adding (or moving) trunks to the remote site provides local site-survivable call control and trunks.

The quantity of ShoreGear resources and the quantity of trunk resources has not materially changed – they have just been distributed. For little to no cost difference, you get complete site-survivable trunking and call control.

But Voice Mail and other Application Services, including Communicator control, are not available in this WAN failure scenario.

REMOTE SITE SURVIVABILITY – APPLICATION SERVICES

To enable local, site-survivable Auto Attendant menus and Voice Mail services, we can change the ShoreGear-50 switch to a ShoreGear-50V switch.

The V-switch has all the same call control and trunking features of a standard ShoreGear switch but adds an internal solid-state hard drive for storage, and adds Application Services capabilities including Auto Attendant Menus, Voice Mail and Account Code Collection.

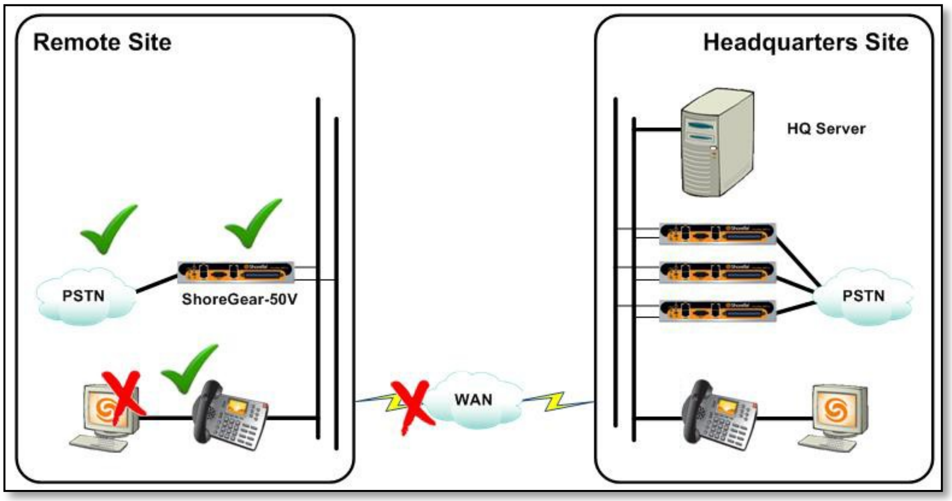

Changing to a V-switch adds survivable Auto Attendant menus and Voice Mail.

We now have locally survivable Voice Mail, Auto Attendant menus, and Account Code Collection services. But we have not yet restored Communicator services to our remote site during this extended WAN outage.

To do that, we need to deploy an additional component: a Mitel Distributed Voice Server, or DVS.

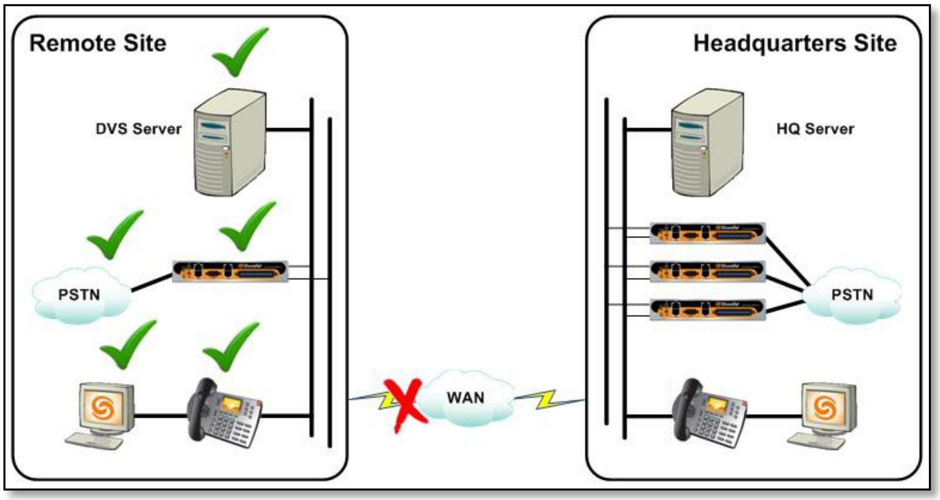

Adding a DVS server provides survivable Communicator services.

With the addition of a DVS, during an extended WAN outage the remote site has local survivability for 100% of its services including Communicator client control and other server- based features such as Workgroups.

Note: For more information on DVS servers see the section Application Services on page 45 above.

Note: Distributed Databases (page 82) and Distributed Workgroups (page 50) are mutually exclusive. You cannot enable both in the same single-image Mitel system.

REMOTE SITE SURVIVABILITY – REDUNDANT CALL CONTROL

We have made the remote site fully survivable in case of a WAN outage. It has local call control, local trunks, local Application Services like Workgroups, Auto Attendant menus and local Communicator client control. All of this is achieved with minimal added costs and/or components.

Note: All of the best practices for a single site also apply to the remote site. Features such as connecting both LAN ports, choosing to deploy several smaller ShoreGear switches instead of one large switch, distributing your trunks to multiple switches, adding analog trunks and analog phones for power fail emergency services, etc. should be followed for all sites.

But, what if the service disruption is not a WAN outage but rather a ShoreGear switch outage?

If you have a large remote site (75 or more users), you might consider adding extra IP phone capacity by deploying additional ShoreGear switch resources (aka N+1 Spare). That would provide full call control redundancy in case of a ShoreGear switch failure.

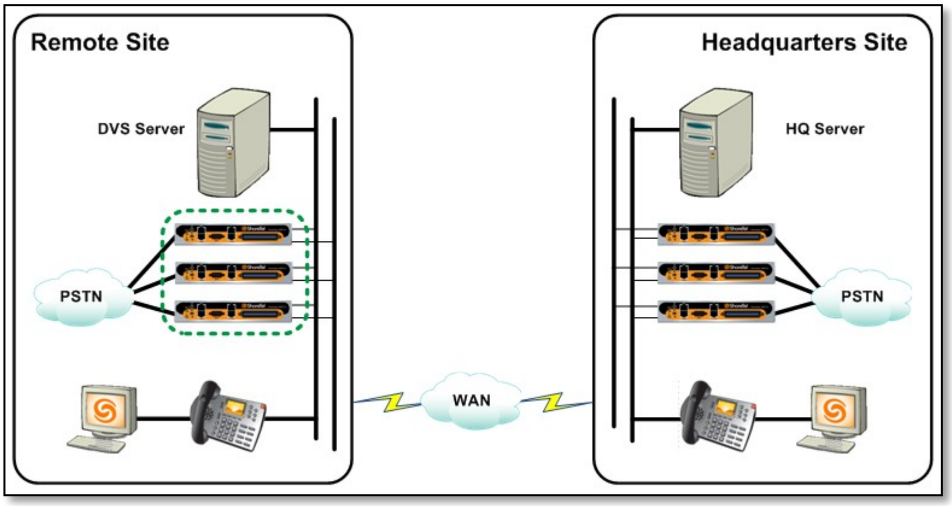

Adding extra ShoreGear capacity provides IP Phone failover and Call Control Redundancy.

This is an excellent solution. However, it can become cost prohibitive if you have many sites *or* if each remote location has a relatively small number of users. For example, if you only need one ShoreGear switch at each of 20 small sites, adding one more ShoreGear switch *per site* would increase your costs dramatically.

Mitel has addressed this with a feature called the 'Spare' switch.

Instead of adding an extra ShoreGear switch as an *active*ShoreGear switch to a single site, you can add an extra ShoreGear switch designated as an *inactive* 'Spare' switch. When a ShoreGear switch fails, or a site adds additional IP phones that extend beyond its local capacity, the Mitel system will automatically assign the 'Spare' switch to the site that needs the extra IP Phone resources.

The inclusion of Virtual Mitel IP Phone switches in Mitel 14 allow you to deploy as many 'Spare' switches as needed without incurring additional costs (that is as long as you have the sufficient Virtual Host resources to accommodate). The Virtual IP Phone switches only require licenses for attached IP phones, meaning when they're defined as 'Spares' there aren't any IP phones being managed by the Virtual switch, thus there aren't any licensing requirements. Once the 'Spare' Virtual IP Phone switch becomes active and IP phones register to it, then you will have 45 days to get into license compliance, either by correcting the 'issue' that caused the 'Spare' Virtual IP Phone switch to become active or purchasing the appropriate number of 'Mitel Virtual Phone Licenses'.

Note: A 'Spare' switch will be activated whenever a site needs additional IP Phone resources. This might be due to a ShoreGear switch outage, but it might just as easily be caused by adding more IP Phones than you had originally planned for. Be careful not to trigger the activation of your 'Spare' switch accidentally by adding a few extra IP Phones at a remote site. Consider adding the 'Spare Switch Activation' event to your email alert notification list using Mitel Event Filters (see page 139).

When activated, a spare switch is assigned, in its entirety, to the site in need. If you have multiple spare switches, they will be used one by one as needed to accommodate multiple failures across your enterprise. Spare switches are commonly located at the Headquarters site but can be physically located, and added to, any level of the Mitel Site Tree Hierarchy.

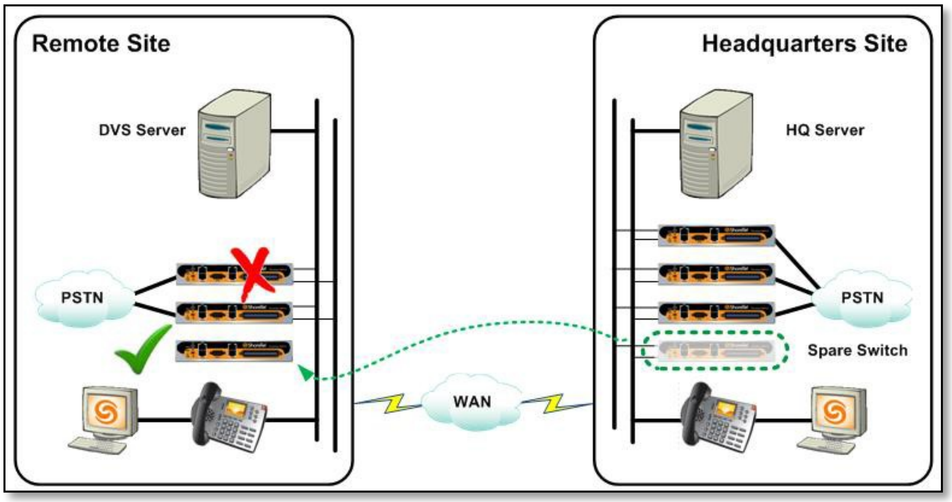

Adding an extra ShoreGear as a ‘Spare’ switch works across sites like an N+1 switch works for the local site.

Spare switches are activated according to the Site Tree Hierarchy. When a site needs additional IP Phone resources and IP Phone Failover is enabled, the Mitel system will look at that local site, and then continue up the tree to the parent site, then the grandparent site, and so on. until the root of the tree (HQ) is reached. The first available spare switch found will be assigned to the site in need.

REMOTE SITE SURVIVABILITY – IP PHONE CONFIGURATION SWITCHES

Refer to *Application Note 10298 ‘Mitel IP Phone Failover Features’* for additional details on the selection and placement of IP Phone Configuration Switches in a multi-site deployment.

MULTI-SITE REDUNDANCY: SUMMARY

In summary, the resiliency and high-availability options for a multi-site deployment include:

1. Distributing your call control ShoreGear appliances to the remote sites for local site- survivable call control.
2. Distributing your trunking to the remote sites for local site-survivable PSTN access.
3. Deploy trunks from different PSTN circuit providers to achieve carrier diversity.
4. Adding a V-switch provides local, survivable Auto Attendant menus and Voice Mail services.
5. Adding a DVS server renders all remaining Application Services locally survivable too, such as Workgroups and Communicator control.
6. Add extra IP phone capacity at your larger remote sites to achieve full call control redundancy (N+1).
7. Add extra IP phone capacity as a ‘Spare’ switch to serve as a more economical backup for call control redundancy for any/all remote sites.

DATABASE CHANGES AND SERVER REDUNDANCY

The Mitel Headquarters (HQ) server is the primary administration point for making changes to the master configuration database.

Examples of changes that need to be written to the configuration database include:

- Adding or editing a user (or any other entity) in Mitel Connect Director.
- A user changing their Call Handling Mode.
- A user adjusting the programming or settings for their Call Handling Mode.
- The system altering the mode of a Hunt Group due to a Schedule.
- An Agent logging into their Workgroups.
- A new IP phone attempting to register with the Mitel system for the very first time.

If the HQ server is the only Mitel server in your topology then, in each of these cases, the requested change must be able to be communicated through the chain of Mitel devices (IP Phones, Communicator clients, ShoreGear switches) all the way to the HQ server. If a remote site has lost connectivity to the HQ server (and has no DVS server), these changes cannot be completed. Attempting to make such a change without connectivity to the HQ server will result in the change *not* being implemented and an error message being displayed on the user's phone or Communicator client.

This makes the HQ server, when operating as the sole server in a Mitel system, a single point of failure.

Several best practices can be brought to bear to improve, and even eliminate, this single point of failure, including:

1. Enable local 'Call Forward Always' overrides (*72/*73)
2. Enable Distributed Database features (Mitel 11 and higher)
3. Enable a server redundancy mechanism using one of the following:

- Server redundancy via Double-Take
- Server redundancy via VMware
- Server redundancy via Hyper-V (Mitel 14.2)

Each of these topics is discussed below.

‘CALL FORWARD ALWAYS’ OVERRIDE (*72/*73)

When a user changes their Call Handling Mode (CHM) from one mode to another, this change alters the main Mitel database stored on the HQ server. Therefore, a user can only change their CHM when connectivity to the HQ server is available. Changes to a user’s Call Handling Mode can be initiated from the user’s IP Phone (by pressing the ‘Mode’ soft key, or by dialing into your mailbox and selecting the Options menu) or from the user’s Communicator client (by using the Call Handling Mode drop down menu or from within the Options/Preferences dialog window).

A situation can occur when a user’s current CHM mode is set with a ‘Call Forward Always’ (CFA) condition. If the user sets their CHM mode to a CFA setting and, afterwards the user loses connectivity to the HQ server, the user has no way to change their CHM mode (or the settings configured for that mode) and can be ‘stuck’ in a ‘Call Forward Always’ condition.

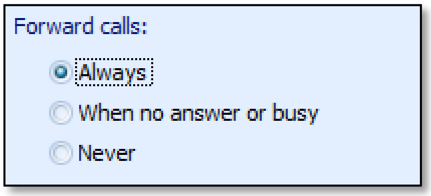

Forward calls:

● Always
○ When no answer or busy
○ Never

‘Call Forward Always’ setting in the Communicator client

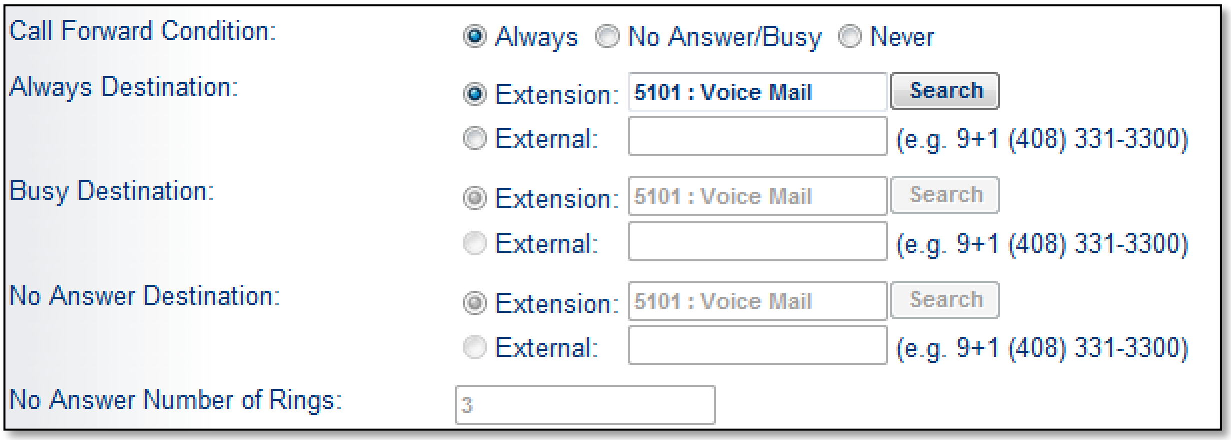

Call Forward Condition: ● Always ○ No Answer/Busy ○ Never

Always Destination: ● Extension: 5101 : Voice Mail [Search]
○ External: (e.g. 9+1 (408) 331-3300)

Busy Destination: ● Extension: 5101 : Voice Mail [Search]
○ External: (e.g. 9+1 (408) 331-3300)

No Answer Destination: ● Extension: 5101 : Voice Mail [Search]
○ External: (e.g. 9+1 (408) 331-3300)

No Answer Number of Rings: 3

‘Call Forward Always’ setting in the Mitel Connect Director

To remedy this, a feature can be enabled that allows a user to override the 'Call Forward Always' condition in their local ShoreGear switch and allow calls once again to be routed through and ring the user's phone.

Configuring a registry setting on the HQ and all DVS servers enables this capability. This setting, once added, will be immediately communicated to all ShoreGear switches managed by that server. No reboot of the server or its managed switches is required. If an HQ connectivity issue occurs (after implementing this registry change) the user can override a 'stuck' Call Handling Mode 'Call Forward Always' condition to allow calls to ring through.

From the user's extension (the actual handset of the extension that is in a CFA state) do the following:

- Press '*73' – This overrides a Call Forward Always condition (if it exists) and allows calls to ring through.
- Press '*72' – This cancels the effect of *73 (do this when connectivity to the HQ server is restored and you wish to return to normal operation).

These feature codes are *only*usable from the actual phone that is 'stuck' in a CFA condition due to lack of connectivity to the HQ server. These feature codes only affect 'Call Forward Always' conditions, and have no effect for users set in a 'Call Forward Busy/No Answer' mode.

These feature codes *only* work if the ShoreGear switch's managing Mitel server hashad the registry setting changed *prior* to the HQ connectivity outage.

This CFA override only works on user extensions and cannot be used for other Mitel entities such as Hunt Groups, Workgroups or Auto Attendants.

Note: See the section 'Application Services' on page 45 for a discussion of explicit Backup Extensions and failover mechanisms for Hunt Groups, Workgroups or Auto Attendants.

This feature was added in Mitel 5.

Note: Use caution when editing a server's registry. Editing a server's registry should only be performed by trained personnel. The Mitel Technical Assistance Center (TAC) will not assist with registry changes.

To configure this registry setting, do the following:

1. On the Mitel server open REGEDIT.EXE

   This setting must be made to all Mitel servers individually – HQ and DVS).

2. Navigate to *HKEY_LOCAL_MACHINE\SOFTWARE\Shoreline Teleworks\Telephony Management server\Settings*]. On some servers, it might be *HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Shoreline Teleworks\Telephony Management Server\Settings*].
3. If it does not already exist, add a new key named 'SwitchDebug'.
   - Key Type: String
   - Key Name: SwitchDebug
4. Edit the value of the new key, adding 'debug options', the parameter 'allow_cfa_deactivation' and a value of '1' followed by a trailing period.

   String Value: *debug_options allow_cfa_deactivation 1.*

Note: The syntax of this string is important: it *must* have the term 'debug_options', the parameter name 'allow_cfa_deactivation', the value (1 to activate, 0 to deactivate), and then a period.

Note: If the SwitchDebug registry entry already exists, multiple settings are delimited by spaces, with a single terminating period.For example, '*debug_options existing_setting 2 allow_cfa_deactivation 1*.'

DISTRIBUTED DATABASE

Prior to Mitel 11, the only writable copy of the master configuration database was located on the solitary HQ server. All other end points, including standard ShoreGear switches, V-switches and DVS servers, obtained a read-only copy of relevant portions of the database from the HQ server.

Starting with Mitel 11, the ability to distribute a limited *writable* copy of the master database to DVS servers was added. This enables more than one server to manage changes to the master database. Synchronization of changes is managed and performed automatically by the multiple, distributed database servers.

Only three specific database changes are allowed to be made by a DVS server that has a writable, distributed copy of the database. These include:

- Changing a user's Call Handing Mode
- Changing a Workgroup Agent's login/logout state
- Updating a SIP device's registration

While these are the only writable functions allowed by a DVS server, with the distributed database parameter enabled, the DVS server can also provide functionality for server required features (that is, Visual Voice Mail and History), thus alleviating the HQ server from these tasks and allowing for more distributed functionality.

In addition, enabling a DVS to host a local, writable copy of the database enables the DVS server to reboot without needing connectivity to the HQ server. A DVS without a local distributed copy of the database, if rebooted, requires connectivity with the HQ server to return to functionality. Rebooting a standard DVS without HQ server connectivity will cause the DVS to become inoperable.

Enabling a DVS to host a local, writable copy of the configuration database requires a reboot of the DVS server.

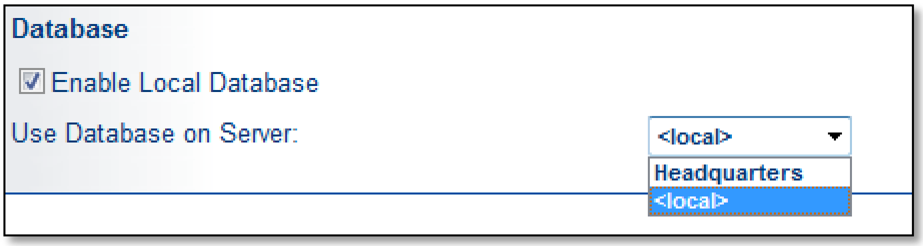

Database

☑ Enable Local Database

Use Database on Server: <local>

Enabling a DVS to use a distributed (local) copy of the master database.

Configuration changes via Mitel Connect Director must still be performed on the HQ server, even if the Distributed Database features have been enabled on one or more DVS servers.

Administration and CDR report generation are still dependent on the HQ server.

The Distributed Database feature is supported on DVS servers only. It is not supported on V-switches. V-switches can be configured to obtain their local (read-only) copy of the configuration database settings from either the HQ server (the default) or from a Distributed Database copy on a DVS. This is configured within the V-switch's edit page in Mitel Connect Director.

CHANGING CALL HANDLING MODES WITH DISTRIBUTED DATABASE ENABLED

Consider the following:

- A user is assigned to an IP Phone
- That IP Phone is managed by a ShoreGear switch
- That ShoreGear switch is managed by a DVS server

If that DVS server has been enabled with a writable copy of the Distributed Database, then that user can change from one Call Handling Mode to another. The change will be serviced entirely by the local DVS. Connectivity to the HQ server is irrelevant.

If, on the other hand, the DVS had *not* been enabled with a Distributed Database, and the DVS had lost connectivity with the Master Database on the HQ server, the local user would not be able to change from one Call Handling Mode to another. Such a request would require a change to the master database on the HQ server, which is unreachable, and would fail.

Changes to the programming/settings for the user's Call Handling Modes are still controlled by the HQ server. Enabling the Distributed Database on the DVS server only enables switching *between* modes not changing the details of a CHM mode.

DISTRIBUTED DATABASE VS. DISTRIBUTED WORKGROUPS

It is important to note that Distributed Databases and Distributed Workgroups are mutually exclusive.

If any server in the enterprise-wide, single-image Mitel system has a Distributed Database enabled, you will not be able to assign a Workgroup to be managed by a DVS server.

Conversely, if any server in the enterprise-wide, single-image Mitel system has a Workgroup assigned to a server other than the HQ server, you will not be able to deploy a Distributed Database on any DVS server.

WORKGROUP AGENT LOGIN STATE

If a Workgroup Agent's phone is serviced by a ShoreGear switch that is managed by a DVS with a writeable copy of the Distributed Database, then changes to the Agents Workgroup login state is serviced entirely by their managing DVS server. Connectivity to the HQ server is not required.

This allows an Agent to change between the Logged-in state and the Logged-out state.

Note: The Wrap-up state is simply a transition to different aspects of the Logged-In state and is always handled by the server controlling the Workgroup (whether HQ or a distributed Workgroup on a DVS).

Note: Since Distributed Databases (DDB) and Distributed Workgroups (DWG) cannot be combined in the same enterprise, the benefit of a DDB-enabled DVS being able to support a Workgroup Agent login state change is of little practical value. If the HQ server is down then the Workgroup is down and Agent state becomes irrelevant. This feature will become a more valuable benefit in a future release.

SERVER REDUNDANCY

The distributed nature of applications and services in the Mitel architecture creates a robust, resilient and highly fault-tolerant UC system.

Yet, there are still features and capabilities that are server dependent. These include:

- Communicator client services
    - If the Mitel server that controls your Communicator client is unreachable, the Communicator client becomes inoperable
- Voice Mail retrieval
    - The ‘Store-and-Forward’ features of the distributed Voice Mail service ensure that new messages are captured and retained, but access to those messages is delayed until your assigned voice mail server is restored
- Collecting of CDR records
    - If a server is unreachable, CDR information for all ShoreGear switches managed by that server are lost
- Adding new components
    - To add a new ShoreGear switch or IP phone, the server intended to manage that device must be reachable
- Workgroups and Route Points
    - Both Workgroups and Route Points require connectivity to a specific server for complete functionality

      Additionally, the HQ server plays a unique role in that it is the only server that runs Mitel Connect Director and is the only server that manages changes to the master configuration database (other than those few items allowed via the Distributed Database feature as described on page 82 above).

      Specifically, if the HQ server is offline, an organization cannot:
- Add or change any configuration programming such as users, Auto Attendant menus or Workgroup configuration
    - The HQ server runs the only instance of the Administration portal, Mitel Connect Director
- Run Call Detail Record (CDR) reports
    - Report generation is performed within Mitel Connect Director running on the HQ server
- Add new equipment such as IP Phones, ShoreGear switches or DVS servers
    - New additions are prohibited
- The rebooting of many entities is not allowed when the HQ server is offline because this function is performed via Mitel Connect Director.
- Re-launch the Mitel SoftPhone
    - No launching or re-launching of any SoftPhones
    - Previously running SoftPhones remain functional
- Invoke schedule changes for Hunt Groups

Only the HQ server can initiate a Hunt Group schedule change.

For some organizations, these tasks are considered mission critical and need to be made redundant. In these circumstances the Mitel servers can be made fully redundant.

Note: For a complete list of features and functions of the HQ server and DVS servers see Appendix A: HQ, DVS and V-switch Features.

Server redundancy is achieved for organizations with a single server (HQ only) through the use of solutions involving 3rd party offerings from Stratus, Double-Take and VMware.

For enterprises with multiple Mitel servers, all servers can be made redundant through the use of virtualization (VMware and Microsoft Hyper-V solutions).

Server redundancy comes in several flavors and each approach has its benefits. You should review and test each option thoroughly before deciding on a server redundancy solution.

HQ SERVER REDUNDANCY WITH DOUBLE-TAKE

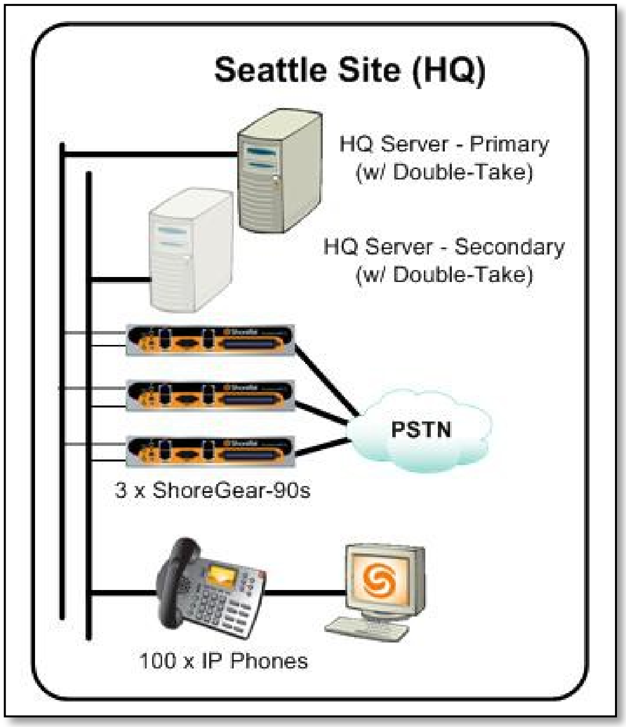
Seattle Site (HQ)
HQ Server - Primary (w/ Double-Take)
HQ Server - Secondary (w/ Double-Take)
PSTN
3 x ShoreGear-90s
100 x IP Phones

For an organization that requires redundant Application Services and has deployed a single HQ server, Double-Take offers a very good solution. Double-Take is a software replication tool that allows the replication of hard drive content and the monitoring of a primary server by a secondary server.

The failure of the primary server will invoke the activation of the secondary server. After a short start-up period, the secondary server will come on-line and begin providing the identical services that the primary server had been providing. This is perfect for providing redundancy for the Mitel HQ server.

Double-Take is also beneficial in that, with adequate bandwidth between servers, you can deploy the primary and secondary HQ servers in different IP subnets and even in different geographical locations. This enhances your survivability by enabling automatic recovery if the primary site is stricken by a major disaster.

Note: See the version specific Mitel documentation for Double-Take that corresponds to your Mitel server version for complete details on implementing Double-Take.

VMWARE HIGH AVAILABILITY (HA)

For Enterprises that have VMware, High Availability is a useful strategy for HQ server and DVS server application availability.

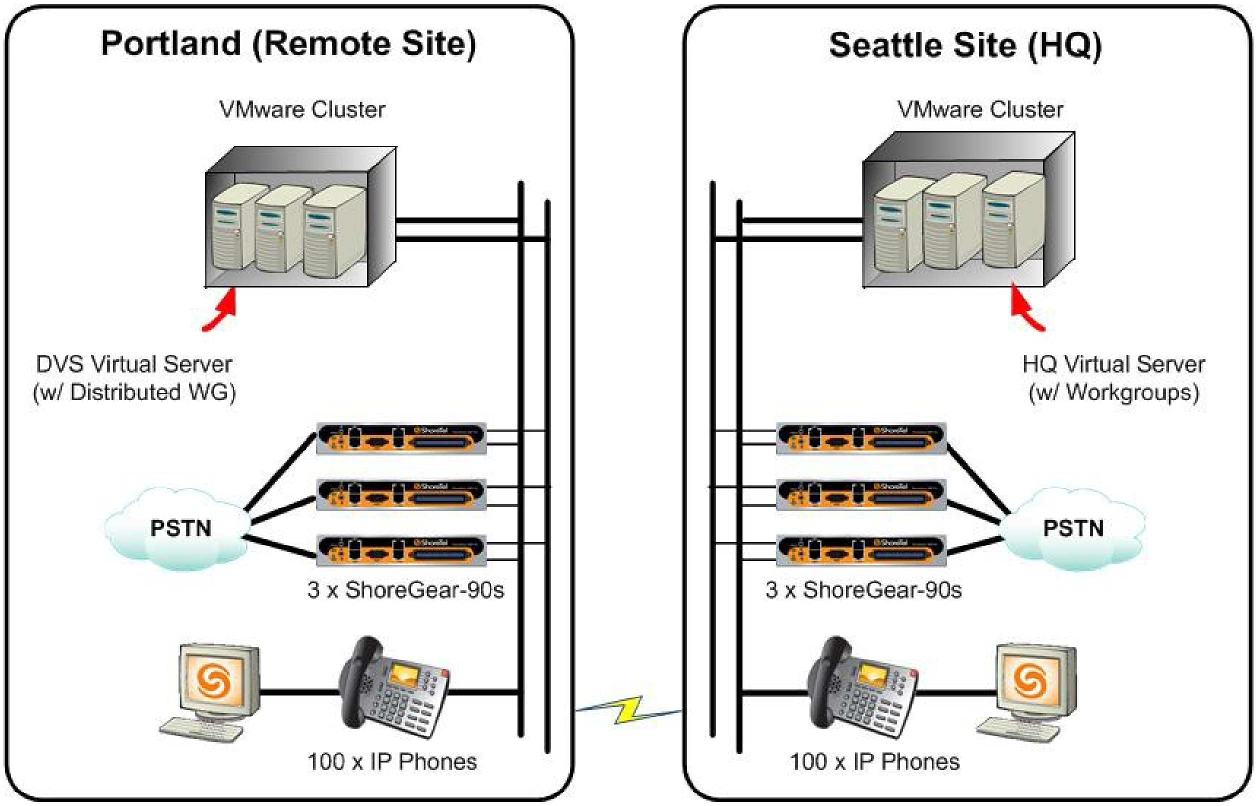

Example: A two-site topology with redundant servers in both locations.

If the VMware host server that the HQ server (or DVS server) are virtualized on fails, HA will bring them back up on another server within the cluster.

MULTIPLE SERVER REDUNDANCY WITH MICROSOFT HYPER-V

Mitel supports the deployment of Mitel servers, both Headquarter (HQ) servers and DVS servers, in Microsoft Hyper-V virtual machines (VMs) running on the Hyper-V 3.0 platform. Support for Mitel servers on Hyper-V 3.0 requires Mitel 14.2 or later.

A failover cluster is a group of two or more computers working together to increase the availability of clustered services or applications. To make virtual machines highly available in Hyper-V environment, you must implement failover clustering on the Hyper-V host servers. Windows Server 2012 Hyper-V comes with a number of new features and improvements for Hyper-V high availability and virtual machine mobility.

Microsoft Windows Server provides 'traditional' host-based clustering with options Hyper-V makes available for failover and load balancing.

Again, by deploying each server as a virtual server on a Hyper-V cluster, we can implement the Failover Clustering (high availability) feature of Microsoft's Hyper-V to achieve a highly available application environment for the HQ and DVS servers.

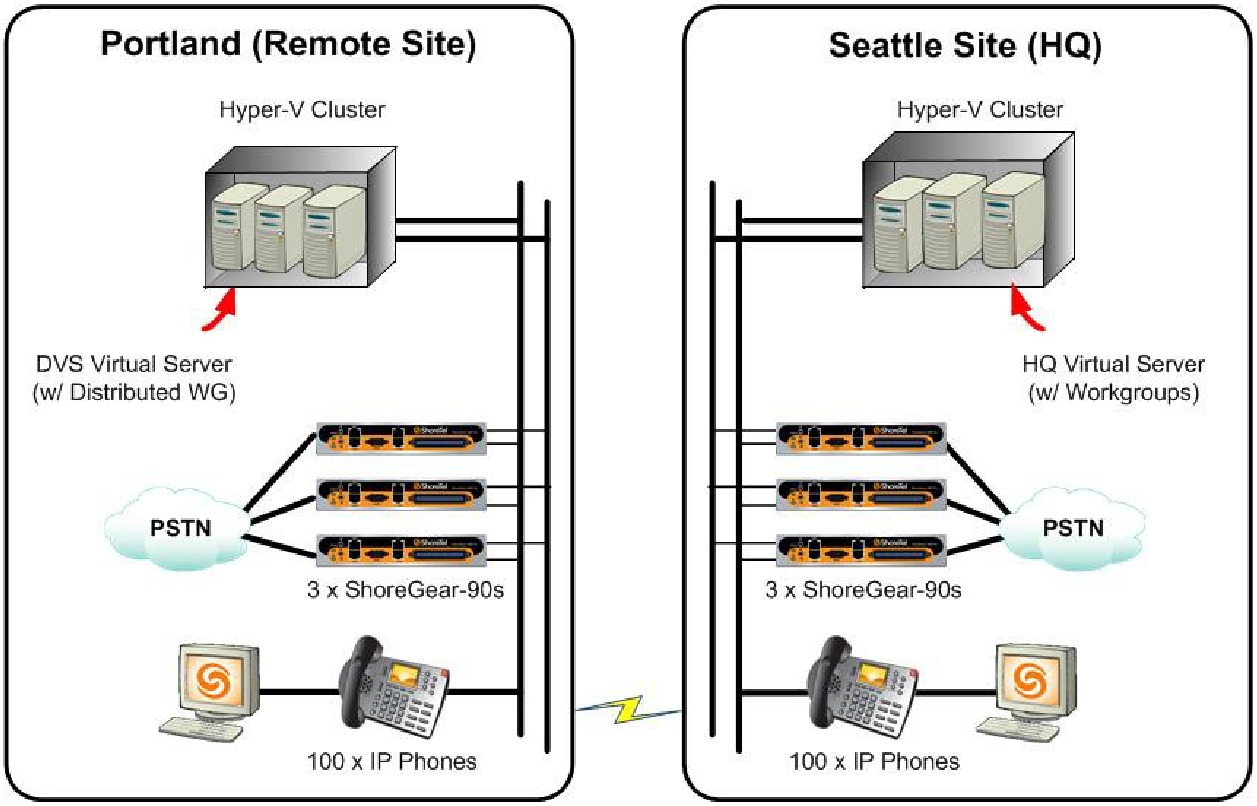

Example: B two-site topology with redundant servers in both locations.

There are some important differences between these two server redundancy solutions:

- Double-Take is officially supported for the HQ server only. VMware HA and Microsoft Hyper-V Failover Cluser are officially supported for both the HQ server and for DVS servers.
- Double-Take supports cross-site failover using different IP addresses. VMware or Hyper- V does not.

Note: See the Mitel Application Note 10259: 'Deploying Mitel on VMware' and Mitel Application Note 14008 'Deploying Mitel Servers under Microsoft Hyper-V 3.0' for complete details.

By making the HQ server redundant (through any method described above), the need to enable the writable, Distributed Database features on a DVS server become less critical. Any hardware or LAN-oriented failure of the primary HQ server will result in only a momentary outage while the secondary HQ server is activated. If your Business Continuity planning includes a server redundancy solution, you might choose to deploy Distributed Workgroups rather than Distributed Databases.

Be aware that a WAN outage will still cause a complete loss of connectivity to the HQ server for a remote site, and a Distributed Database enabled on the DVS server at that remote site would still allow for Call Handling Mode changes (for users) and allow the DVS to be rebooted during the WAN outage.

Recall that the Distributed Database features and the Distributed Workgroup features *cannot* be enabled at the same time on the same single-image Mitel system. You must choose between the two services.

DATABASE CHANGES AND SERVER REDUNDANCY: SUMMARY

In summary, the resiliency and high-availability options for the HQ server, DVS servers and the master configuration database include:

1. Enabling the local 'Call Forward Always' override features by adding the necessary registry change to all Mitel servers.
2. Strategically installing DVS servers and enabling *either*Distributed Workgroups or Distributed Database features on the servers to enhance local remote-site survivability.
3. If adds, moves and changes or running reports is critical, even during an HQ server outage, consider one of the supported server redundancy methods for the HQ server.
4. If Communicator client functionality, voice mail retrieval, or CDR retention iscritical, even during a DVS outage, consider one of the supported server redundancy methods for all servers.
5. Supported Server redundancy solutions include those from 3rd party vendors such as Double-Take,VMware and Microsoft Hyper-V.

FAILURE SCENARIOS

Having discussed single site, multi-site, and Application Services redundancy and fail-over options, let us now take a moment to summarize the actual effects of some specific outages and identify how a well-designed Mitel system will react.

We will use the two-site diagram below for these sample scenarios.

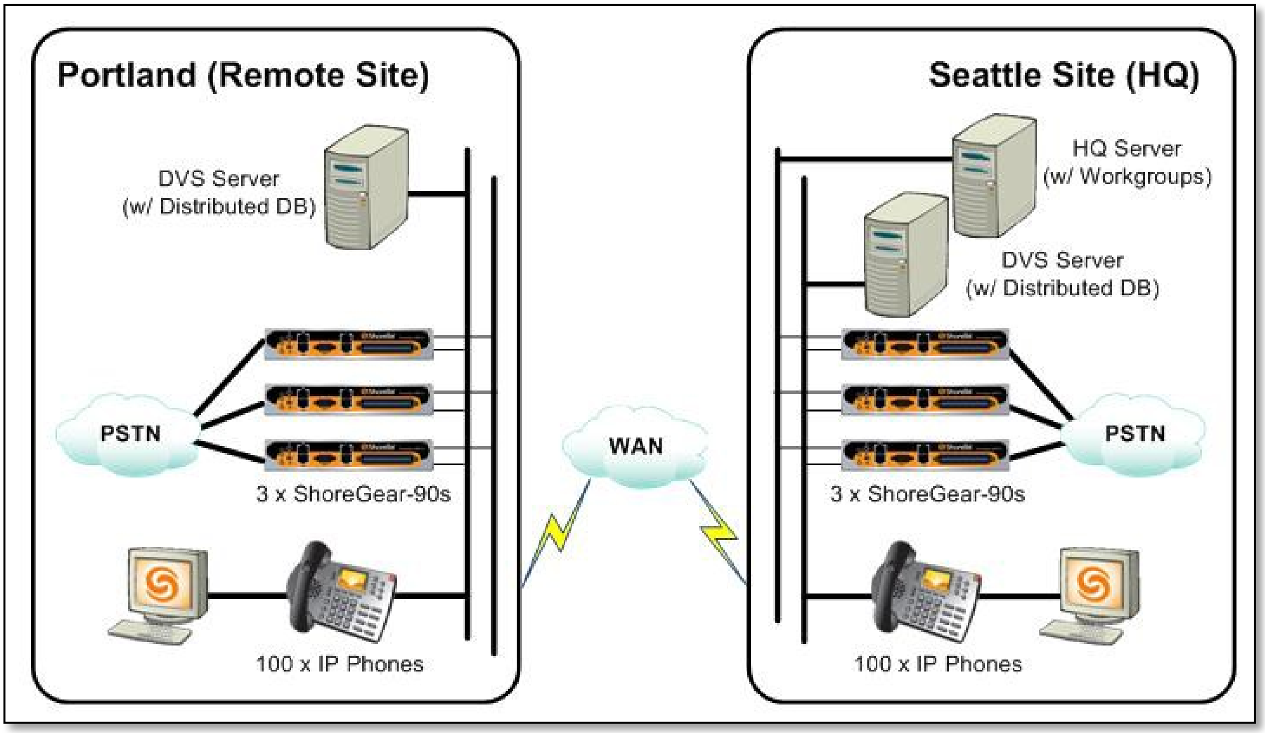

Sample two-site Mitel environment.

TOPOLOGY

Seattle (Headquarters) Site:

- One HQ server (running a Seattle Workgroup and a Portland Workgroup)
- One DVS server with Distributed Database enabled
- All Seattle users have been assigned to the Seattle DVS as their Voice Mail server
- All Seattle switches are managed by the Seattle DVS
- The HQ server is used for Workgroups, Administration and CDR Reporting

Note: No Server redundancy solution has been deployed.

Portland (Remote) Site:

- One DVS server with Distributed Database enabled
- All Portland users have been assigned to the Portland DVS as their Voice Mail server
- All Portland switches are managed by the Portland DVS

FAILURE SCENARIOS

IF THE HQ SERVER GOES OFF-LINE, WHAT *STOPS* WORKING?

- Mitel Connect Director would not be accessible
- No administrative changes can be made
- No reports can be run
- The Master Database on the HQ is unreachable
- No new equipment can be added (IP phones, ShoreGear switches)
- No changes to user options/settings could be made (within Communicator or from an IP Phone)
- Workgroups would stop
- Since Distributed Databases are enabled, distributed Workgroups cannot be enabled. Only the HQ server is running Workgroups
- Workgroups would failover to their explicit Backup Extension such as Hunt Groups
- Hunt Groups would not automatically change via Schedules
  - Hunt Group schedules are initiated by the HQ server

IF THE HQ SERVER GOES OFF-LINE, WHAT *CONTINUES* TO WORK?

- Communicator for Windows, the Web and Mac clients continue to function
  - They are pointed to their local DVS for call control and Voice Mail control
  - Call Recording can still be initiated from Communicator
- IP Phones can be powered off and restarted and will continue to function
  - IP Phones cache their controlling switch, which is still available
- ShoreGear switches can be powered off and restarted and will continue to function
  - ShoreGear switches cache their managing server, which is still available
- Mitel DVS servers with Distributed Database enabled can be powered off and restarted and will continue to function
  - A DVS with Distributed Database enabled stores a copy of the *complete* database, and will continue to function (using the local database copy) even if the HQ server is unreachable

    Note: Rebooting a *standard* DVS server (without DDB enabled) without connectivity to the HQ server will render its services non-functional
- All users can still change their Call Handling mode
  - Both DVS servers have a distributed, writable copy of the database
- Call routing continues to function
  - The Network Call Routing algorithm doesn't rely on servers
- Calls routed to Workgroups would be serviced via backup Hunt Groups

    - Calls would be successfully routed to the same agents using similar call patterns
- All Auto Attendant menus would still work
    - Auto Attendant menus are fully distributed to all DVS servers
- All Account Code Collection would still work
    - Account Code Collection is fully distributed to all DVS servers
- Call Detail Records would be retained
    - For up to 2 hours (by default) before timing out
- Directory buttons on IP Phones continue to work
    - Directory features are serviced by the DVS servers

IF THE WAN BETWEEN SITES FAILS, WHAT *STOPS* WORKING?

- Mitel Connect Director would be accessible only to the Seattle site
- No administrative changes could be made *from Portland*
- No reports could be run *from Portland*
- Any configuration changes made in Seattle would not propagate to Portland
- The Master Database on the HQ is unreachable to the Portland site
- No new equipment can be added *in Portland*
- Workgroups would be inaccessible to calls from Portland
- Calls received in Portland, destined for the Portland Workgroup, would failover to a Hunt Group running on a ShoreGear switch in Portland
- Calls received in Seattle, destined for the Portland Workgroup, would be routed to the 'No Agents' destination of the Portland Workgroup
- Network Call Routing across sites would be unavailable (see page 99)
- Trunks in Portland are not reachable by Seattle users, and vice-versa

IF THE WAN BETWEEN SITES FAILS, WHAT *CONTINUES* TO WORK?

- Communicator for Windows, the Web and Mac clients continue to function
    - They are pointed to their local DVS for call control and Voice Mail control
- IP Phones can be powered off and restarted and will continue to function
    - IP Phones cache their controlling switch, which is still available
- ShoreGear switches can be powered off and restarted and will continue to function
    - ShoreGear switches cache their managing server, which is still available in Portland
- Mitel DVS servers with Distributed Database enabled can be powered off and restarted and will continue to function
    - A DVS with Distributed Database enabled stores a copy of the *complete* database, and will continue to function (using the local database copy) even if the HQ server is unreachable

Note: Rebooting a *standard* DVS server (without DDB enabled) without connectivity to the HQ server will render its services non-functional

- All users can still change their Call Handling mode
    - The Portland DVS has a distributed, writable copy of the database and will allow the changes. The DVS will sync-up with the HQ server when connectivity is restored
- Call routing continues to function
    - The Network Call Routing algorithm doesn't rely on servers.
    - Portland users will use Portland trunks. Seattle users will use Seattle trunks.
- Site-to-Site calls would route over the PSTN
    - PSTN Failover, when configured, allows PSTN trunks to route extension-to- extension calls over PSTN trunks (see page110)
- Calls routed from Portland to the Portland Workgroup would be serviced by the Workgroup's explicit Backup Extension, likely a backup Hunt Group assigned to a Portland-based ShoreGear switch
    - Calls would be successfully routed to the same agents using similar call patterns
- All Auto Attendant menus would still work
    - Auto Attendant menus are fully distributed to all DVS servers and V-switches
- All Account Code Collection would still work

- Account Code Collection is fully distributed to all DVS servers and V-switches
- Voice Mail recording and retrieval would still work for intra-site calls. For inter-site calls, voice mail recordings would be stored on an alternative Application Services end point, but the recordings could not be retrieved until WAN service is restored
- Call Detail Records in Portland would be retained for up to 2 hours (by default) before timing out

- Directory buttons on IP Phones continue to work (driven by the DVS servers)
- Auto Attendant Schedules would still be adhered to (driven by the local servers/V-switches) (see page 118)
- Workgroup Schedules would still be adhered to (driven by the server managing the Workgroup, that is, the HQ server) (see page 118)

IF A DVS FAILS, WHAT *STOPS* WORKING?

- Communicator services for all users assigned to that server would stop working
- Communicator clients assigned to other servers will continue to function
- Call Handling Mode changes for all users assigned to that server would not be changeable
- 'Call Forward Always' modes could be overridden using *72/*73
- CDR records for all switches managed by that server would be lost
- IP Phone Directory buttons for users who's phones/switches are managed by that server would not function
- Schedules assigned for Call Handling Mode changes might be affected.
- See the Schedules section on page 118 for details

- New equipment could not be added to the site with the failed DVS
- New devices (IP phones, ShoreGear switches) at the Portland site should be managed by the Portland DVS. If the Portland DVS is down no new equipment can be added that is managed by the Portland DVS.
- If the Seattle DVS is down, you could add new switches and assign them to be managed by the Seattle HQ server. Then new IP Phones could be added and would be assigned to the HQ-managed switches.

IF A DVS FAILS, WHAT *CONTINUES* TO WORK?

- Auto Attendant menus continue to work
  - Calls will be diverted to an alternative Application Services end point
- Workgroups continue to work
  - Workgroups run off the HQ server
- Mitel Connect Director is still accessible
  - Administrative changes could be made
  - Reports could be run
- The Master Database on the HQ server is up and reachable
  - New equipment can be added (to the unaffected site)
  - All devices, servers and clients can reboot and restart successfully
- The other DVS is still up and running
  - Clients at the other site still have Communicator services
- Voice Mails will be recorded on another Application Services end point but retrieval of the recording will not be possible until the DVS server is restored
- Call routing continues to function
  - The Network Call Routing algorithm doesn't rely on servers (see page 99)

WHAT IF A SHOREGEAR SWITCH FAILS?

- Phones that are orphaned by the failed switch will reconnect to available IP Phone resources at the same site (aka, N+1)
- If a Spare switch is added to the HQ site, it will be automatically assigned to the site in need. Remember there is no additional Mitel hardware cost to add a Mitel Virtual IP Phone Switch as a Spare switch.

WHAT IF A TRUNK CIRCUIT FAILS?

- Outbound calls will be re-routed (by the Network Call Routing process) out the trunks that are still available (see page 99)
- Inbound calls can be re-routed (by the Telco) to enter on other trunks that are still available (see page 103)

CAN I IMPROVE ON THIS DESIGN?

- Yes
    - Consider making your HQ and DVS servers redundant with VMware or Microsoft Hyper-V
    - Consider adding a Spare switch to the Seattle HQ site
    - Consider making the HQ site a ‘standalone’ site

NETWORK CALL ROUTING

Common questions that often arise at this juncture are:

- What if a trunk circuit fails?
- How does the Mitel system work around a PSTN outage?

Thankfully, the Mitel system implements a distributed, network call routing algorithm that intelligently works around nearly any trunk selection issue.

OUTBOUND CALL ROUTING

When you create the components of a distributed Mitel system, you define exactly what each component can do. Let's use the simple, two-site example below.

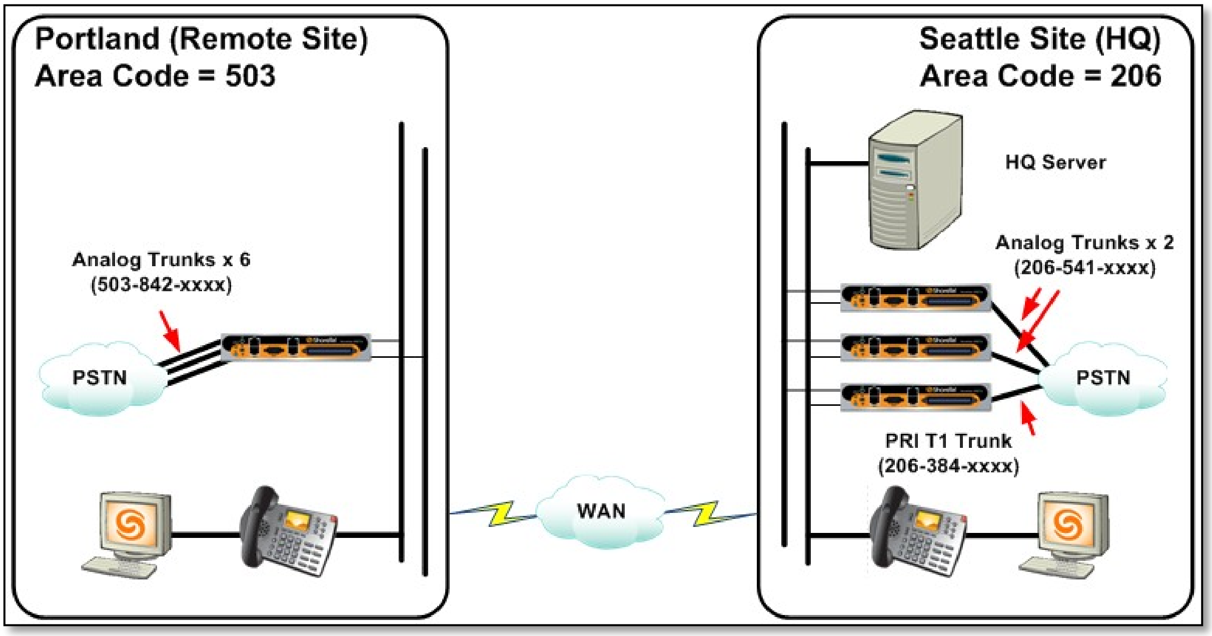

Sample two-site topology

In this diagram we have the following:

- One Headquarter Site: Seattle, WA in area code 206
- One Remote Site: Portland, OR in area code 503
- 3 ShoreGear switches in Seattle
- 3 trunks in Seattle:
- Two analog trunks servicing numbers 206-541-xxxx
- T1 PRI digital trunks servicing numbers 206-384-xxxx
- 1 ShoreGear switch in Portland
- A set of 6 analog loop start trunks in Portland:

- Servicing numbers 503-842-xxxx

When you create the Seattle and Portland sites, you explicitly define what area codes are associated with the sites. More accurately, you define what area codes each site can call as local numbers. For example, Seattle is in the local area code 206, but it is also able to dial some 425 and some 253 area code numbers as local, non-toll numbers. Portland is in the 503 area code, but can also dial some 971 area code numbers as local, non-toll numbers.

When you create the trunk groups (one in Portland and two in Seattle) you define not only what area codes are local, but what prefixes are also local (non-toll) calls.

Note: Creating local prefix lists is made much easier by using online databases such as those found at www.localcallingguide.com. See Knowledge Base Article KB11321: Creating Local Prefixes Lists for more information.

All this information is configured in Mitel Connect Director, added to the HQ server's configuration database and then distributed to each and every DVS server and ShoreGear switch.

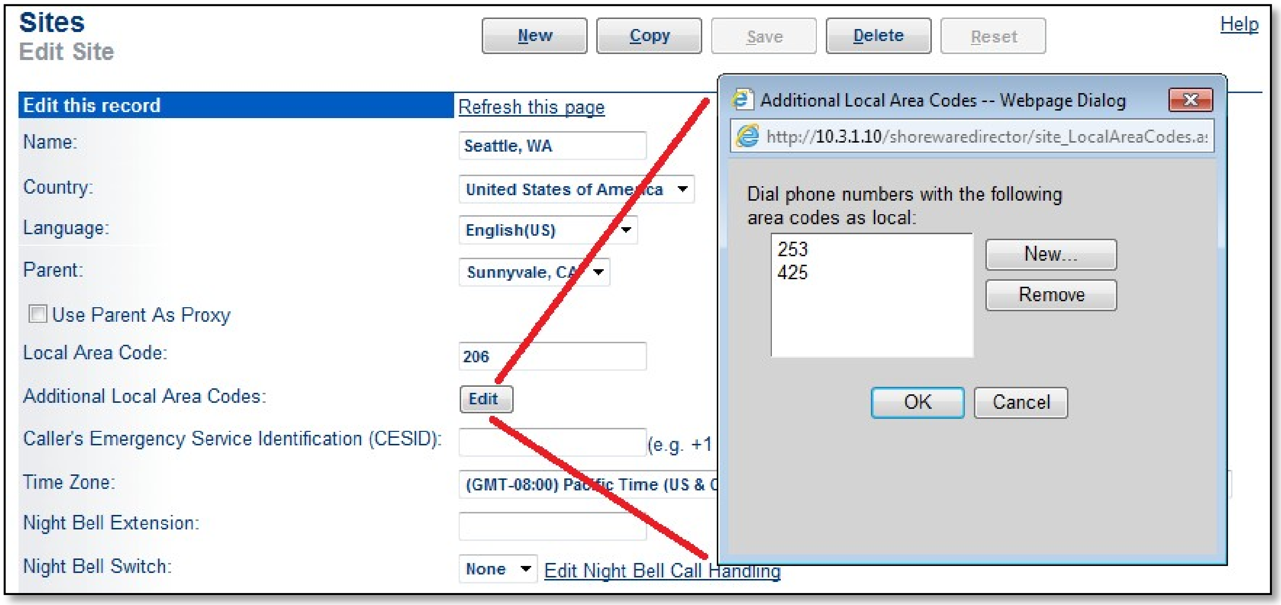

Sites
Edit Site

Edit this record

Name: Seattle, WA
Country: United States of America
Language: English(US)
Parent: Sunnyvale, CA
□ Use Parent As Proxy
Local Area Code: 206
Additional Local Area Codes: Edit
Caller's Emergency Service Identification (CESID): (e.g. +1
Time Zone: (GMT-08:00) Pacific Time (US & C
Night Bell Extension:
Night Bell Switch: None Edit Night Bell Call Handling

Additional Local Area Codes -- Webpage Dialog
http://10.3.1.10/shorewaredirector/site_LocalAreaCodes.a:

Dial phone numbers with the following area codes as local:
253
425

The Seattle site with a Local Area Code and two Additional Local Area Codes.

Every switch now knows about least cost routing and dial plan permissions, and the types of calls that can be placed as non-toll calls out each trunk group at all sites. All the switches exchange details of their local users and trunks (via the LSP protocol, see page 21 and Appendix D) so each switch knows which ShoreGear owns each user and each individual trunk.

CALL SCENARIOS

Example 1: A user in Portland dials a local 503 or 971 number.

Call flow 1: The local ShoreGear switch in Portland will determine that the best trunk group to use is the one it owns itself, and will place the call using a local Portland analog trunk.

Example 2: A user in Seattle dials a local 206, 425, or 253 number.

Call flow 2: The local ShoreGear switch in Seattle determines that the best trunk group to use is the T1/PRI and will place the call using the local PRI trunk.

Example 3: A user in Portland dials a long distance 206, 425, or 253 number.

Call flow 3: The local ShoreGear switch in Portland determines that the best trunk group to use would be one of the two trunk groups in Seattle, since it can place the call as a non-toll call. Digital trunks are preferred over analog trunks, so the Portland switch will contact the Seattle ShoreGear switch that owns the PRI trunk group and will direct the call out the PRI in Seattle.

Example 4: A user in Portland dials a 206 number but all the Seattle T1/PRI trunks are in use.

Call flow 4: The local ShoreGear switch in Portland reaches out to the T1 switch in Seattle but the T1 switch indicates it has no available trunks/channels. The Portland switch then communicates with a ShoreGear switch that owns an analog trunk in Seattle. The trunk is free and the call is placed.

Example 5: A user in Portland dials a 206 number but the WAN is down between the two sites.

Call flow 5: The local ShoreGear switch in Portland would have preferred to use any of the Seattle trunks, but none of the Seattle-based ShoreGear switches are reachable and all Seattle trunks have been marked as Inactive in the LSP table in the Portland switch. The Portland ShoreGear places the call as a 1+10-digit long distance call out the local Portland trunks.

Example 6: A user in Portland dials a number in the 212 area code (New York).

Call flow 6: The call will be a long distance call no matter what trunk group is used. The local ShoreGear chooses to place the call out the local Portland trunks saving WAN bandwidth. If the local Portland trunks had all been down or busy, the call would have been routed out a Seattle Trunk Group.

As these call flow examples show, every ShoreGear switch uses its full knowledge of every trunk group and every permission setting, along with its direct communication relationship with all other ShoreGear switches to place each and every call out the most appropriate trunk at the moment the call is placed by the user. No server intervention is required.

To route a call, each ShoreGear switch starts with the complete list of enterprise-wide trunk groups. The list is winnowed down by selecting just those trunk groups that are reachable, that match the user-dialed trunk access code (e.g. ‘9’), that the user has permission to use, and those trunk groups that have been given permission to place such a call (for instance, allowing long distance but disallowing international calls).

Armed with the list of all trunk groups that would be usable to place the call, the ShoreGear switch selects the *best* trunk group according to the following criteria:

- Prefer a non-toll call over a toll call
- Prefer a local-site’s trunk group over a remote-site’s trunk group
- Prefer SIP trunks over digital trunks (T1/E1/BRI)
- Prefer digital trunks over analog trunks

Note that the Mitel system’s network call routing process is not based on the Site Tree Hierarchy. Nor is it based on Application Servers. It is based solely on User Group permissions, Trunk Group capabilities, and ShoreGear connectivity and availability. Using all the trunk groups across the entire organization, each call will be placed out the most appropriate and available trunk.

If a particular ShoreGear switch is off-line or unreachable, its trunks won’t be selected. If a WAN link is down, any ShoreGear switches that cannot be reached will be marked as Inactive and removed from the selection process.

INBOUND CALL ROUTING

It has been described above that you can make your trunking more resilient by:

- Splitting the trunks of a single circuit into trunks on multiple circuits

- Dividing your multiple trunks across multiple ShoreGear switches
- Decentralizing your trunks by distributing them to multiple sites
- Creating carrier diversity by deploying trunks from two or more discrete carriers

These steps help make outbound calling far more resilient and tolerant to faults, outages and accidents. But what about inbound calling? How do you design for the most survivable delivery of inbound calls?

ANALOG TRUNKS

If a site has a quantity of analog trunks from a PSTN provider, those trunks are usually grouped in an inbound Telco Hunt Group. For example, if you have four analog lines at one location, the Telco might configure them as follows:

- 206-555-6111: Primary line; published in the phone book; if busy, roll call to 206-555- 6112
- 206-555-6112: Second line; *not* published; if busy, roll inbound call to206-555-6113
- 206-555-6113: Third line; *not* published; if busy, roll inbound call to206-555-6114
- 206-555-6114: Last line; *not* published; if busy, *play 'busy signal' to caller*

If all trunks are operating correctly and calls are active on all four trunks (inbound or outbound) the next inbound caller will receive a busy signal from the telephone company. And no more outbound calls can be placed. An attempt at placing an outbound call will result in a busy signal played from the Mitel system. (In such cases, an entry will be posted in the Mitel server logs. You should respond by ordering more trunks.)

But what if the lines are not busy, but instead have been disconnected through a construction mishap on your street? As before – no calls can be received or placed. To remedy this you need an alternative trunk entry point.

Options to consider are:

- Forward the last line (206-555-6114) to a cell phone number
- You would only be able to take a few calls but that might be better than none
- Install and deploy alternative trunks
- Some PSTN providers offer specially-priced emergency trunks for a lower monthly fee with a higher per-use cost; perfect for occasional or rare use as a disaster recovery plan
- Forward the last line (206-555-6114) to the alternative trunk's pilot number
- Plan carefully: If the construction worker digs up both your primary Telco lines and your alternative trunk lines (at the same site) with the backhoe – your alternative trunks will not be functional either
- Deploy trunks at another site
- If the quantity of trunks is adequate, split them up and move half to another site
- If the other site is within the same Telco service area they should be able to re-design the Hunt Group so that calls hunt' across both groups of trunks

DIGITAL TRUNKS

If your primary trunk group is a PRI (T1/E1) or BRI circuit, inbound calls will 'hunt' from channel to channel instead of from line to line. On a North American T1/PRI with 23 bearer (or 'B') channels, it is common for the Telco to 'hunt' starting at channel 1 and proceed 'down' towards channel 23. The Mitel system, by default, will place outbound calls starting on channel 23 and hunt 'up' towards channel 1.

If all channels are busy, the next inbound call will receive a busy signal from the Telco. And the next outbound call will receive a busy signal from the Mitel system.

To provide redundancy you need an alternative trunk connected to the Mitel system. The alternative trunks can be delivered on another digital circuit (T1/BRI), via SIP or via analog trunks. The additional trunks can be:

- Connected to the same switch (e.g. a ShoreGear-220T1a)
- Located at the same site but connected to a different ShoreGear switches, or
- Located at another site

For the best survivability, it is recommended to deploy adequate trunking at two different sites.

For example, it is very common for a local municipality to deploy a Mitel system with one (or more) PRI trunks at the City Hall location and deploy another PRI at their largest Police Department site. Mitel will use *all* trunks at *all* sites according to the normal Network Call Routing mechanism (see page 99).

For inbound redundancy you need to work with your trunk provider(s) to create either an automatic, or a manually initiated, failover process in case the primary trunks are full or out of service. If the secondary trunks are from the same provider, you likely will have options to include the secondary trunks in the Telco's inbound PRI 'Hunt Group' and to present the same inbound DID/DDI numbers on the secondary trunk (or span).

In such a scenario, any inbound direct inward dial (DID/DDI) numbers that are assigned to one 'span' will 'hunt' and route into the Mitel PBX from the Telco on *any* of the configured spans – even those at the other site. This scenario is most desirable because the originally dialed DID is still received whether the call comes in on the City Hall trunks or the Police Station trunks. The programming necessary on the Mitel system to accommodate these cross-site inbound calls requires that you:

- Add the necessary DID range to the primary Trunk Group at City Hall
- Add an identical DID range to the secondary Trunk Group at the Police Station
- Assign DID numbers from the primary Trunk Group to your users/entities

Inbound DID calls will be routed to the correct destination regardless of which trunk group they are received on.

Similarly, if DNIS Mapping is used, you should replicate the DNIS map entries on *both*trunk groups.

If the two sites are *not* serviced by the same trunk provider (e.g. one is a PRI and one isSIP; or the two sites are in different locations; or are serviced by different Telco providers) you might need to create a 'Forward when Busy' route from the first Telco Hunt Group to the pilot number of the second trunk group.

For example, say you have PRI's in Seattle and backup PRI's in Phoenix, AZ. You would need to work with your Seattle trunk provider to route all calls to a Phoenix pilot number if the Seattle channels are in a 'busy/out of service' state. Conversely, you would need to work with your Phoenix trunk provider to route calls to a Seattle pilot number if the Phoenix trunks were unavailable.

This scenario is less desirable because the originally dialed DID might be discarded when the call is forwarded to the alternative number. The diverted calls would need to be answered by an Automated Attendant menu or a receptionist at the alternative site.

Note:

- Routing calls across long distance boundaries can invoke long distance charges.
- You should discuss costs and capabilities with your PSTN trunk providers

PARENT AS PROXY

As stated earlier, the trunk selection algorithm uses the following decision tree:

- Prefer a non-toll call over a long distance call
- Prefer a local-site trunk group over a remote-site trunk group
- Prefer SIP trunks over digital trunks (T1/E1)
- Prefer digital trunks over analog trunks

It is also important to realize that the trunk selection process and call routing mechanism is used only for 'routable' numbers – those numbers that conform to a standard PSTN format.

Non-routable numbers *do not* follow the network call routing rules and *must* be placed using a local trunk. Non-routable numbers include emergency numbers (e.g. 911), local informational number such as n11 (e.g. 411, 511, 811, and so on.) and operator assisted calls.

But what if a site has no trunks? Or what if those trunks are restricted to 'local only' and do not allow 911 calls?

Consider a manufacturing plant in Seattle, WA with five buildings: A, B, C, D and E. All five buildings are built on the same city block and Telco services enter only into Building 'A'.

You could define the entire block as one 'Seattle, WA' Headquarters site, but you wish to add overhead paging in each building. So, you must split up the organization into building- specific sites in the Mitel site hierarchy: building 'A' becomes the headquarters site (named 'Seattle, WA') and buildings B, C, D and E become child sites of the headquarters site.

Having done so, with no configuration changes from the default, if someone in Building 'B' dials 911 they will get a fast-busy signal from the Mitel system because there are no local trunks at that site and 911 calls *must* be placed out of local trunks.

To accommodate this situation, you should check the 'Use Parent As Proxy' check box on the site edit page for site's B, C, D and E.

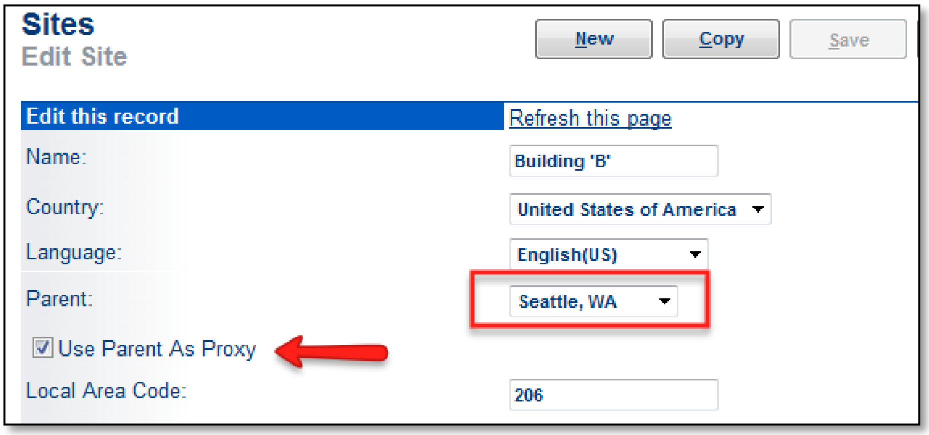

Sites
Edit Site

Edit this record

Name:	Building 'B'
Country:	United States of America
Language:	English(US)
Parent:	Seattle, WA
☑ Use Parent As Proxy	
Local Area Code:	206

Site "Building 'B'" uses its parent site, 'Seattle, WA', as a Proxy site.

Checking ‘Use Parent As Proxy’ alters the trunk selection algorithm and will treat all trunks at the parent site as if they were local to the original site. Now, a 911 call will be successfully placed out the trunks at the ‘Seattle, WA’ (or Bldg. ‘A’) site.

Note: It is *critical* that you manage Public Safety Location Databases (PSAP ALI) so that the call, when received by the emergency dispatcher, displays address information specific to Building ‘B’ instead of the parent site. See the resources section at the end of this document (page 143) for more on 911 and emergency services.

You should *not* to use ‘Parent As Proxy’ unless the two sites are geographically next to each other. If the emergency vehicles accidentally arrive at the main building, the actual emergency should be in an adjacent building.

Do not use ‘Parent as Proxy’ for geographically disperse locations. Always deploy local trunks to locations that are geographically dispersed.

For example, if you have a multi-site school district with schools and administrative buildings spread throughout a small city, you should *not* use parent as proxy. If a caller at an elementary school dials 911, and parent as proxy is checked, the call could be placed out a trunk at the administrative building (the parent site). This might inadvertently route emergency vehicles to the wrong location. Instead, ensure there are ample trunks at each local site to place emergency calls.

CALL COST PROMOTION

By design, call routing is performed to conserve costs and to conserve resources. This includes avoiding long distance when possible and avoiding the use of WAN bandwidth when possible.

A situation arises when a local user places a local call but all local trunk resources are in use or unavailable. In such cases, the default behavior is to disallow the call. Technically, elevating the call from a ’local, non-toll call’ to a ‘remote-site, long distance, WAN bandwidth-using call’ is an elevation of ’cost’ that is too high and is therefore prohibited.

For example, consider a user in Portland who is ordering lunch from the local pizza parlor. If the local Portland trunks are all busy, or out of service, the call will be disallowed and the Mitel user will hear a busy signal from the Mitel system.

This might seem trivial when you consider the minimal costs of long distance between states in the United States but consider the following two scenarios:

1. What if the situation was reversed and all the Seattle trunks were out of service? Would you want all outbound Seattle user calls to consume the scarce 6 analog trunks in Portland?
2. What if the alternative site wasn’t Seattle but was Hamburg, Germany?

   If WAN bandwidth is minimal or very expensive, if your trunk resources are scarce, or if long distance costs (e.g. across the Atlantic) are cost prohibitive, then you do *not* want to allow such cost promotion of calls. It is for these cases that ‘Call Cost Promotion’ is inhibited by default.

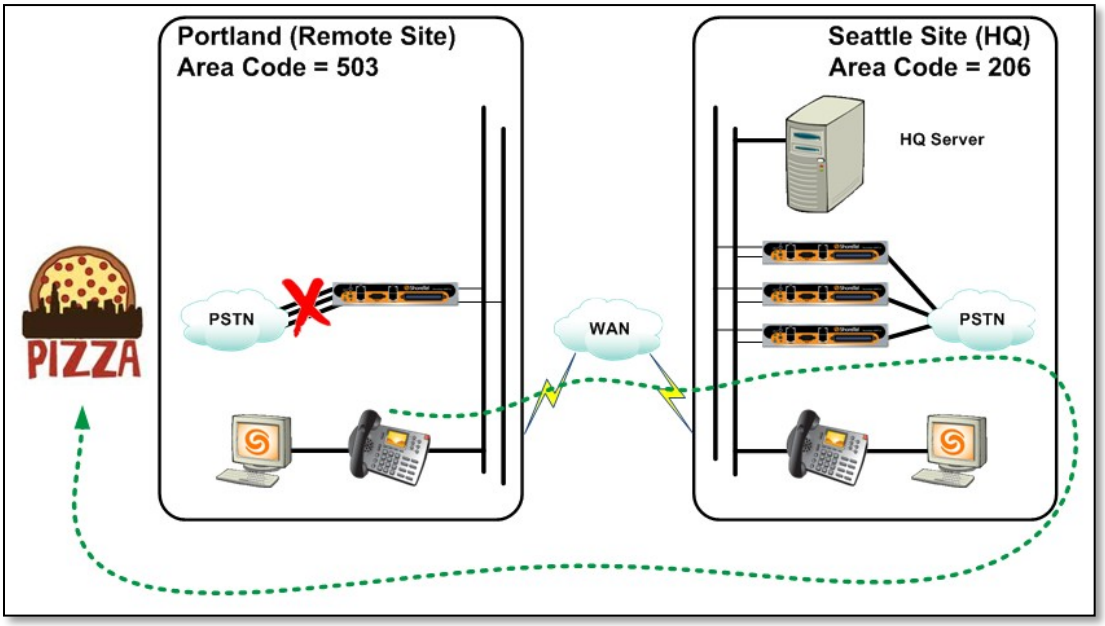

The Portland site has been customized to allow local numbers to be dialed out another sites trunk as a long distance number

Call ‘Cost Promotion Inhibition’ is defined, essentially, as ‘don’t make a long distance call out of a local call’.

You can alter this behavior, and allow a local call to be promoted and placed as a long distance call out an alternative site if desired. This change is made on a site-by-site basis:

You could enable this for calls placed by users in the Portland site but not for calls placed by users in the Seattle site.

The change is made by modifying the custom dial plan parameters for the site that is *placing* the call.

To enable the Portland users to dial a Portland number out a Seattle trunk, you would modify the Portland site’s Custom Dial Plan.

Note: Editing a site’s custom dial plan should only be performed on the advice of the Mitel Technical Assistance Center (TAC). For details on configuring the necessary parameters to enable call cost promotion contact your local Mitel Sales Engineer or the Mitel TAC.

Note: Call cost promotion is enabled based on the calling party’s site, not the destination trunk group. If you enable call cost promotion you will allow the local call to be placed as a long distance call out *any* trunk group at *any* site.

This might be useful if your alternative site is within your country (like Seattle) but undesirable if you have sites in other countries (like Hamburg, Germany). There is no way to restrict a promoted call. If all Seattle trunks are busy, the call will be placed using a German trunk to order pizza in Portland.

PSTN FAILOVER

PSTN Failover is a Mitel feature that enables site-to-site, extension-to-extension dialing, even if there is no IP bandwidth available between sites.

By default, if a user at site ‘A’ dials a user at site ‘B’, the Source ShoreGear switch will contact the Destination ShoreGear switch to establish the call. If the WAN is down, or there is already too much Voice over IP/RTP traffic on the WAN segment, the call would be denied and the caller would be routed to the destination user’s forwarding destination, likely to a local Voice Mail end point.

To counter this, you can enable ‘PSTN Failover’ for the dialing user and specify a dialable PSTN number for the destination user.

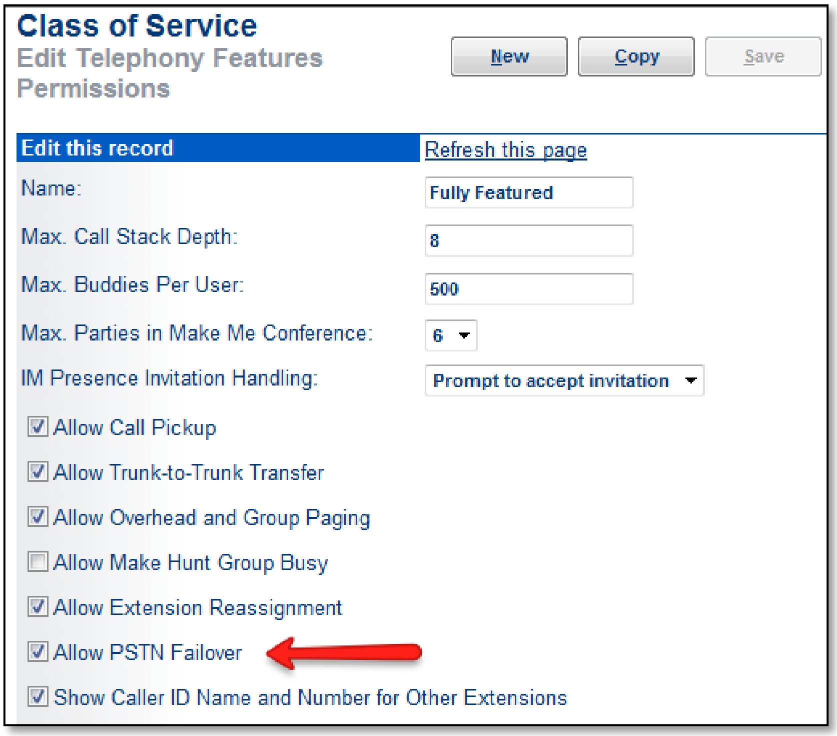

Class of Service
Edit Telephony Features Permissions

Edit this record

Name: Fully Featured

Max. Call Stack Depth: 8

Max. Buddies Per User: 500

Max. Parties in Make Me Conference: 6

IM Presence Invitation Handling: Prompt to accept invitation

☑ Allow Call Pickup

☑ Allow Trunk-to-Trunk Transfer

☑ Allow Overhead and Group Paging

☐ Allow Make Hunt Group Busy

☑ Allow Extension Reassignment

☑ Allow PSTN Failover

☑ Show Caller ID Name and Number for Other Extensions

Enabling ‘PSTN Failover’ in the Telephony Class of Service applied to a User Group.

The dialing user must have this permission for PSTN Failover to be invoked.

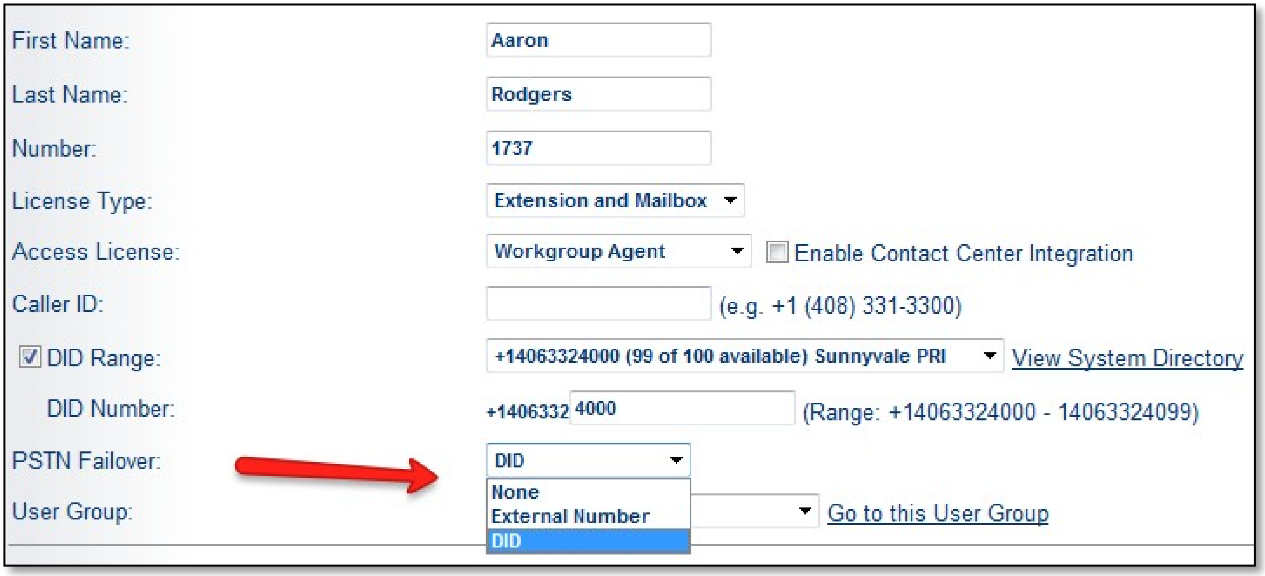

First Name: Aaron

Last Name: Rodgers

Number: 1737

License Type: Extension and Mailbox

Access License: Workgroup Agent ☐ Enable Contact Center Integration

Caller ID: (e.g. +1 (408) 331-3300)

☑ DID Range: +14063324000 (99 of 100 available) Sunnyvale PRI

DID Number: +1406332 4000 (Range: +14063324000 - 14063324099)

PSTN Failover: DID

User Group:

The destination user must have a dialable number assigned as their ‘PSTN Failover’ number

If the dialing user has the ‘PSTN Failover’ permission assigned to their Telephony Class of Service as part of their User Group assignment, and the extension they are calling has a valid dialable number set for their PSTN Failover number, then the Mitel system will use a trunk to place the PSTN call to the destination user’s PSTN Failover number. This re- routing will occur if site-to-site IP connectivity is down or full (as determined by the two sites’ Admission Control Bandwidth settings).

By enabling PSTN failover you are acknowledging that some calls are important enough to use a trunk on both the source site and the destination site and, potentially incur long distance charges. You enable this capability user-by-user via User Group settings. And you select which users are important enough to receive these re-routed calls by adding a PSTN Failover number on their User Edit page.

Any destination user that does *not* have a PSTN Failover number specific in their User Edit page will not invoke PSTN Failover. Any dialing user that does *not* have the PSTN Failover permission will not invoke PSTN Failover.

Note: The PSTN Failover number assigned to a user should be the destination user’s DID whenever possible. If the destination user does not have a DID, an alternative PSTN number can be entered. Realize that the dialing user’s call will be unexpectedly answered by a destination other than the originally dialed destination. An Auto Attendant main menu or a receptionist number would be appropriate alternative numbers for users without DIDs.

Note: PSTN Failover is invoked only for user-initiated calls between sites. External calls that come in on a trunk at one site and are routed to a user at another site will *not* invoke PSTN Failover. Calls routed by Hunt Groups and Workgroups will not invoke PSTN failover. Transfers by the Voice Mail system or an Automated Attendant menu will *not* invoke PSTN Failover.

However if a trunk-based call is answered by a user (not a system extension) and the user does a blind or consultative transfer (or conference) with a user at another site, PSTN failover *will* be invoked.

SITE TREE HIERARCHY

When a ShoreGear switch needs to route a call to a trunk it uses the Network Call Routing mechanism (see page 99).

When a ShoreGear switch needs a server-based resource, it goes directly to the server that owns the resource (e.g. Workgroups) or to its primary (or secondary) Application Services end point for the resource (e.g. Auto Attendant or Voice Mail).

But, there are some services for which the requesting ShoreGear switch will ‘hunt up the tree’ to select the ‘closest’ resource. For those resources, the Mitel Site Tree Hierarchy is used.

The Site Tree Hierarchy represents an ordered list, and relative positioning, of all the sites as configured within Mitel Connect Director.

When you create a site you must specify which site will act as the ‘parent’ of the site you are creating. There is only one site at the root of the tree. This root site is originally named ‘Headquarters’ but can be renamed as desired. All other sites are children or grandchildren (or great grandchildren, and so on.) of the root (HQ) site.

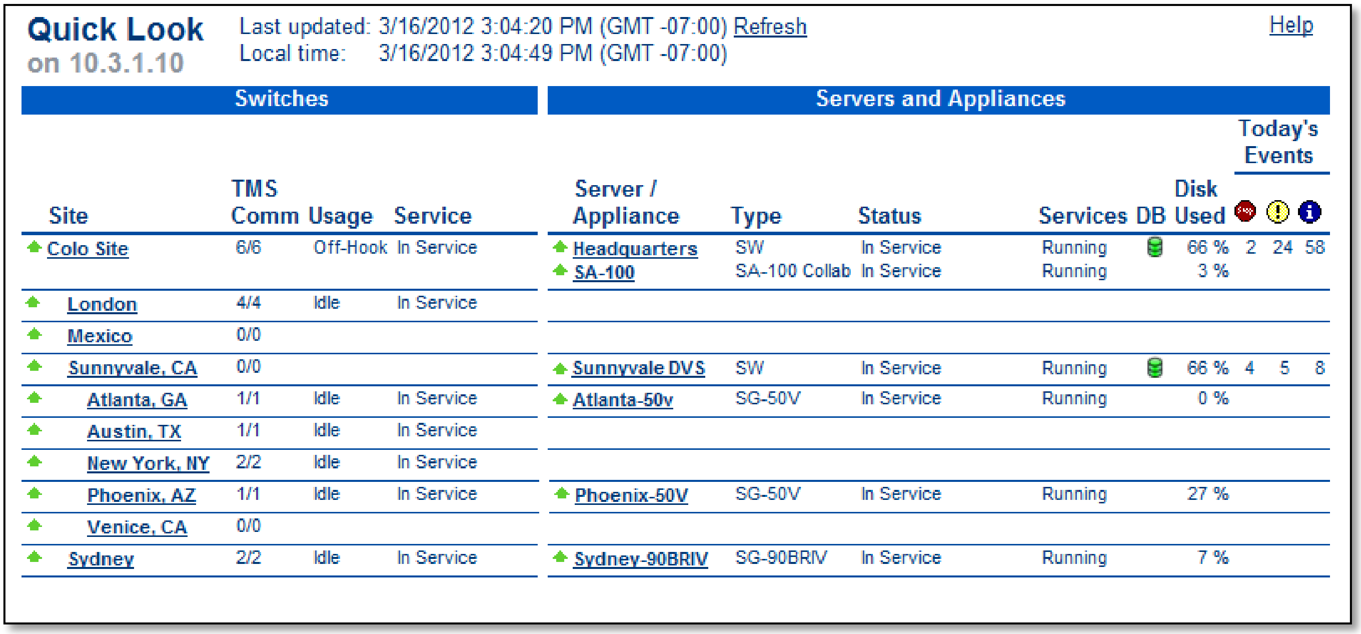

Quick Look on 10.3.1.10

Last updated: 3/16/2012 3:04:20 PM (GMT -07:00) Refresh
Local time: 3/16/2012 3:04:49 PM (GMT -07:00)

Help

Switches				Servers and Appliances								
										Today's Events		
Site	TMS Comm	Usage	Service	Server / Appliance	Type	Status	Services	DB	Disk Used	⛔	❗	ℹ
Colo Site	6/6	Off-Hook	In Service	Headquarters SA-100	SW SA-100 Collab	In Service In Service	Running Running	🛢	66 % 3 %	2	24	58
London	4/4	Idle	In Service									
Mexico	0/0											
Sunnyvale, CA	0/0			Sunnyvale DVS	SW	In Service	Running	🛢	66 %	4	5	8
Atlanta, GA	1/1	Idle	In Service	Atlanta-50v	SG-50V	In Service	Running		0 %			
Austin, TX	1/1	Idle	In Service									
New York, NY	2/2	Idle	In Service									
Phoenix, AZ	1/1	Idle	In Service	Phoenix-50V	SG-50V	In Service	Running		27 %			
Venice, CA	0/0											
Sydney	2/2	Idle	In Service	Sydney-90BRIV	SG-90BRIV	In Service	Running		7 %			

The Site Tree Hierarchy as seen in Quick Look in Mitel Connect Director.

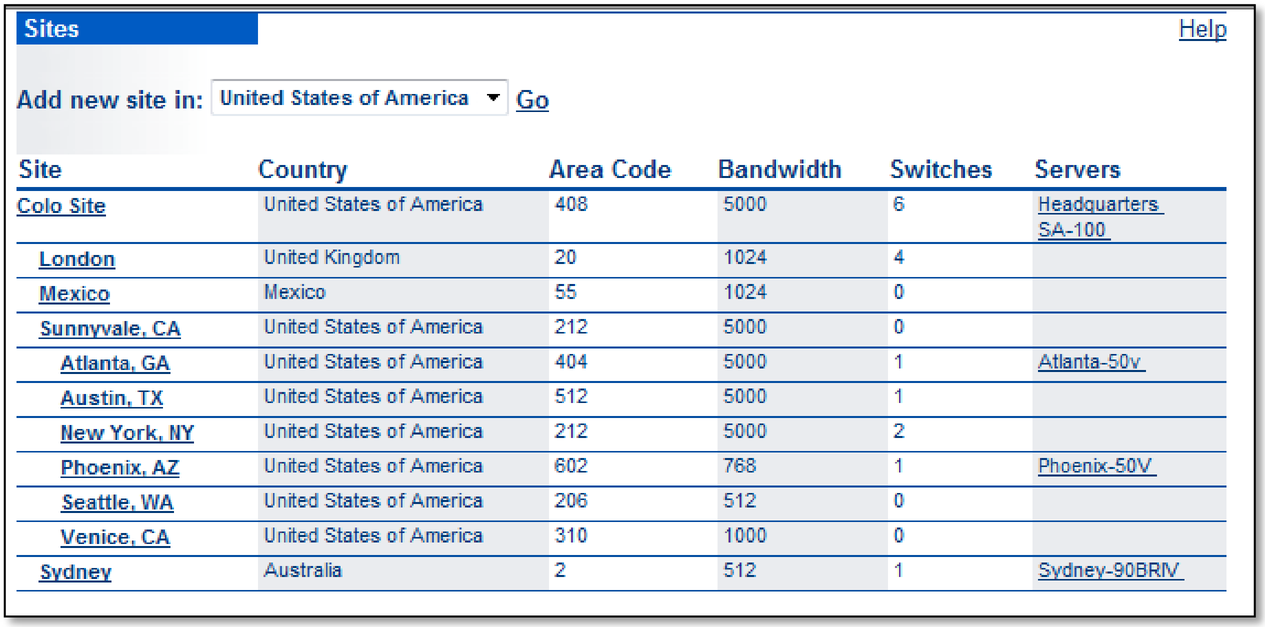

Site	Country	Area Code	Bandwidth	Switches	Servers
Colo Site	United States of America	408	5000	6	Headquarters SA-100
London	United Kingdom	20	1024	4	
Mexico	Mexico	55	1024	0	
Sunnyvale, CA	United States of America	212	5000	0	
Atlanta, GA	United States of America	404	5000	1	Atlanta-50v
Austin, TX	United States of America	512	5000	1	
New York, NY	United States of America	212	5000	2	
Phoenix, AZ	United States of America	602	768	1	Phoenix-50V
Seattle, WA	United States of America	206	512	0	
Venice, CA	United States of America	310	1000	0	
Sydney	Australia	2	512	1	Sydney-90BRIV

The Site Tree Hierarchy as configured using ‘Administration > Sites’ in Mitel Connect Director.

The Site Tree Hierarchy is used when a request for the following types of resources is made:

- Requesting ‘Make-Me’ Conferencing ports (that is, 4, 5, 6-party conferencing)
- ‘Spare’ Switch assignment
- Selection of DRS Servers
- Locating Conference Bridge resources on a Service Appliance (for example, the SA-100)

Note:

- The site tree is not used when making a trunk selection. That uses the Network Call Routing algorithm. (see page 99).
- The site tree is *not* used when accessing an Auto Attendant or the Voice Mail system. Those calls are routed to the primary and secondary Application Services end points. The selection of the two ‘closest’ Application Services end points is related to the site tree hierarchy). (See page 26).
- The site tree is *not* used when a needed resource is unavailable and has an explicit Backup Extension (like Workgroups or Hunt Groups). (See page 46).
- Resources that are hunted for ‘up the tree’ will only be found at the site that is looking for the resource and its parent/grandparent sites. Resources will never be found, or used, at sibling (same level) sites or children sites.

4, 5, AND 6-PARTY CONFERENCING

When a user attempts to join a 4th party to a conference, the ShoreGear switch managing that user’s phone will search ‘up the tree’ looking for a ShoreGear with adequate ‘Make Me’ conference port resources. The Source switch will look in its local site and its parent site. It will not look farther up the tree.

Note: Ad-hoc ‘Make Me’ conferencing is the only ‘hunt up the tree’ mechanism that has such a depth restriction.

As an example, consider the following scenario when planning the deployment of 4, 5, and 6-party conference ports:

Your company wants to have the ability to have up to four simultaneous 6-party conference calls. This requires that 24 ‘Make Me’ conference ports be assigned on ShoreGear switches. A ShoreGear-24a is added to your topology to supply the 24 conference ports. But where should that ShoreGear-24a be located?

If it is placed at the top of the tree (at the Colo Site, in our sample tree hierarchy) then a request by a user assigned to a switch at London, Mexico, Sunnyvale or Sydney would be able to use the resources. But users at the grandchildren sites (below Sunnyvale) would not. Make Me conference resource selection only uses the local site and the parent site: Make Me conference resources are not used at grandparent sites or above.

Placing the ShoreGear-24a at the Sunnyvale site allows users in Sunnyvale and all its children-sites to initiate a 4, 5 or 6-party conference but would prevent the users in London, Mexico, Sydney and the Colo Site from doing so.

Also, consider your WAN bandwidth usage. What if all 6 parties on the conference call are users (or trunks) located at a child site, like the Phoenix, AZ site? All six calls, that previously were local to the Phoenix, AZ site, will be re-directed over the WAN to the ShoreGear-24a with the conference resources located at the Sunnyvale site. This might be an inefficient use of bandwidth. You might want to consider re-locating the ShoreGear-24a to the Phoenix site to save WAN resources.

But, if the 24a is located at the Phoenix site, what happens if a user in New York, NY attempts to add a 4th party to a call? Their local ShoreGear switch, in need of Make Me conference resources, will hunt ‘up the tree’ looking first at the local New York site, then the parent Sunnyvale site. No conference resources will be found, and the addition of the fourth call will fail.

Moving the conference resources too far ‘down’ the tree has prevented it from being used by sibling sites and parent sites. Moving the conference resources too far ‘up’ the tree prevents grandchildren sites from access the resources.

You must compromise.

Consider the following when deciding on the site placement of Make Me Conference Port resources:

- Who will be the primary users of the resources? Will they be primarily from one location? If so, it might make sense to locate the resources at their local site even though this excludes the use of the resources by sibling and parent sites.
- If the users will be spread evenly across all sites, then it makes more sense to locate the resources at, or near, the top of the tree, even though this means using more WAN bandwidth for the calls; and possibly excluding grandchildren sites.
- Consider splitting up the conference resources. For example placing a ShoreGear-90 with 12 ports of conferencing in Phoenix and another ShoreGear-90 with 12 ports of conferencing at the Colo Site. This provides local resources for the heavy conference users in Phoenix and provides 12 additional ports for general use by many other sites.
- Consider using Virtual Mitel IP Phone Switches instead, these virtual switches provide 6 Make Me conference resources per 100 users, which means a 1000 phone virtual switch will have 60 built-in Make Me conference resources, 30 for a 500 phone virtual switch and 15 for a 250 phone virtual switch.

SPARE SWITCH SELECTION

The Site Tree Hierarchy is also used when a Spare switch needs to be assigned to a site in need of IP Phone resources.

Note: See the section 'Remote Site Survivability – Redundant Call Control' on page 73 for more information on Spare switches.

When a site runs out of IP Phone resources (either due to the loss of a ShoreGear switch, or due to the addition of more IP Phones than there are IP Phone resources at that site), and 'IP Phone Failover' is enabled, the site that is in need of IP Phone resources will hunt 'up the tree' looking for a site that has an available Spare switch.

The first Spare switch that is found will be assigned, in its entirety, to the site in need. If no spare switch is found, the extra (or orphaned) IP Phones will not function and will display 'Requesting Service'.

Note: Spare switch selection is not limited to specific levels of the Mitel Tree Hierarchy. Any child level will hunt 'up the tree,' all the way to the root (HQ site), until an available 'spare' switch is found. If the 'spare' switch found is a Virtual IP Phone Switch you might now be out of license compliance. If you are, be certain to rectify this issue prior to 45 days otherwise you will get locked out of administering the Mitel system, until you are back-in license compliance.

SELECTION OF DRS SERVERS

Each ShoreGear switch will locate and select up to two Mitel servers to use when the Distributed Routing Service is enabled. Selection of the two DRS servers is not limited to specific levels of the Mitel Tree Hierarchy. Any child level will hunt 'up the tree,' all the way to the root (HQ site), until the two 'nearest' servers are found.

Note: See the section on the Distributed Routing Service on page 134 for more details.

SCHEDULES

Many Mitel entities can be assigned schedules. Schedules are created in Mitel Connect Director and assigned to a specific entity to alter call routing and call handling mechanisms based on time of day, day of week and with customized hourly and daily overrides.

There are three types of schedules in the Mitel system:

- On-Hours/Off-Hours schedules
- Holiday Schedules
- Custom Schedules

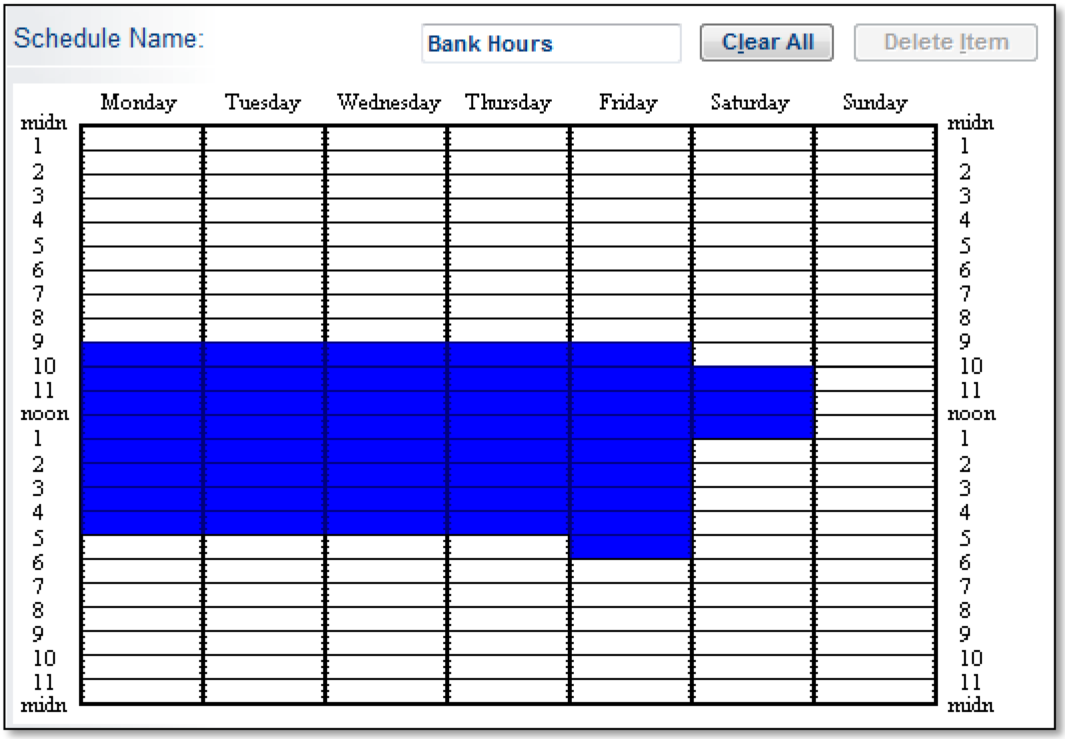

Schedule Name: Bank Hours

Sample of an On-Hours schedule called 'Bank Hours'.

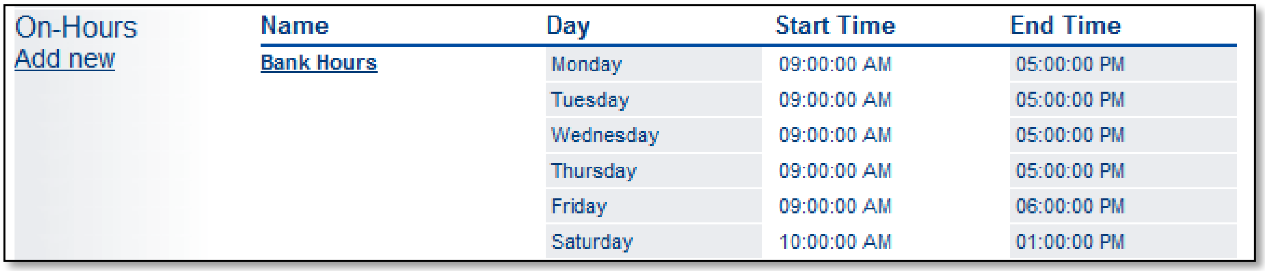

On-Hours
Add new

Name	Day	Start Time	End Time
Bank Hours	Monday	09:00:00 AM	05:00:00 PM
	Tuesday	09:00:00 AM	05:00:00 PM
	Wednesday	09:00:00 AM	05:00:00 PM
	Thursday	09:00:00 AM	05:00:00 PM
	Friday	09:00:00 AM	06:00:00 PM
	Saturday	10:00:00 AM	01:00:00 PM

The same 'Bank Hours' schedule displayed in table format.

On-Hours/Off-Hours schedules define a week's worth of time that covers one or more start and stop times for each day of the week.

Holiday Schedules specify complete days that override On-Hours/Off-Hours schedules.

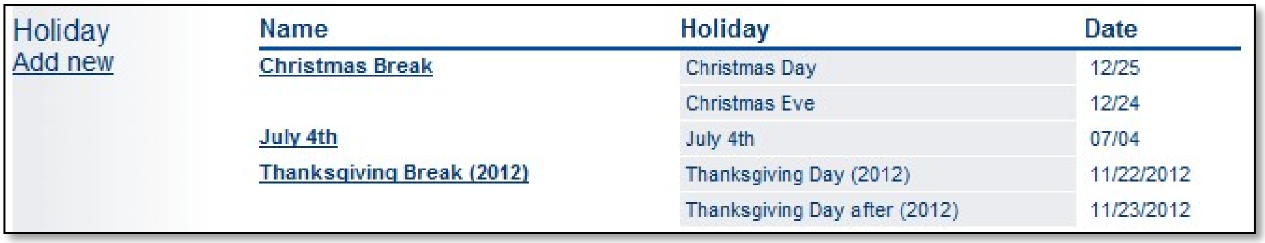

Holiday
Add new

Name	Holiday	Date
Christmas Break	Christmas Day	12/25
	Christmas Eve	12/24
July 4th	July 4th	07/04
Thanksgiving Break (2012)	Thanksgiving Day (2012)	11/22/2012
	Thanksgiving Day after (2012)	11/23/2012

Several Holiday Schedules.

Custom Schedules specify a small range of hours that override Holiday and On-Hours/Off- Hours schedules.

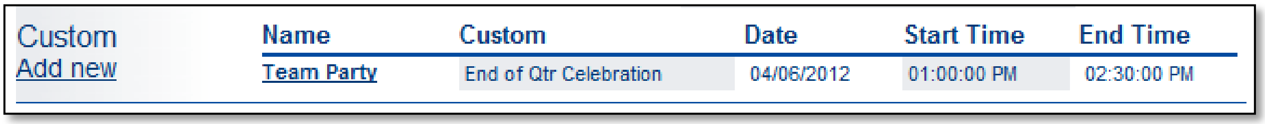

Custom
Add new

Name	Custom	Date	Start Time	End Time
Team Party	End of Qtr Celebration	04/06/2012	01:00:00 PM	02:30:00 PM

Custom Schedule for a short (1.5 hours) team-building event on a single day (Friday, April 6, 2012)

The following entities can have schedules applied to them:

Entity	Schedule(s) Used	Effect
User: Standard CHM	On-Hours/Off-Hours	Changes between Standard CHM
User: Extended Absence CHM	Holiday	Activates Extended Absence CHM
User: Custom CHM	Custom	Activates Custom CHM
Hunt Group: On Hours	On-Hours/Off-Hours	Toggles between answering calls and forwarding all calls to 'Off-Hours / Holiday' destination
Hunt Group: Holiday	Holiday	Forwards all calls to 'Off-Hours /
Holiday' destination		
Route Point	On/Off/Holiday/Custom	On/Off/Holiday/Custom call flow
Workgroup	On/Off/Holiday/Custom	On/Off/Holiday/Custom call flow

THE APPLICATION OF SCHEDULES

Schedules are activated, or triggered, differently for each type of Mitel entity:

- Hunt Group schedule changes are triggered by the HQ server.
- Workgroup schedule changes are triggered by the server that manages the Workgroup.
- Auto Attendant schedule changes are triggered by the server, or V-switch, that is playing the menu to the caller.

The time (and time zone) that is used to trigger a schedule change from one mode to another also varies by entity:

- Hunt Groups always change according to the time (and time zone) of the HQ server.
- Workgroups change according to the time (and time zone) of the server that manages the Workgroup.
- Auto Attendants change according to the time (and time zone) of the server or V-switch that is playing the menu to the caller.

> **Note:** The time used by a server to initiate a Schedule change is based on the time of the Operating System of that server.
>
> **Note:** The time used by a V-switch to initiate a Schedule change comes from the NTP server assigned to the V-switch and is modified by the Time Zone selected for the switch's Mitel site. For example, a V-switch in Chicago might get its GMT time from an HQ server in Seattle but will adjust its local time according to the Time Zone selected for the 'Chicago' Mitel site.

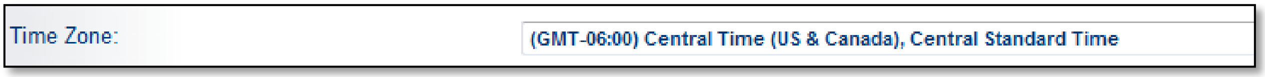

Time Zone: (GMT-06:00) Central Time (US & Canada), Central Standard Time

The control of the schedule change also varies by entity:

- Hunt Group schedule changes are always initiated by the HQ server:
- For switches managed by the HQ server: Loss of connectivity to the HQ server will cause the Hunt Group to remain in its currently active mode indefinitely. Once connectivity is restored, the Hunt Group will be changed to the proper mode.
- For switches managed by a DVS server: Loss of connectivity between the HQ server and the DVS server, or loss of connectivity between the DVS server and the ShoreGear switch, will cause the Hunt Group to remain in its currently active mode indefinitely. Once connectivity is restored, the Hunt Group will be changed to the proper mode.
- Workgroup schedule changes are controlled solely by the server that manages the Workgroup. Connectivity to any other server is not required.
- Auto Attendant changes are controlled solely by the server or V-switch that is playing the menu to the caller. Connectivity to any other server is not required.

> **Note:** On a V-switch that has lost connectivity to the HQ server, Auto Attendant menus with Schedules will change properly but Hunt Groups with Schedules will not.

HUNT GROUP SUMMARY:

The mode that is followed for all Hunt Groups is determined by the time of the Operating System on the Mitel HQ server. This is true regardless of what ShoreGear switch the Hunt Group is configured on, what site the switch is located at, or what Mitel server the ShoreGear switch is being managed by.

For example, with an HQ server in Austin (Central time), if a Seattle-based user dials a Hunt Group that is serviced by a Seattle-based ShoreGear switch (Pacific time), and that switch is managed by a Mitel server based in New York (Eastern time), the HQ server time (Austin time zone) would be used to determine the proper call handling for the Hunt Group.

For a Hunt Group to change from one mode to another based on a schedule, the HQ server needs to 'trigger' the change. If the HQ server cannot communicate to the switch that owns the Hunt Group, or if there is a loss of communication between the HQ server and a DVS server (or a DVS server and a switch) for those switches managed by a DVS, the Hunt Group will not change modes. When connectivity is restored the Hunt Group will then change to the proper mode.

WORKGROUP SUMMARY:

Schedule changes for Workgroups are based on the time of the Operating System on the Mitel server that the Workgroup is running on.

For example, a Workgroup assigned to a Mitel DVS server located in Barcelona, Spain will change based on the local time of the server in Barcelona regardless of where the call comes from.

AUTO ATTENDANT SUMMARY:

If an Auto Attendant menu is played from a Mitel server, schedule changes are based on the local time of that server.

If an Auto Attendant menu is played from a V-switch, schedule changes are based on the local time of the V-switch.

For example, if a call is received on a Seattle-based trunk and is directed to a New York- based Mitel server to play an Auto Attendant menu, the call will be routed based on the Schedule according to New York time.

If a call is received on a Seattle-based trunk and is directed to a Seattle-based V-switch to play an Auto Attendant menu, the call will be routed based on the Schedule according to the time of the V-switch (Seattle time).

The time-zone of the site the call entered is irrelevant.

ACCOMMODATING DIFFERENT TIME ZONES

When an entity needs to change from one mode to another, based on a Schedule, the time of the change is dependent on the time (and time zone) of the Mitel server, or V-switch, as indicated above.

Based on these feature differences, it can be important to create separate Schedules, and separate entities, for alternative time zones based on your desired behavior.

HUNT GROUPS:

For Hunt Groups, you should create a separate schedule for each time zone in which you have a Hunt Group configured. Each Schedule will need to be ’offset’ to accommodate the time zone difference from the site hosting your Hunt Group and the time zone of the HQ server.

If you originally had only one Hunt Group servicing multiple time zones, you should replicate the Hunt Group (use the ‘Copy’ button in Mitel Connect Director) and create separate Hunt Groups and separate Schedules for each time zone.

WORKGROUPS:

For Workgroups, you should create a separate schedule and separate Workgroup on the managing server (HQ and/or DVS) using an ‘offset’ for any site that uses a Workgroup that is in a different time zone.

For example, if you expect inbound calls to enter in one time zone but be routed to a Workgroup assigned to a Mitel server in another time zone you need to create a new (or replicate an existing) Workgroup so that it can be assigned the modified Schedule with an appropriate time-zone offset.

Consider an organization with two sites, Seattle and New York, but only one Mitel server located in Seattle.

The Sales Department wants all sales calls received in Seattle to be handled solely by Seattle representatives and all sales calls received in New York to be handled solely by New York representatives. To do this you need to create two separate Workgroups, each with its own agents, and each using its own, separate Schedule. The first Schedule, called ‘Seattle On-Hours’, would list 9:00am-5:00pm as its on-hours. The second Schedule, named ‘New York On-Hours’, would list 6:00am-2:00pm as on-hours. This accommodates the three hour time offset between the two locations.

In that same organization, the Service Department wants to use one Workgroup for all calls that are received from either location, and has multiple agents distributed at both sites. Agents log in from 9:00am - 5:00pm according to their respective local time zone. Calls that enter before 9:00am Eastern time or after 5:00pm Pacific time should be routed to the Off-Hours prompt and sent to voice mail. In this instance you can create just one Workgroup and use one schedule defining on-hours from 6:00am-5:00pm. This one schedule will route calls to any available agent from 9:00am New York time through 5:00pm Seattle time. Calls received outside these times will be routed to voice mail.

AUTO ATTENDANT MENUS:

For Auto Attendants, understand that the ‘nearest’ server or V-switch will always be used to play the Auto Attendant menu. If you expect inbound trunk-based calls to enter in one time zone but be routed to a V-switch or Mitel server in another time zone, you must create a new (or replicate an existing) Auto Attendant menu so that it can be assigned a modified Schedule with an appropriate time-zone offset.

For example: Consider the same organization as above with two sites, Seattle and New York, and only one Mitel server located in Seattle.

The administrator desires to use the same Auto Attendant menu for calls that are received from either location and offer an ‘On Hours’ prompt from 9:00am - 5:00pm local time, and present an ‘Off Hours’ prompt for calls received after-hours. You must create two Auto Attendant menus and two Schedules to accomplish this. Define the first Schedule as ‘Seattle On Hours,’ listing 9:00am-5:00pm as on-hours. Create the second named ‘New York On Hours,’ listing 6:00am-2:00pm as on-hours. This accommodates the three hour time different between the two locations.

Now create one Auto Attendant menu according to your needs, call it ‘Seattle Main AA Menu’ and apply the ‘Seattle On Hours’ Schedule to it. Make a second menu by copying the first menu (use the ‘Copy’ button within Mitel Connect Director), call it ‘New York Main AA Menu’ and change it to use the ‘New York On Hours’ schedule.

Finally, edit your Seattle trunk groups to route calls to the Seattle menu and edit your New York trunk groups to route calls to the New York AA menu.

If you were to add a DVS server to the New York site you would need to change the ‘New York Main AA Menu’ to use the 9:00am-5:00pm schedule. In fact, you can delete the New York menu entirely and direct all calls to the remaining ‘Seattle Main AA Menu’.

This is because all incoming calls from the New York trunks will be handled by switches in New York. Those New York-based switches will pick the DVS as their ‘closest’ Application Services end point and direct all calls to the DVS for Auto Attendant playback. The DVS in New York will use its local time zone for Schedule changes.

Note: Be aware that offset schedules do not accommodate failovers gracefully. For example, after adding the DVS in New York and making the changes listed above, consider the following: If the New York server goes offline and inbound calls are redirected to the HQ server to play the Auto Attendant main menu (remember that all Auto Attendants are fully distributed to all servers and all V-switches) the New York caller would be routed to the correct Auto Attendant Menu but it would be controlled by the Seattle Server’s time – which is three hours different.

SCHEDULES APPLIED TO CALL HANDLING MODES

Starting in Mitel 11 the time used to apply schedules to users changed due to the support of the new Client Application Services protocol (CAS) (see page 60). User-based Call Handling Mode changes are handled in the following fashion, based on Mitel version number:

MITEL 11 AND HIGHER

Call Handling Mode changes in Mitel 11 and above are driven by the Client Application Services protocol (CAS).

CAS only runs on Mitel servers, not on V-switches. Therefore all schedule changes for Call Handling Modes are driven by the servers in Mitel 11.

- If a user has their voice mail box assigned to a Mitel server, their Call Handling Mode will change, due to a Schedule assignment, according to the time of that server.
- If a user has their voice mail box assigned to a V-switch, their Call Handling Mode will change, due to a Schedule assignment, according to the time of the server that manages that V-switch.

MITEL 8.1, 9.X, 10.X

Prior to Mitel 11, Call Handling Mode changes were handled by the Voice Mail service instead of CAS. This behavior is similar to the way Auto Attendant menu Schedules are handled.

Since the Voice Mail service runs on V-switches (but CAS does not), a user in Mitel versions 8.1, 9.x, and 10.x. who has their voice mail box assigned to a V-switch would have their CHM changed according to the time of the V-switch (not its managing server).

Specifically:

- If a user has their voice mail box assigned to a Mitel server, their Call Handling Mode will change, due to a Schedule assignment, according to the time of that server (same as in Mitel 11).

- If a user has their voice mail box assigned to a V-switch, their Call Handling Mode will change, due to a Schedule assignment, according to the time of that V-switch, not according to the time of the server that manages the switch. This is different from Mitel 11.

DISASTER RECOVERY PLANNING

In addition to designing your multi-site Mitel deployment to be as resilient, fault- tolerant, and redundant as possible, you should also consider some worst-case scenarios such as natural disasters.

Large and/or extended site-wide outages can be classified into three different categories:

1. The site's equipment and circuits are functioning but the location isinaccessible. This

   scenario could arise due to many conditions such as:

- Poor weather – Snow prevents your workforce from coming into the office
- Police activity – You are unable to reach your location due to a parade
- Safety issues – A gas leak prevents you from occupying the building

2. The site's equipment has stopped working but the WAN/data circuits are still functioning.

   In this scenario, all Mitel UC equipment at a site (possibly even the main HQ site) is out of service and unreachable, but all other sites remain connected to each other and trunks are distributed to the other site.

   This scenario could occur due to circumstances such as:

   - Power outage – The HQ site is without power and all on-site equipment has stopped working, but all remaining locations still have data connectivity (e.g. MPLS circuits) and access to trunks at alternative sites
   - Fire, earthquake, tornado or other natural disaster

3. The site's equipment *and* all trunking and data circuits are down simultaneously.

   This can occur when you have designed your WAN in a hub-and-spoke topology and centralized your trunking. The loss of the central hub site can render the following inoperable:

   - All data circuits to all remote sites are disabled
   - All PSTN trunking at the central hub are out of service
   - All centralized equipment is unavailable and unreachable

We will discuss and address each of these disaster scenarios below.

ACCOMMODATING INACCESSIBLE SITES

In scenario #1, all services and all capabilities of the Mitel system remain functioning, but your staff is unable to reach their desks. Mitel is well suited to accommodate such a condition using features such as Extension Reassignment and the Mitel SoftPhone.

If your staff is unable to reach the office, they can use a variety of means to re-route incoming calls that normally would ring at their desk phone, to ring at an alternative location or device. This could be their cell phone, their home phone, a Mitel phone at another nearby site or location, or their computer itself via the Mitel SoftPhone.

To reassign their extensions, a user could use any of these methods:

- Communicator for Web (via a remote desktop, Citrix or Windows Terminal server connection)
- Communicator for Mac (via a software VPN)
- Communicator for Windows (via a software VPN)
- Communicator for Mobile/iPhone
- Mitel Mobility Client for iOS, Android or Blackberry
- Via the keypad of a Mitel phone at another site/location
- Via dialing into the Mitel voice mail system and re-assigning their desk phone to their *last* external destination

In each of these scenarios, the user can reroute calls bound for their extension to ring to at an alternative device or destination: a Mitel phone, the SoftPhone, or a PSTN number.

Call routing mechanisms, including Hunt Groups, agents that are part of Workgroups and even the Mitel Enterprise Contact Center, can route calls to users and agents wherever they are and whatever device they have assigned themselves to.

The Mitel architecture allows for all users, agents, supervisors and administrators to be completely virtual and will route all calls correctly and successfully.

Considerations:

- Be advised that external assignment to PSTN numbers uses additional trunks: one trunk for the original inbound call and another trunk for the re-directed outbound call to the user's external PSTN number. Proper trunk capacity planning is required.
- VPN connectivity (for Windows and Mac users) must have been previously configured and tested.
- Remote Desktop, Citrix and/or Windows Terminal Server connections must have been previously configured and tested.

ACCOMMODATING DOWN SITES

In scenario #2, if you lose an entire site due to a catastrophe or natural disaster, it is important not to lose functionality at your other locations. This can be accomplished, primarily, by designing WAN circuits that will continue to provide connectivity to all other sites if a primary location is incapacitated; and by distributing your call control, PSTN trunking and Application Services end points rather than centralizing them.

Additionally you could deploy a backup Mitel server by utilizing Double-Take software toaccomplish geographic redundancy for your HQ server thus alleviating you from completely losing HQ server functionality.

PREVENTING TOTAL LOSS OF SERVICES

To avoid scenario #3, a total loss of all services, Mitel strongly recommends the following:

- Avoid hub-and-spoke WAN topologies, especially those without alternative network routes and/or circuits
- Avoid centralizing your entire call control components in one location
- Avoid centralizing your PSTN trunks in one location

Instead, you should consider the following recommendations:

- If your call control equipment must be centralized, ensure you have redundant WAN circuits to connect remote sites to the central hub

- Consider adding a Disaster Recover (DR) site (described below) with redundant servers (using Double-Take,VMware or Hyper-V replication services, page 78) and redundant components
- Distribute your PSTN circuits to at least two major locations. Ensure that the total loss of one location (due to a disaster) does not prevent the remaining remote sites from accessing the alternative location. Calculate the minimum amount of PSTN trunking your business requires and ensure that *both*locations provide that minimum amount of trunking
- Distribute call control, PSTN trunks and Application Services whenever feasible

EXTENSION REASSIGNMENT AND SOFTPHONE CONSIDERATIONS

Extension reassignment and the SoftPhone are powerful tools to use on a daily basis, but can also be extremely important tools when making preparation for disaster recovery. Any user on the system can reassign their extension (given the correct User Group permissions). Any user with Communicator for Windows and Mac along with a Professional Access license can redirect their extension to their computer to use the SoftPhone.

Note:

- Reassigning a user from one phone to another (SoftPhone or otherwise) requires a change to the Mitel database on the HQ server. If the HQ server is inaccessible, reassignment to another phone (or SoftPhone) will not work.
- Ensure your disaster recovery plan includes HQ server redundancy.

DISASTER RECOVERY (DR) SITES

Many enterprises consider a Disaster Recovery (DR) site as part of their Business Continuity planning.

A DR site can be defined as a location that will provide an alternative set of services if the primary location for those services is unreachable. In a traditional data processing or non- Mitel telephony environment, that might mean duplicating or replicating equipment, services and servers in another location such as a data center, a colocation facility, or another site of reasonable size within the organization.

Since Mitel uses a distributed architecture by design, it is possible to make each site fully survivable and independent as a normal matter of best practice design. If any one site is isolated from the rest, or if your Headquarter site becomes separated from the other locations, the Mitel system, with its distributed call control and network call routing, can make such events nearly non-impactful.

The best practice when considering a DR site in a Mitel environment is to ensure that there is no single point of failure for WAN outages (avoid hub-and-spoke WAN topologies) and there is access to adequate trunks at multiple locations. There is no need for exhaustive duplication of equipment or servers – just proper placement of distributed equipment and distributed trunking.

HOT-STANDBY SITES

Some Mitel customers are vitally dependent on their phone system to remain in business, especially call centers. If the primary location used by agents becomes inaccessible, a nearby site with a set of at-the-ready, hot-standby Mitel phones, allows for an immediate and seamless mechanism to restore productivity to their agents.

If your main site is inaccessible, send your call center agents to the hot-standby location. An agent can sit at any desk with any Mitel phone and reassign their extension to the phone they are sitting at. All calls in, and out, are managed seamlessly by the Mitel system. All reporting and all supervisory and management features are identical and fully functional.

This ‘feature transparency’ allows any agent to operate with 100% of their capabilities and maintains 100% of all monitoring, reporting and supervisory features, no matter where in the organization the agent or the supervisor is located.

Note: For additional information on call center redundancy, see Application Note 10406

ECC REDUNDANCY BEST PRACTICES

To implement a hot-standby site you would:

- Add the site to the Mitel tree hierarchy as a standard site
- Deploy an appropriate amount of Mitel IP phones to accommodate the minimum number of relocated users/agents
- Deploy an appropriate amount of ShoreGear resources (switches) to provide call control for the IP Phones
- If centralized trunking will be used, no additional trunks need to be added to the hot- standby site
- If redundant trunking is required, order and deploy the necessary amount of PSTN trunks to accommodate the outbound calling needs of the relocated users/agents

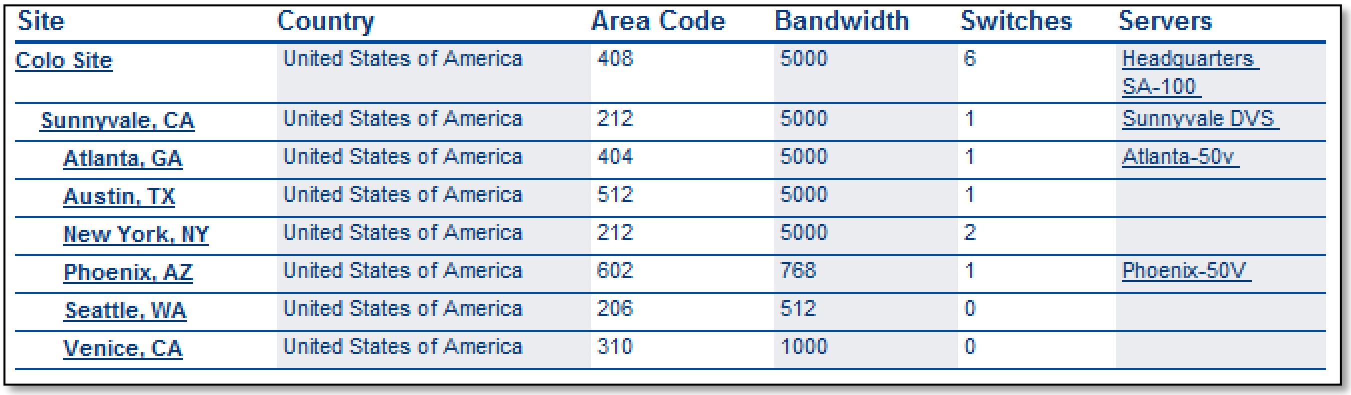

Site	Country	Area Code	Bandwidth	Switches	Servers
Colo Site	United States of America	408	5000	6	Headquarters SA-100
Sunnyvale, CA	United States of America	212	5000	1	Sunnyvale DVS
Atlanta, GA	United States of America	404	5000	1	Atlanta-50v
Austin, TX	United States of America	512	5000	1	
New York, NY	United States of America	212	5000	2	
Phoenix, AZ	United States of America	602	768	1	Phoenix-50V
Seattle, WA	United States of America	206	512	0	
Venice, CA	United States of America	310	1000	0	

Colo Site is the root of the tree. The main ‘Sunnyvale, CA’ site is the only child, and all other sites aregrandchildren

All other recommendations (server redundancy, N+1, and so on.) still apply to all sites in your design.

Note: While these recommendations are directed for a Colocation facility, it is recommended that you adhere to these recommendations regardless of whether a Colocation facility or data center will be utilized. These are considered best practices for a Mitel system.

DISTRIBUTED ROUTING SERVICE

The Distributed Routing Service (DRS) allows a large Mitel system to scale beyond 100 voice switches, up to a total of 500 switches (including SoftSwitches). In deployments with 10,000 users or less, DRS must be enabled on systems with 101 or more switches. In deployments larger than 10,000 users (up to 20,000 users), DRS must be enabled on systems with 61 or more switches.

When the Distributed Routing Service is *disabled*, Mitel switches in a system build an internal routing database (based on the Location Service Protocol, LSP, page 21) from the peer-to-peer communication with all other switches at all other sites. Every Mitel switch (ShoreGear and SoftSwitch) will learn and build routing information for all dialable end points in the system, including information regarding trunk location to be used for outbound calls. When calls are placed from any extension, each switch is able to route the call to the correct Mitel switch, at any site, based on its internal LSP routing table.

When the Distributed Routing Service is *enabled*, Mitel switches only exchange routing information (LSPs) with other switches at their same site, rather than with all switches at all sites. Therefore, each ShoreGear switch only maintains routing information for devices and trunks within its own local site.

All Mitel servers (HQ and all DVS servers) on the other hand, run an instance of the Distributed Routing Service, which maintains complete, system-wide LSP routing knowledge and information.

Note: V-switches do not run an instance of the Distributed Routing Service. Only HQ and DVS servers. For call routing, V-switches act the same as a standard ShoreGear switch.

EXTENSION CALL ROUTING WITH DRS

When DRS is enabled, extension call routing behavior is modified as follows:

When a call is placed to an *internal extension*, the Source ShoreGear switch that owns the call will look in its local LSP routing table to see if the destination extension is known. If the extension is found, the call will be routed to the local ShoreGear switch that owns that destination. If the destination extension is *not* found, the Source switch will contact a Mitel server asking for more information on the destination extension. The DRS service running on the server will respond with the necessary LSP details for the destination extension and the Source switch will then reach out to the Destination switch to place the call.

For extension calls, the Source switch uses a ‘Contact DRS if needed’ approach.

If the Source switch cannot contact a DRS service on a Mitel server the call will be routed to the switch’s internal Backup Auto Attendant.

PSTN CALL ROUTING WITH DRS

When DRS is enabled, PSTN call routing behavior is modified as follows:

When a call is placed to a *PSTN number*, the Source switch will immediately contact a Mitel server and ask for more information on routing the call. The DRS service running on the server will communicate the necessary information on all relevant trunk resources at all other sites to the requesting switch. The Source switch, now armed with complete trunk options, will reach out to the best and most appropriate trunk (based on the same trunk group prioritization rules listed in the Network Call Routing section on page 99 above) to place the call.

For PSTN calls, the Source switch uses an ‘Always contact DRS first’ approach.

If the Source switch cannot contact a DRS service on a Mitel server the call will be routed out a local trunk, if possible.

SELECTING DRS SERVERS

The Distributed Routing Service runs on all Mitel servers. Each ShoreGear switch learns of its two closest DRS servers from its managing server. Each ShoreGear switch will learn up to two Mitel servers to use for DRS queries, and will always contact its primary ('closest') DRS service first. If that server is unreachable the switch will contact its secondary ('next closest' DRS server.

The assignment of the two closest DRS servers is separate and unrelated to the selection process for the primary and secondary Application Services end points (page 26) although both selections use the Site Tree Hierarchy to determine relative proximity to establish the two 'closest' DRS servers and two 'closest' Application Services end points.

EXTENT OF DRS INVOLVEMENT IN CALL CONTROL

In general, if connection between a ShoreGear switch and another resource involves discovering the resource via the ShoreGear's local LSP table, then it will be affected by enabling the Distributes Routing Service.

For example, if a resource at another site is normally discovered and learned via the LSP process (such as user extensions, trunk group locations, and so on.) then the local ShoreGear switch's view of those resources will be restricted when DRS is enabled.

On the other hand, if connectivity to a resource at another location is *not* based on the local LSP table (such as the primary and secondary Application Services end points used for Voice Mail and Auto Attendant menus, the location of DRS servers, the activation of a Spare switch, the use of Service Appliance resources, and so on.) then enabling or disabling DRS does not change the behavior of the local ShoreGear switch.

MAINTENANCE PAGES WHEN DRS IS ENABLED

The Switch Connectivity display under the Maintenance menu in Mitel Connect Director lists all Mitel voice switches when DRS is disabled. When DRS is enabled, the switch connectivity table is organized by site.

DESIGN CONSIDERATIONS WITH DRS ENABLED

Consider the following multi-site topology:

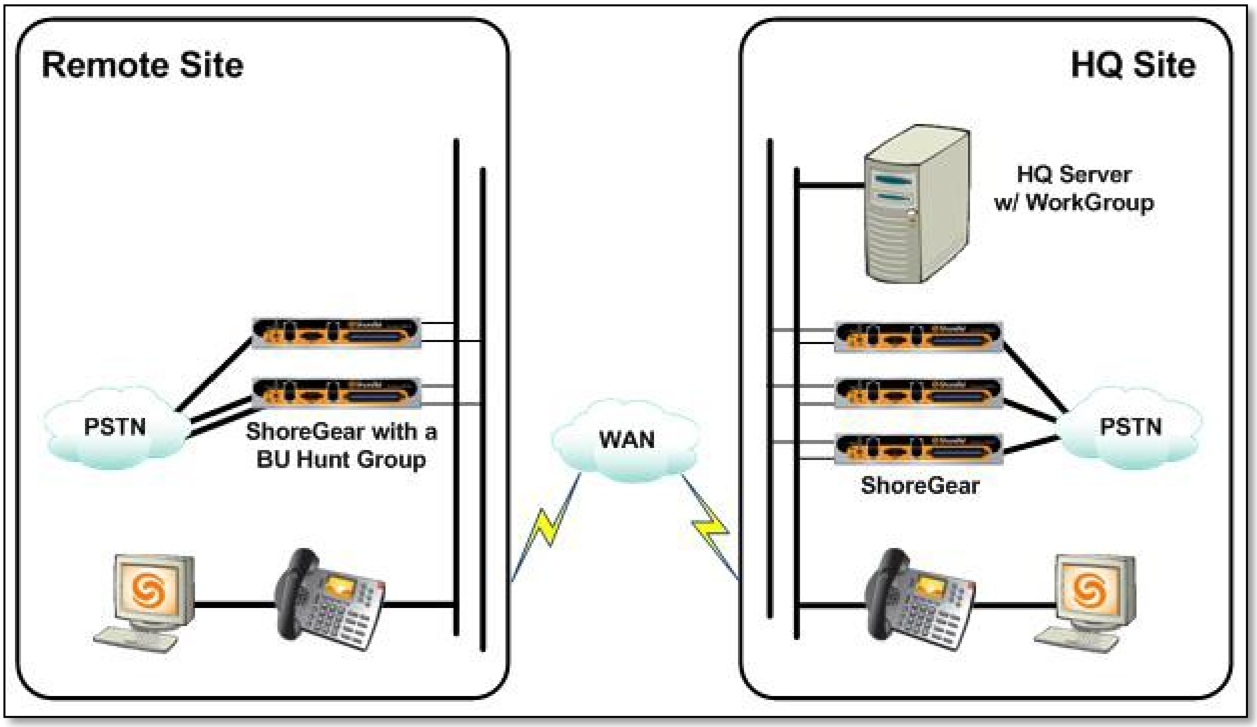

The ‘Workgroup’ (running on the HQ server) is configured with the ‘BU Hunt Group’ (running on a ShoreGear switch at the Remote Site) as its explicit Backup Extension.

When DRS is *disabled*, the ShoreGear switches at the Remote Site learn about all the extensions and entities at the HQ site through the exchange of LSP packets. They also learn the ‘forwarding destination’ for each of those entities. They know that the HQ server’s SoftSwitch owns the Workgroup extension and if that destination is not reachable they know that the forwarding destination for the Workgroup is the ShoreGear switch at the Remote Site, the owner of the Hunt Group.

With DRS *disabled*, if the WAN were to fail, an inbound call entering at the remote site destined for the Workgroup extension would be routed according to the following:

- The local ShoreGear switch handling the call needs to route the call to the Workgroup’s extension
- That Workgroup extension is found in its local LSP table
- That extension is marked as ‘Inactive’ (since the destination SoftSwitch is unreachable)
- The call will be routed to the Workgroup’s ‘forwarding destination’ which is listed as the Hunt Groups extension
- The Hunt Group extension is found in its local LSP table
- The Hunt Group extension is owned by a ShoreGear switch that is reachable
- The call will be sent to (and answered by) the Hunt Group service running on the local ShoreGear switch at the Remote Site

Now we *enable* DRS. The LSP tables on the remote-site's ShoreGear switches are flushed and rebuilt. But LSP packets are only exchanged between switches at the *same* site. And the switches are aware of the DRS service running on the HQ Server at the Headquarters site.

The switches at the Remote Site do *not* learn about the Workgroup extension.

With DRS *enabled*, if the WAN were to fail, an inbound call entering at the remote site destined for the Workgroup extension would be routed according to the following:

- The local ShoreGear switch handling the call needs to route the call to the Workgroup's extension
- That Workgroup extension is *not* found in its local LSP table
- The local ShoreGear will reach out to its DRS server, but the DRS service is unreachable due to the WAN failure
- Since there is no local entry for the Workgroup extension there is no known 'forwarding destination' either
- Since the DRS service is unreachable, the local ShoreGear switch will route the call to its internal Backup Auto Attendant
- The call will *not* reach the 'BU Hunt Group'

To remedy this situation, it is recommended that you modify your configuration according to the following:

- Create an additional 'Pilot Hunt Group' assigned to a local ShoreGear switch at the Remote site
- Configure the 'Pilot HG' with the 'Service Workgroup' as its primary member and the 'BU Hunt Group' as its secondary member
- Route calls intended for the Workgroup to the 'Pilot HG' instead of directly to the Workgroup

Now, with DRS enabled *or* disabled, calls will reach the Workgroup when the WAN is up and will reach the BU Hunt Group when the WAN is down.

Adding a DVS server at the Remote Site will also remedy the situation. If there had been a DVS server at the Remote Site, the local ShoreGear switch would have been able to successfully reach the DRS service running on the DVS server. The DRS service, being fully aware of all LSP information at all sites, would have told the local ShoreGear switch about the Workgroup extension and the forwarding destination for the Workgroup.

The local ShoreGear switch would see that the owner of the Workgroup is unreachable and would send the call to the Workgroup's forwarding destination: the 'BU Hunt Group'.

Note: There is a limit of 20 DVS servers in a single-image Mitel system. In addition, you could have also enabled Distributed Workgroup and had calls function normally, since the local ShoreGear switches would not need to reach the HQ server.

ADDITIONAL CONSIDERATIONS

In addition to the exhaustive coverage of features and recommendations above, there are several other best practices worth mentioning. These include:

EVENT FILTERS AND ALARM NOTIFICATION

As outages occur, the Mitel system can notify you, via email, of important failures, connectivity issues or device outages. ShoreGear switches can even provide notification when temperatures rise or a fan begins to slow.

Note: See Application Note 129 – 'Common Event Filters (Best Practices) for additional details.

ENTERPRISE CONTACT CENTER

The Mitel Enterprise Contact Center has a complete set of built-in redundancy features that overlay on top of the distributed Mitel architecture.

Note: See Application Note 10406 – 'ECC Redundancy Best Practices for additional details.

SERVICE PPLIANCES

The Mitel Service Appliance is a server-based appliance that is used for hosting some select Mitel UC and Application Services. Starting in Mitel version 12, the SA-100 was introduced as a collaboration appliance to provide audio conferencing, web-based conferencing, desktop sharing, instant messaging (IM/Chat) and presence server capabilities to the Mitel UC environment. Mitel version 12.3 introduced the SA-400, adding additional capacity and Mitel version 14.2 introduced the Mitel Virtual Service Appliance (which allows you to deploy IM with no physical hardware cost from Mitel).

ShoreGear switches use the Site Tree Hierarchy for locating Service Appliance resources. Similar to how a ShoreGear switch selects the two 'closest' Application Services end points, each ShoreGear switch will select the two nearest Service Appliances.

During the selection of the nearest Application Services end points, there is a guaranteed Mitel server located at the root of the tree; the HQ server. But there is no guarantee that an SA appliance has been added to the root of the site hierarchical tree. To ensure that every ShoreGear switch at every site will be able to locate at least one SA appliance a new 'Service Appliance Conference Backup Site' configuration option has been created.

This setting is configured in the Site Edit page of the Headquarters (or root) site. If an SA appliance has been added to a child site you can select that site as the 'Service Appliance Conference Backup Site' for the HQ site. This backup site is considered 'above' the root/HQ site, but only for purposes of a ShoreGear switch selecting the two closest SA appliances.

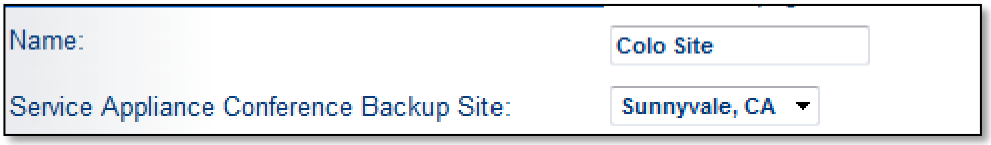

Name: Colo Site

Service Appliance Conference Backup Site: Sunnyvale, CA

Selecting the child site 'Sunnyvale, CA' as the backup site since the sole SA appliance has been deployed at the Sunnyvale location and not the Colo Site.

REDUNDANT POWER SUPPLIES

As discussed in the single site and multi-site redundancy sections above, it is recommended to deploy a few medium sized ShoreGear switches, or multiple smaller-sized ShoreGear switches, instead of deploying one large ShoreGear switch. For little or no increased cost, this “Don’t put all your eggs in one basket” recommendation distributes the risk and minimizes the impact of a device outage.

In doing so, the Mitel architecture inherits an additional redundancy attribute: distributed power supplies.

Many competitive products offer a chassis-based IP telephony system, involving a large chassis containing multiple slots, cards and modules. Those chassis-based systems can only be made survivable by adding a secondary (and often external and often expensive) power supply.

Instead of having a large single chassis with one large power supply, Mitel has chosen to add a much smaller and more efficient power supply to each of our ‘modules.’ This design element further separates the Mitel devices from reliance on a centralized component.

Additionally, just as you can add one more ShoreGear switch to achieve full device redundancy (aka N+1), adding that one additional switch effectively adds a redundant power supply. If any one power supply fails, causing a complete outage of a ShoreGear switch, the affected IP Phones will automatically and transparently reassign themselves to other resources at that same site.

Losing a power supply is no longer a critical event; it is simply a brief hiccup that is quickly and automatically worked around by the distributed intelligence of the ShoreGear switches.

Note: Refer to the definition of Resiliency and Redundancy at the start of this document.

CONCLUSION

Every Unified Communications environment is different. Every network topology is unique. Organizational budgets will vary drastically and the impact of outages will be highly dependent on each individual customer situation.

The information, suggestions and best practices described in this advanced Application Note should be reviewed and customized, as needed, by each individual customer and situation.

Rarely does an organization implement *all* of these recommendations, however a concerted effort to understand your business continuity needs and the potential disasters you desire to protect against will help determine which of these recommendations should be incorporated into your Mitel design.

RESOURCES, REFERENCES AND FURTHER READING

In general, refer to the three main documentation resources for your version of Mitel. These include:

- Mitel System Administration Guide
- Mitel Installation and Planning Guide
- Mitel Maintenance Guide

Also, refer to these specific additional resources:

- IP Phone Failover features:
  - Application Note: 10298: Mitel IP Phone: Failover Features
  - Administration Guide, Section 5.5: 'Failover for IP Phones and Spare Switch
- SoftPhone and Communicator failover features:
  - Application Note: 10338: Mitel SoftPhone: Features, Functions and Details
- Call Handling Mode 'Call Forward Always' Override:
  - Knowledge Base Article KB11777: 'Call Forward Always' Override
- Distributed Database:
  - Installation and Planning Guide, Section 4.5: 'Mitel Distributed Database'
  - Maintenance Guide, Section 2.4.2: 'Distributed Database'
- Server Redundancy:
  - Application Note: 10058 - Stratus Fault Tolerance
  - Application Note: 10109 - Double-Take with Mitel
  - Application Note: 10259 - Deploying Mitel on VMware
  - Application Note: 14008 – Deploying Mitel Servers under Microsoft Hyper-V 3.0
- Importing Prefix lists:
  - Knowledge Base Article KB11321: Creating Local Prefixes Lists
- 911 and Emergency Calling Resources:
  - System Administration Guide, Appendix A: 'Emergency Dialing Operations'
  - Mitel's 'Emergency Notification Service' Professional Services Application
- Backup and Restore Procedures:
  - System Administration Guide, Chapter 20: 'System Recovery'
- Schedules:
- Mitel System Administration Guide, Chapter 15: Configuring Schedules

- Application Note: 00129 - Common Event Filters (Best Practices)
- Distributed Routing Service:
  - Administration Guide, Section 9.13.1: ‘Distributed Routing Service’

APPENDIX A: HQ, DVS AND V-SWITCH FEATURES

In any Mitel deployment you must have one primary, or root, server called the Headquarters server. This is often referred to as the Director server or the HQ server. In addition to the single, mandatory HQ server, you can optionally add up to 20 additional Application Servers. These secondary servers are referred to by many synonymous names including: Distributed Voice Servers (DVS), Distributes Application Servers (DAS), Remote Application Servers (RAS) or Distributed Voice Mail (DVM) servers.

We use the term Distributed Voice Server, or DVS, throughout this document.

Starting in Mitel 8.1 a third device was added that supports some Application Services, called the Voice Switch, or V-switch. V-switches (ShoreGear-50V, 90V, 90BRIV, and so on.) have identical features and functionality as a standard ShoreGear appliance but add an internal, solid-state drive for on-board storage, and add additional processing capabilities to offer many of the services previously delegated solely to Windows-based Mitel servers.

The features and failover mechanisms of the Mitel Application Services end points (the HQ server, DVS servers, and V-switches) have been discussed earlier (see the section 'Application Services' on page 45 above).

The chart below will enumerate the complete list of features and functions performed by each of these entities.

Features of the Mitel HQ server, Distributed Voice servers (DVS), V-switches and standard ShoreGear switches:

Feature	HQ server	DVS servers	V-switches	Standard Switches	Comments
Management	✓				
Mitel Connect Director	✓				
FTP Services	✓	✓			Usable for upgrades
SMTP Services	✓	✓	✓		Used for Voice Mail delivery
Manages ShoreGear	✓	✓			
Database					
Master Database	✓				
Distributed Database		✓*			* For DVS with Distributed DB
Read-only subset of Database		✓*	✓	✓	* For DVS without Distributed DB

Call Detail Records					
Report Generation	✓				
Master CDR Database store	✓				
Store & Forward backup		✓	✓		
Computer Tel. Integration (CTI)					
Communicator Client	✓	✓			
3rd Party CTI	✓	✓			
Voice Mail Services	✓	✓	✓		
Store & Forward backup	✓	✓	✓		
Recorded Names	✓	✓			For non-assigned users
Recorded Name retrieval			✓		For non-assigned users
Workgroups	✓	✓			Selectable on a DVS server
Group Paging	✓	✓			Selectable on a DVS server
Auto Attendant	✓	✓	✓		
Acct. Code Collection	✓	✓	✓		
Applies Schedule changes	✓	✓			
IP Phone Options and Directory	✓	✓			
Call Control			✓	✓	

Bridged Call Appearances			✓	✓	
Hunt Groups			✓	✓	
Backup Auto Attendant			✓	✓	

APPENDIX B: VOICE MAIL PROMPT BEHAVIOR

The Voice Mail functions of the different Application Services end points are as follows:

Voice Mail Features	Mitel server (HQ or DVS)	ShoreGear V-switch
Local content stored for assigned users	Recorded name, 5 CHM greetings,	Local content stored for assigned users
Local content stored for other users	Recorded Names	None
Ability to dynamically retrieve content		Ability to dynamically retrieve content
Ability to 'voice announce' destination's extension	Yes	Yes
Ability to 'voice announce' destination's CHM	Yes	Yes

VOICE MAIL PROMPTS

If a caller is sent to a user's *assigned voice mail end point*, the caller will hear the following:

1. If the destination user *has not* recorded a 'Recorded Name' prompt and *has not* recorded a Call Handling Mode (CHM) greeting (for the active mode):
   - The caller will hear a built-in Mitel prompt, followed by the system reading the extension number, followed by the current Call Handling Mode, followed by the built-in 'instructions prompt':
     - "You are being forwarded to a Mitel Voice Mail System."
     - "Extension 1…2…3…4…"
     - "Is in a meeting"
     - "Please leave a message. When you have finished recording you might..."
2. If the destination user has recorded a 'Recorded Name' prompt but has not recorded a Call Handling Mode greeting (for the active mode):
   - The caller will hear a built-in Mitel prompt, followed by the Recorded Name prompt, followed by the current Call Handling Mode, followed by the built-in 'instructions prompt':
     - "You are being forwarded to a Mitel Voice Mail System."
     - "Benjamin Franklin" (in the end user's actualvoice)
     - "Is in a meeting"

    - "Please leave a message. When you have finished recording you might..."

3. If the destination user has recorded a Call Handling Mode greeting:
    - The caller will hear the user's greeting followed by the built-in 'instructions prompt':
        - "This is Ben, and you've reached my voice mail. Press 1 to find me or leave a message after the beep." *(in the end user's actual voice)*
        - "Please leave a message. When you have finished recording you might

    Note: The status of the Recorded Name prompt is irrelevant in this case.

If a call is answered by a Mitel *server* as an alternative Application Services end point, the caller will hear the following:

1. If the destination user *has not* recorded a 'Recorded Name' prompt:
    - The caller will hear a built-in Mitel prompt, followed by the system reading the extension number, followed by the current Call Handling Mode, followed by the built-in 'instructions prompt':
        - "You are being forwarded to a Mitel Voice Mail System."
        - "Extension 1…2…3…4…"
        - "Is in a meeting"
        - "Please leave a message. When you have finished recording you might..."

    Note: The Call Handling Mode (CHM) greetings are irrelevant because the alternative Application Services server does not store CHM greetings for non-local users.
2. If the user *has* recorded a 'Recorded Name' prompt:
    - The caller will hear a built-in Mitel prompt, followed by the Recorded Name prompt, followed by the current Call Handling Mode, followed by the built-in 'instructions prompt':
        - "You are being forwarded to a Mitel Voice Mail System."
        - "Benjamin Franklin" (in the end user's actualvoice)
        - "Is in a meeting"
        - "Please leave a message. When you have finished recording you might..."

If a call is answered by a *V-switch* as an alternative Application Services end point, the caller will hear the following:

1. If the V-switch *has* connectivity to a Mitel server (HQ or DVS):
    - The V-Switch will retrieve the destination user's Recorded Name from the Mitel server
    - The caller will hear a built-in Mitel prompt, followed by the 'Recorded Name' prompt, followed by the current Call Handling Mode, followed by the built-in 'instructions prompt':
        - "You are being forwarded to a Mitel Voice Mail System."
        - "Benjamin Franklin" (in the end user's actualvoice)
        - "Please leave a message. When you have finished recording you might..."

    Note: The Call Handling Mode (CHM) greetings are irrelevant because the V-switch does not store CHM greetings for non-local users.

2. If the V-switch *does not* have reachability to any Mitel servers, or if the destination user *has not* recorded a ‘Recorded Name’ prompt:
    - The caller will hear a built-in Mitel prompt, followed by the system reading the extension number, followed by the current Call Handling Mode, followed by the built-in ‘instructions prompt’:
        - “You are being forwarded to a Mitel Voice Mail System.”
        - “Extension 1…2...3…4…”
        - “Is in a meeting”
        - “Please leave a message. When you have finished recording you might...”

APPENDIX C: QUESTIONS AND ANSWERS

REBOOTING QUESTIONS

Q: Can a Standard ShoreGear switch reboot without connectivity to its managing Mitel server?

A: Yes. A standard ShoreGear switch, that has previously been added and connected to a Mitel system, can reboot without the need for connectivity to any Mitel server. It will boot up and load the locally-stored database content from its NVRAM FLASH storage. The rebooted switch will then contact all previously managed IP phones and begin exchanging LSP packets with all other previously known switches.

Previously assigned IP phones will receive dial tone and call control and be able to take and place calls.

Q: Can a ShoreGear V-switch reboot without connectivity to its managing Mitel server?

A: Yes - but only the ShoreGear switch features and functions of the V-switch will operate. All Application Services of the V-switch, such as Voice Mail, Auto Attendants and the Account Code Collection service, will remain inoperable until connectivity with the HQ server is restored.

Q: Can an IP Phone reboot without connectivity to a Mitel server?

A: Yes. A Mitel IP Phone, that has previously been added and connected to a Mitel system, can reboot without the need for connectivity to any Mitel server. Each IP Phone caches knowledge of its assigned resource switch and will attempt to contact that switch upon reboot. If the previously assigned ShoreGear switch is reachable the IP Phone will boot up and become functional.

If the previously assigned ShoreGear switch is not reachable, and IP Phone failover has been enabled, the IP phone will attempt to contact the IP Phone Configuration switches to be reassigned. But since the HQ server is unavailable no database changes can be made (including the reassignment of a phone to a new ShoreGear switch) and the IP Phone failover will fail.

Q: Can a Distributed Voice Server (DVS) reboot without connectivity to the Mitel HQ server?

A: No. Rebooting a DVS server without a copy of the Distributed Database (DDB) will cause the SoftSwitch on the server to become enabled but all local services like Voice Mail, Auto Attendants, the Account Code Collection service, Communicator client control, DRS, and so on. will be non-functional.

Connectivity to the HQ server is required for the local services of a DVS to become active after a reboot.

Q: Can a Distributed Voice Server (DVS) enabled with a local copy of the Distributed Database reboot without connectivity to the Mitel HQ server?

A: Yes. A DVS with DDB enabled stores a fully replicated copy of the database on its local hard drive. After a reboot, the DVS with DDB will load its local copy of the configuration database and then attempt to synchronize with the HQ server. All services on the DVS with DDB will become fully operational even if connectivity with the HQ server is down.

COLOCATION SITE QUESTIONS

Q: Are ShoreGear switches required at the site at the 'top of the tree'?

A: No. You do not need any ShoreGear switches added to the site which is the top of the site hierarchical tree

Q: Are ShoreGear switches recommended at the site at the 'top of the tree'?

A: Yes. It is recommended that you deploy two ShoreGear switches, configured as the IP Phone Configuration Switches, at the Colo site. IP Phone Configuration switches must be managed by the HQ server. These switches *can* be located in other (children) sites, if desired, but placing them at the top of the tree can make them better able to withstand a natural disaster and be more 'highly available' than if they were located at a 'less-protected' Admin building.

Having ShoreGear switches at the top of the tree is also recommended for use as 'Spare' switches that can be accessible to all child(ren) sites.

COMMUNICATOR QUESTIONS

Q: What Mitel server should the Communicator for Windows, Mac or Web be pointed to?

A: Any server. The Communicator client for Windows, for the Mac and for the Web can be pointed at any Mitel server. Each user's IP Phone is controlled by a ShoreGear switch and each ShoreGear switch is managed by a specific Mitel server. When a Communicator client tries to authenticate to any Mitel server it will be automatically redirected to the server that manages the switch controlling that user's phone.

Therefore, any server is a valid starting point for authentication of the Communicator clients. This is true for the very first launching of the client and for all subsequent re-launchings of the Communicator client.

With that said, best practice for Windows and Mac Communicators is to configure them to the 'HQ' server. If there is a local DVS server, 'HQ' will redirect the client to the nearest server. This will allow for redundancy, if there should ever be a local DVS server failure.

HUNT GROUP QUESTIONS

Q: When a Hunt Group is 'Made Busy' by dialing *18, is connectivity to a server required?

A: No. A user, with the proper permissions, can dial *18 to make a hunt group route calls to its 'Call Stack Full/Busy' destination without the need for the HQ server to be up and reachable. The change of forwarding destination (from routing calls to the ShoreGear switch that owns the Hunt Group to routing calls to the configured Busy destination) is a local setting that is controlled by the ShoreGear switch that owns the Hunt Group. Using *18 to change the Hunt Group State between the 'Made Busy' and 'Normal' states is managed by the switch and propagated to other switches using LSP packets. No server intervention is required.

APPENDIX D: LSP TABLES

Each ShoreGear switch is aware of all the local resources (user extensions, devices, trunks) that it is controlling. It is assigned and given this information from its managing Mitel server. This information, along with other call routing and permissions details are part of the relevant subset of the configuration database that each ShoreGear switch receives from its managing Mitel server. Each ShoreGear switch stores this database information in non-volatile FLASH memory so it is accessible after a reboot.

Each ShoreGear switch also exchanges information about its locally owned devices and trunks with other switches using the Location Service Protocol (LSP). Via the exchange of LSP packets, Each ShoreGear switch dynamically learns the whereabouts of all resources through the enterprise. LSP information is dynamic and stored in DRAM. If a ShoreGear switch reboots it loses all learned LSP information and must begin the process of exchanging LSP data with other switches again.

As other resources throughout the organization become unavailable or unreachable, the learned LSP entries for those extensions and trunks on those devices will be changed from an 'Active' state to an 'Inactive' state. For example, if a Mitel server is disconnected and a ShoreGear switch attempts to contact the Voice Mail or Auto attendant service on that server the connection will fail and the LSP entries for those end points will be marked as 'Inactive'. When the server connection is restored the server's SoftSwitch will begin propagating LSP packets and the local ShoreGear will update its LSP table with the new 'Active' state of those extensions.

LSP COMMANDS

To view the LSP tables in a ShoreGear voice switch you first Telnet to the switch (or SSH to the switch for V-switches) and run commands from the command line.

LSP commands include:

Command	Description
lsp_ping	Tests the LSP UDP communication to a far end switch.
lspConList	Displays switch connectivity to other switches.
lspTelList	Displays local and remote contacts.
lspTelList 1	Displays detailed information about local contacts.

Note: See the Mitel Maintenance Guide for details on switch command line usage and commands.

LSP EXAMPLES

The LSP commands and content in the following examples are taken from a ShoreGear 60/12 in a Mitel 12.2 environment using the following topology:

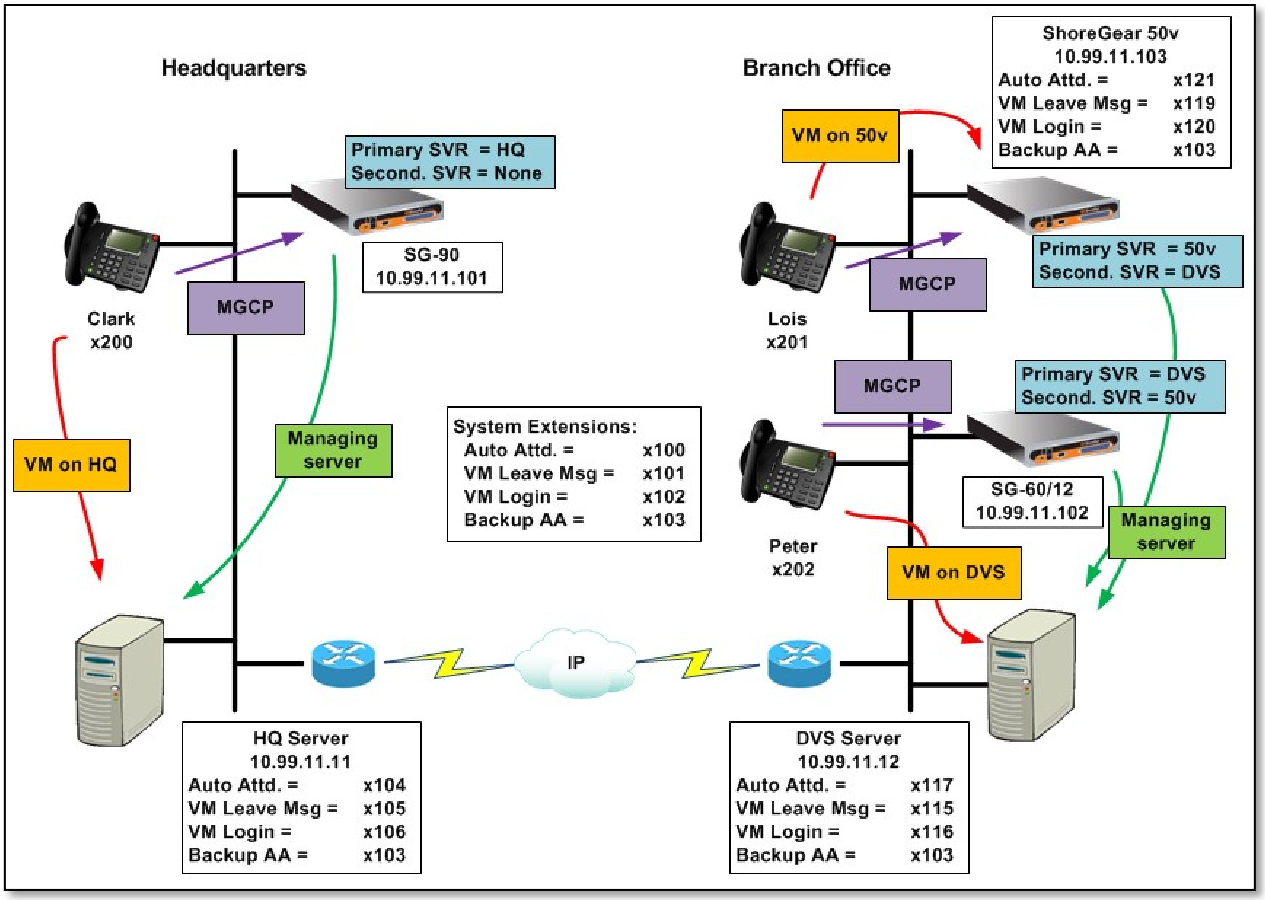

Sample network topology and Mitel devices, including assigned system extensions and server/V-switch extensions

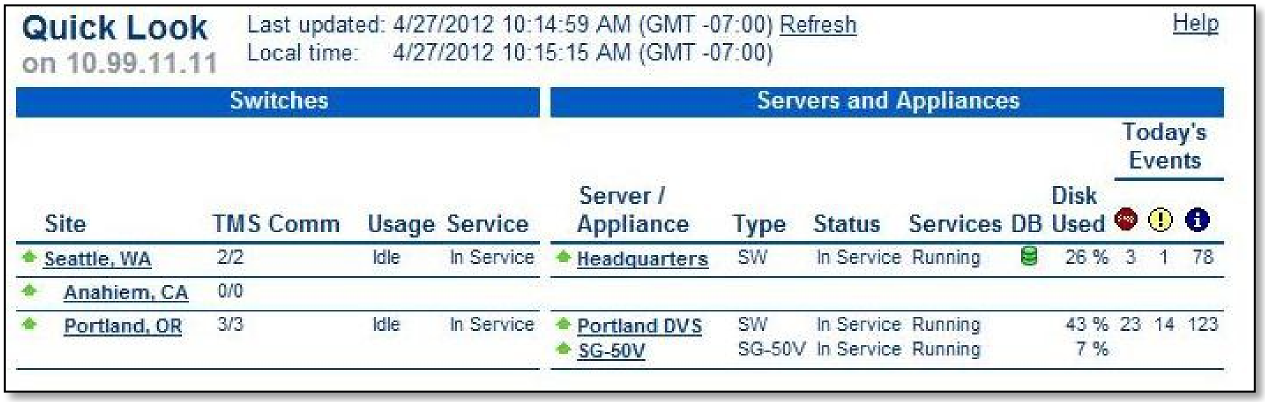

Quick Look on 10.99.11.11

Last updated: 4/27/2012 10:14:59 AM (GMT -07:00) Refresh
Local time: 4/27/2012 10:15:15 AM (GMT -07:00)

Help

Switches				Servers and Appliances								
										Today's Events		
Site	TMS Comm	Usage	Service	Server / Appliance	Type	Status	Services	DB	Disk Used	🛑	⚠	ℹ
▲ Seattle, WA	2/2	Idle	In Service	▲ Headquarters	SW	In Service	Running	🟩	26 %	3	1	78
▲ Anahiem, CA	0/0											
▲ Portland, OR	3/3	Idle	In Service	▲ Portland DVS	SW	In Service	Running		43 %	23	14	123
				▲ SG-50V	SG-50V	In Service	Running		7 %			

The Site Tree Hierarchy from Quick Look in Mitel Connect Director.

With all servers, switches and devices up and running we connected to the ShoreGear 60/12 using telnet to run commands.

LSPCONLIST

The '*lspConList*' command will display what other switches (ShoreGear and SoftSwitches) the ShoreGear 60/12 is communicating with and the state of that connection. All four other switches are listed as 'ACTIVE' and in sync.

```
-> lspConList

LIST OF CONNECTIONS

 session for <HQ Server>        10.99.11.11    - ACTIVE, STATE_IN_SYNC
 session for <SG-50V>          10.99.11.103   - ACTIVE, STATE_IN_SYNC

 session for <Portland DVS>     10.99.11.12    - ACTIVE, STATE_IN_SYNC
 session for <SG-90>           10.99.11.101   - ACTIVE, STATE_IN_SYNC
```

LSPTELLIST

The '*lspTelList'* command shows all the extensions learned from LSP packets from each of the other four switches.

Observe that the switch knows about its own 'LOCAL' extensions and trunk groups (learned from its managing server) and all the extensions and trunk groups learned from the four 'REMOTE' switches via LSP packets:

Note: Information in <COLORS> is added or highlighted for visual clarity purposes only and is not found in the normal command output.

```
-> lspTelList


LOCAL LIST

202 202 103 TGrp_6 107 100 101 102 DRS 107 DRS 100 101 102 107 100 101 102 Media Proxy


REMOTE LIST

 Switch 10.99.11.11 <HQ Server>

    204 204 114 114 104 104 105 105 106 106 109 109 203 203 205 205

 Switch 10.99.11.103 <SG-50V>

    ProxyVIP TGrp_5 103 119 119 120 120 121 121 122 122 201 201

 Switch 10.99.11.12 <Portland DVS>
    115 115 116 116 117 117 118 118

 Switch 10.99.11.101 <SG-90>

    TGrp_1 TGrp_4 299 299 103 108 200 200 ProxyVIP MOH_1
```

Also, note that additional LSP information is learned and display that is important, but not relevant to this example, including:

- SIP Proxy switches ('ProxyVIP')
- Music on Hold sources ('MOH_1')
- Media Proxy details for transcoding purposes ('Media Proxy')

LSPTELLIST 1

The '*lspTelList 1*' command displays detailed information about the local users, devices, extensions and trunk groups.

Notice that each dialable extension can be listed more than once, each entry with a different preference. Higher preference values are preferred. Each local extension has details about the proper call handling for that end point.

For instance, x202 (marked in **<red>** below) is the local user Peter who's IP Phone has been assigned to be controlled by this ShoreGear switch. The highest preference (preference = 1000) is the most preferred destination and indicates:

- 'Ringall' 'Calling Additional Phones' is enabled
- 'delay 3' The Additional Phones 'Number of Rings' is set to 3
- 'LSP - sip:202@10.99.11.102' The owning switch is at IP address 10.99.11.102

  The less preferred entry for x202 (preference = 1) indicates the 'forwarding destination' that Peter is set to: x101 which is the system extension for the Voice Mail system.

Also notice that the ShoreGear switch has:

- Two entries for DRS servers (marked in **<green>** below):
- The 'closest' being the DVS server with a higher preference of 12
- The 'next closest' being the HQ server with a lower preference of 11
- An entry for the local instance of the Backup Auto Attendant owned by itself:
- x103, preference = 1000 (marked in **<purple>** below)
- Three entries for each system extension. The first two are mapped to the primary and secondary Application Services end points. The last is mapped to the internal Backup Auto Attendant (x103):
- Auto Attendant (x100 – marked in **<blue>** below) maps to:
- x117 on the DVS server – the primary (preference = 12)
- x121 on the 50v – the secondary (preference = 11)
- x103 – the internal Backup Auto Attendant extension (preference = 1)
- Voice Mail 'Leave a Message' (x101) maps to:
- x115 on the DVS server – the primary
- x119 on the 50v – the secondary
- x103 – the internal Backup Auto Attendant
- Voice Mail 'Login' (x102) maps to:
- x116 on the DVS server – the primary
- x120 on the 50v – the secondary
- x103 – the internal Backup Auto Attendant
- Account Code Collection service (x107) maps to:
- x118 on the DVS server – the primary
- x122 on the 50v – the secondary
- x103 – the internal Backup Auto Attendant

```
-> lspTelList 1
LOCAL LIST


Contact:
  DN - 202 tShell:  LSP - pref 1     site_id 2 tel:101;
  phone_context=NPlan_0tShell: LSP -tel:101;phone_context=NPlan_0
Contact:
  DN - 202 tShell: LSP - pref 1000 site_id 2 Ringall delay 3 sip:202@10.99.11.102;
  phone_context=tShell: LSP - sip:202@10.99.11.102;phone_context=
Contact:
  DN - 103 tShell: LSP - pref 1000 site_id 2 sip:103@10.99.11.102;
  phone_context=tShell: LSP - sip:103@10.99.11.102;phone_context=

Contact:

  TRKGRP - TGrp_6 tShell: LSP - pref 1000 site_id 2 sip:@10.99.11.102;
  phone_context=TGrp_6tShell: LSP - sip:@10.99.11.102;phone_context=TGrp_6

Contact:

  DN - 107 tShell: LSP - pref 11 site_id 2 sip:122@10.99.11.103;
  phone_context=tShell: LSP - sip:122@10.99.11.103;phone_context=

Contact:

  DN - 100 tShell: LSP - pref 11 site_id 2 sip:121@10.99.11.103;
  phone_context=tShell: LSP - sip:121@10.99.11.103;phone_context=

Contact:

  DN - 101 tShell: LSP - pref 11 site_id 2 sip:119@10.99.11.103;
  phone_context=tShell: LSP - sip:119@10.99.11.103;phone_context=

Contact:

  DN - 102 tShell: LSP - pref 11 site_id 2 sip:120@10.99.11.103;
  phone_context=tShell: LSP - sip:120@10.99.11.103;phone_context=

Contact:

  DN - DRS tShell: LSP - pref 11 site_id 1 sip:DRS@10.99.11.11:5442;
  phone_context=tShell: LSP - sip:DRS@10.99.11.11:5442;phone_context=

Contact:

  DN - 107 tShell: LSP - pref 12 site_id 2 sip:118@10.99.11.12;
  phone_context=tShell: LSP - sip:118@10.99.11.12;phone_context=

Contact:

  DN - DRS tShell: LSP - pref 12 site_id 2 sip:DRS@10.99.11.12:5442;
  phone_context=tShell: LSP - sip:DRS@10.99.11.12:5442;phone_context=

Contact:

  DN - 100 tShell: LSP - pref 12 site_id 2 sip:117@10.99.11.12;
  phone_context=tShell: LSP - sip:117@10.99.11.12;phone_context=

Contact:

  DN - 101 tShell: LSP - pref 12 site_id 2 sip:115@10.99.11.12;
  phone_context=tShell: LSP - sip:115@10.99.11.12;phone_context=

Contact:

  DN - 102 tShell: LSP - pref 12 site_id 2 sip:116@10.99.11.12;
  phone_context=tShell: LSP - sip:116@10.99.11.12;phone_context=
Contact:

  DN - 107 tShell: LSP - pref 1 site_id 2 tel:103;XNR;
  phone_context=NPlan_0tShell: LSP - tel:103;XNR;phone_context=NPlan_0

Contact:

  DN - 100 tShell: LSP - pref 1 site_id 2 tel:103;XNR;
  phone_context=NPlan_0tShell: LSP - tel:103;XNR;phone_context=NPlan_0

Contact:

  DN - 101 tShell: LSP - pref 1 site_id 2 tel:103;XNR;
  phone_context=NPlan_0tShell: LSP - tel:103;XNR;phone_context=NPlan_0

Contact:

  DN - 102 tShell: LSP - pref 1 site_id 2 tel:103;XNR;
  phone_context=NPlan_0tShell: LSP - tel:103;XNR;phone_context=NPlan_0

Contact:

  DN - MediaProxy tShell: LSP - pref 1000 site_id 2 sip:MediaProxy@10.99.11.102:2727
  tShell: LSP - sip:MediaProxy@10.99.11.102:2727
```

LSP STATUS DURING AN OUTAGE

Now let's look at LSP information when an outage occurs.

When all servers and devices are operation the ShoreGear 60/12 lists all other switches as active and in sync:

```
-> lspConList

LIST OF CONNECTIONS

 session for <HQ Server>        10.99.11.11     - ACTIVE, STATE_IN_SYNC
 session for <SG-50V>           10.99.11.103    - ACTIVE, STATE_IN_SYNC

 session for <Portland DVS>     10.99.11.12     - ACTIVE, STATE_IN_SYNC
 session for <SG-90>            10.99.11.101    - ACTIVE, STATE_IN_SYNC
```

If we disconnect the Ethernet cable from the ShoreGear 50v switch we now notice the ShoreGear 60/12 lists the 50V as 'INACTIVE' and awaiting reconnection.

```
-> lspConList

LIST OF CONNECTIONS

 session for <HQ Server>        10.99.11.11     - ACTIVE, STATE_IN_SYNC

 session for <SG-50V>           10.99.11.103    - INACTIVE, STATE_WAIT_REBOOT ACK

 session for <Portland DVS>     10.99.11.12     - ACTIVE, STATE_IN_SYNC
 session for <SG-90>            10.99.11.101    - ACTIVE, STATE_IN_SYNC
```

All extensions and services in the LSP table for that ‘down’ switch will be treated as inactive forcing the Network Call Routing process to remove any trunks owned by the down switch from the selection process and causing calls destined for extensions on the down switch to be routed to the ‘forwarding destination’ of the dialed extension (e.g. the learned preference of ’1’ pointing to the Voice Mail system extension at x101).