



A MITEL
PRODUCT
GUIDE

MiVoice Business

Resiliency Guidelines

Release 10.0
Document Version 1.0

February 2023

Notices

The information contained in this document is believed to be accurate in all respects but is not warranted by **Mitel NetworksTM Corporation (MITEL[®])**. The information is subject to change without notice and should not be construed in any way as a commitment by Mitel or any of its affiliates or subsidiaries. Mitel and its affiliates and subsidiaries assume no responsibility for any errors or omissions in this document. Revisions of this document or new editions of it may be issued to incorporate such changes. No part of this document can be reproduced or transmitted in any form or by any means - electronic or mechanical - for any purpose without written permission from Mitel Networks Corporation.

Trademarks

The trademarks, service marks, logos and graphics (collectively "Trademarks") appearing on Mitel's Internet sites or in its publications are registered and unregistered trademarks of Mitel Networks Corporation (MNC) or its subsidiaries (collectively "Mitel") or others. Use of the Trademarks is prohibited without the express consent from Mitel. Please contact our legal department at legal@mitel.com for additional information. For a list of the worldwide Mitel Networks Corporation registered trademarks, please refer to the website: <http://www.mitel.com/trademarks>.

[®], TM Trademark of Mitel Networks Corporation

© Copyright 2023, Mitel Networks Corporation

All rights reserved

Contents

1	Introduction.....	1
1.1	About this Document.....	1
1.1.1	Who this Document is For.....	1
1.1.2	For More Information.....	1
1.2	What's New in This Release.....	2
1.2.1	Mitel MiVoice Business Release 8.0.....	2
1.2.2	Mitel MiVoice Business Release 7.0.....	2
1.2.3	Mitel Communications Director MCD Release 6.0.....	2
1.2.4	Mitel Communications Director MCD Release 5.0.....	3
1.2.5	Mitel Communications Director MCD Release 4.1.....	3
1.2.6	3300 ICP Release 9.0.....	3
1.2.7	3300 ICP Release 8.0.....	3
1.2.8	3300 ICP Release 7.1 UR2.....	4
1.2.9	3300 ICP Release 7.0 UR2.....	4
1.2.10	3300 ICP Release 7.0.....	4
1.2.11	3300 ICP Release 6.0.....	5
1.2.12	3300 ICP Release 5.2.....	7
1.2.13	3300 ICP Release 5.0.....	7
1.3	Glossary.....	8
1.4	About Resiliency.....	19
1.4.1	The Advantages of Resiliency.....	19
1.5	Mitel Resiliency Solution Overview.....	21
2	Mitel Resiliency Solution.....	24
2.1	Introduction.....	24
2.2	Device Resiliency.....	24
2.2.1	Resilient Devices.....	25
2.2.2	States of Resilient Operation.....	27
2.2.3	Resilient IP Device Behavior and User Interfaces.....	28
2.2.4	Resilient SIP Device Behavior.....	35
2.2.5	System Support for Resilient Devices.....	46
2.3	Call Resiliency.....	47
2.3.1	Call Survival.....	47
2.3.2	Phone User Interface.....	47
2.3.3	Resilient Call Cleardown.....	48
2.4	Feature Resiliency.....	48
2.4.1	Feature Handling.....	49

2.4.2 Resilient Feature Support.....	49
2.4.3 Administrative Feature Interactions with Resilient Routing.....	83
2.5 Line Appearance Resiliency.....	87
2.5.1 Description.....	87
2.6 Hunt Group Resiliency.....	91
2.6.1 Conditions.....	92
2.6.2 Programming.....	93
2.7 Personal Ring Group and Multi-device User Group Resiliency.....	93
2.7.1 Description.....	93
2.7.2 Conditions.....	95
2.7.3 Programming.....	95
2.8 Ring Group Resiliency.....	96
2.8.1 Conditions.....	96
2.8.2 Programming.....	96
2.9 T1/E1 Trunk Resiliency.....	97
2.9.1 Conditions.....	98
2.9.2 Installation and Configuration.....	100
2.10 SIP Resiliency Support.....	100
2.10.1 Head Office.....	100
2.11 Resilient System Elements and Interactions.....	101
2.11.1 About Resilient Clusters.....	102
2.11.2 Managing Resilient Devices from the MiVoice Business System Administration tool.....	103
2.11.3 Multiple Cluster Interaction.....	104
2.11.4 Interactions Between System and External Applications.....	104

3 ACD Resiliency..... 106

3.1 ACD Resiliency Solution.....	106
3.1.1 Description.....	106
3.1.2 About ACD Agent Hot Desking.....	108
3.1.3 Features Available to Resilient Hot Desk ACD Agents.....	110
3.1.4 Conditions.....	112
3.1.5 Agent Failover and Call Redirection.....	114
3.1.6 Licensing Requirements.....	116
3.1.7 Maximum Number of ACD Agents Agent IDs and Paths.....	117
3.1.8 Maximum Number of Hot Desk User Profiles.....	117
3.1.9 Configuring ACD Resiliency.....	118
3.1.10 Maintaining ACD Resiliency Data.....	132
3.2 Resilient ACD Configurations.....	135
3.2.1 Basic ACD Resiliency.....	136
3.2.2 Advanced ACD Resiliency.....	141
3.2.3 Full ACD Resiliency.....	145
3.2.4 ACD Resiliency Programming Examples.....	149
3.3 Converting a Traditional ACD Site: An Example.....	152
3.3.1 Before the Conversion.....	152

3.3.2 After the Conversion.....	152
3.3.3 Conversion Procedure.....	153
3.4 Resilient Virtual Contact Center Configurations.....	155
3.4.1 Resilient Virtual Contact Center Configuration 1.....	156
3.4.2 Advantages.....	159
3.4.3 Conditions.....	160
3.4.4 Resilient Virtual Contact Center Configuration 2.....	160
3.4.5 Local Call Breakout in a Virtual Environment including E911.....	164

4 Resilient Call Routing..... 167

4.1 About this Chapter.....	167
4.2 Resilient Call Routing ARS.....	167
4.2.1 General Guidelines for Resilient Routing ARS.....	167
4.2.2 Remote Directory Numbers Form.....	168
4.2.3 Resilient Call Routing Operation.....	172
4.2.4 Basic Resilient Call Routing Scenarios.....	175
4.3 Call Forwarding to Resilient Devices.....	177
4.3.1 Call Forwarding Terminology.....	177
4.3.2 Minimal Configuration.....	178
4.3.3 On System Call Forwarding Minimal Configuration.....	178
4.3.4 Non home Call Forwarding.....	180
4.3.5 Off System Call Forwarding.....	182

5 IP Console Resiliency..... 185

5.1 Introduction.....	185
5.1.1 For More Information.....	185
5.2 About Resilient Consoles.....	186
5.2.1 Console Behavior.....	186
5.2.2 User Interface.....	186
5.2.3 Operator Status.....	186
5.2.4 Service Status.....	186
5.2.5 Emergency Call Handling.....	186
5.3 Call Routing.....	187
5.3.1 Queue Handling.....	187
5.3.2 Out of Service Handling.....	191
5.3.3 Known Routing Conditions.....	192

6 Planning Resiliency..... 194

6.1 About this Chapter.....	194
6.2 Network Configuration Types.....	194
6.2.1 Non resilient Configurations.....	194
6.2.2 Enhanced Reliability Configuration.....	197
6.2.3 Resilient Configurations.....	198

6.3 Recommended Resilient Topologies.....	200
6.3.1 Resilient Standalone Environment.....	200
6.3.2 Resilient Distributed Network (Local).....	201
6.3.3 Resilient Hybrid Network.....	203
6.4 IP Device Distribution.....	204
6.5 Planning a Resilient Network.....	205
6.5.1 Familiarizing Yourself with the Engineering Guidelines.....	207
6.5.2 Identifying the Network Topology.....	208
6.5.3 Planning License Requirements.....	214
6.5.4 Planning Resilient Clustered Hot Desking.....	216
6.5.5 Planning VLANs and DHCP Options.....	220
6.5.6 Planning Call Routing.....	220
6.5.7 Planning ARS Routes.....	221
6.5.8 Planning T1 E1 Trunk Resiliency.....	223
6.5.9 Planning Voice Mail and Other Options.....	223
6.5.10 Planning the DECT IP Wireless Solution (EMEA) Networks.....	225
6.5.11 Considering Upgrading and Migration Issues.....	225

7 Implementing Resiliency..... 227

7.1 Implementation Overview.....	227
7.2 Assigning User and Device Licenses.....	230
7.3 Installing or Upgrading the MiVoice Business Systems.....	230
7.4 Programming the Cluster and ARS.....	230
7.4.1 Guidelines for Programming a Resilient Cluster.....	231
7.5 Setting up VLANs and DHCP Options.....	231
7.6 Implementing IP Trunking.....	232
7.7 Configuring T1 E1 Trunk Resiliency.....	232
7.8 Migrating the Cluster to Support RDN Synchronization.....	235
7.9 Provisioning Resilient Devices.....	236
7.9.1 Provisioning SIP Devices.....	236
7.10 Configure Resilient Hunt Groups.....	237
7.10.1 Enabling Resiliency for a Hunt Group.....	238
7.10.2 Adding a Member to a Resilient Hunt Group.....	239
7.10.3 Deleting a Member from a Resilient Hunt Group.....	239
7.10.4 Disabling Hunt Group Resiliency.....	239
7.10.5 Deleting Resilient Hunt Groups.....	240
7.10.6 Changing the Secondary of a Resilient Hunt Group.....	240
7.10.7 Migrating from Pre- 3300 Release 7.0 Hunt Group Configurations.....	240
7.11 Configure Resilient Multi-device User Groups.....	241
7.12 Configure Resilient Personal Ring Groups.....	241
7.13 Configure Resilient Ring Groups.....	242
7.14 Programming Voice Mail.....	242
7.14.1 Embedded Voice Mail.....	243
7.14.2 Non Resilient Centralized Voice Mail.....	247
7.14.3 Resilient NuPoint Unified Messenger Voice Mail Ports.....	248

7.15 Configuring the IP Consoles.....	249
7.16 Programming Listed Directory Numbers.....	250
7.17 Programming Class of Service Options.....	251
7.18 Programming Resilient Keys.....	251
7.19 Verifying Resiliency.....	254
7.20 Maintaining a Resilient System.....	255
7.20.1 Upgrading MiVoice Business System.....	256

8 Appendix..... 258

8.1 About this Appendix.....	258
8.2 General Engineering Guidelines.....	258
8.2.1 Performance.....	259
8.3 Guidelines for Resilient Network Design.....	262
8.3.1 Distribution of Servers Across the Network.....	264
8.3.2 Distribution of Devices Across the Network.....	265
8.3.3 IP Trunking.....	265
8.3.4 TDM Trunking.....	266
8.3.5 Network Power Provisioning.....	266
8.3.6 Providing Adequate Network Bandwidth.....	267
8.4 Resiliency, Functional Description.....	271
8.5 Link Failure Detection and Management.....	278
8.5.1 Link Failure Management Mechanism.....	279
8.5.2 Link Failure Detection by IP Phones.....	282
8.5.3 Detecting an ICP or Link Failure Through the Phone Display.....	282
8.5.4 MiVoice Business Health Check.....	284
8.5.5 MiVoice Business Behavior During Link Failure Detection.....	286
8.5.6 IP Phone Behavior During Resilient Mode Operation.....	286
8.6 Connecting and Configuring IP Network Devices.....	302
8.6.1 Connecting MiVoice Business Systems to L2 Switches.....	302
8.6.2 Connection Scenario for Applications Servers.....	305
8.6.3 Connecting IP Phones and Desktop PCs to Layer 2 Switches.....	306
8.6.4 MiVoice Business Console Registration Limitations.....	308
8.6.5 IP Phone Interaction with Routers, DHCP, and TFTP Servers.....	308
8.6.6 Controller Registry Configuration.....	312
8.6.7 DHCP Servers.....	314
8.6.8 Connection Scenario for Routers.....	325
8.6.9 Spanning Tree Protocol STP and Rapid Spanning Tree Protocol RSTP.....	335
8.6.10 External Applications.....	358
8.7 3300 ICP.....	358
8.7.1 De-rating ICP Capability to Support Resiliency and Clustering.....	359
8.7.2 Provisioning MiVoice Business Resources.....	360
8.8 Software Version Control.....	361
8.8.1 IP Device Interface Protocol Version Control.....	361
8.8.2 Ensuring Phone Application Software Version Consistency.....	364

Introduction

This chapter contains the following sections:

- [About this Document](#)
- [What's New in This Release](#)
- [Glossary](#)
- [About Resiliency](#)
- [Mitel Resiliency Solution Overview](#)

1.1 About this Document

This document provides a comprehensive overview of the Mitel[®] Resiliency Solution and offers customers the tools to understand, plan, and implement a resilient network.

1.1.1 Who this Document is For

This document is for certified Mitel MiVoice Business technicians and system administrators who plan, install, and program a resilient MiVoice Business system network.

To fully understand the information provided in this document, technicians and administrators must have a strong understanding of, or experience with

- Mitel MiVoice Business
- MiVoice Business network planning and configuration
- MiVoice Business System Administration Tool
- MiVoice IP phone functionality
- IP Networking

1.1.2 For More Information

For information about the MiVoice Business system, refer to the following online documents at Document Center:

- *Mitel MiVoice Business Engineering Guidelines*
- *Mitel MiVoice Business Resiliency Guidelines* (this document)
- *Mitel MiVoice Business General Information Guide*
- *3300 ICP Hardware Technical Reference Manual*
- *3300 ICP Technician's Handbook*

- *Mitel MiVoice Business System Administration Tool Help*

For information about any Mitel IP devices mentioned in this document, see Document Center.

1.2 What's New in This Release

This section describes the changes in this document due to new and changed functionality in MiVoice Business Release 10.0.

Table 1: Document Version 1.0

Feature/ Enhancements	Update	Location	Publishing Date
NA	No changes have been made to this document for the 10.0 release.	NA	February, 2023

1.2.1 Mitel MiVoice Business Release 8.0

New Resilient Devices

- Mitel 6920, 6930, 6940 IP Phones

1.2.2 Mitel MiVoice Business Release 7.0

New Resilient Devices and Applications

- MiVoice Conference Unit
- MiVoice Video Unit
- MiVoice Business Console

1.2.3 Mitel Communications Director MCD Release 6.0

New Resilient Devices

- Mitel UC360 Collaboration Point

New Resilient Features

This release adds resiliency support for the following features:

- Direct Transfer to Voice Mail

1.2.4 Mitel Communications Director MCD Release 5.0

- Resiliency handling of Line Appearances on IP devices and Attendant Console softkeys now covers more key configurations, in particular the configuration of multicall keys shared between resilient and non-resilient devices and between devices with different primary and secondary controller pairing.

1.2.5 Mitel Communications Director MCD Release 4.1

New Resilient Devices

- Mitel 5540 IP Console, 5320 IP Phone, 5360 IP Phone

1.2.6 3300 ICP Release 9.0

New Resilient Features

This release adds resiliency support for the following features:

- Phone Lock

New Resilient Devices

- The 5304 IP Phone, 5312 IP Phone, and 5324 IP Phone are supported as resilient devices.

1.2.7 3300 ICP Release 8.0

Resilient Ring Groups

Resiliency is supported for the Ring Groups feature. [Ring Group Resiliency](#)

New Resilient Features

This release adds resiliency support for the following features:

- Call Forward settings
- Call Forward - Follow Me - End Chaining
- Cancellation of extension call forwarding from the attendant console.

- Clear All Features and Remote Clear All Features
- Speed Call Keys (Note that Personal - Speed Calls are not fully resilient)

Resiliency Performance Improvements

The rate at which resilient sets fail over from the primary controller to the secondary controller has been improved. See [Performance](#).

New Maintenance Commands for Resiliency

The following new maintenance commands support T1/E1 trunk resiliency:

- EDT Courtesy Handoff
- EDT Cancel Handoff

The following new maintenance commands allow you to transfer resilient IP phones to the secondary controller when the primary controller needs to be rebooted:

- HST Courtesy Handoff
- HST Cancel Handoff
- HST Status Handoff

New Resilient Devices

- The 5560 IPT is supported as a resilient device.
- The Cordless Handset Module functionality, supported on the 5330, 5340, and 5360 IP Phones, is resilient.

1.2.8 3300 ICP Release 7.1 UR2

Agent Groups have been renamed to Agent Skill Groups within the documentation and the system tools to more accurately reflect the supported functionality.

1.2.9 3300 ICP Release 7.0 UR2

The 5330 and 5340 IP Phones are supported as resilient devices.

1.2.10 3300 ICP Release 7.0

ACD Resiliency

Resiliency is supported for ACD agents and ACD agent groups. See [ACD Resiliency Solution](#) on page 106.

Resilient Clustered Call Pickup

Resiliency is supported for the Clustered Call Pickup feature. See "Call Pickup - Clustered".

Resilient Hunt Groups

Resiliency is supported for voice and voice mail hunt groups. See [Hunt Group Resiliency](#) on page 91.

Synchronization of Resilient User and Device Data

The System Data Synchronization (SDS) feature now synchronizes the data of resilient users and devices between primary and secondary controllers. After you set up data sharing between the primary and secondary controllers, the application automatically maintains the synchronization of user and device data. This feature keeps user and device data, such as feature keys and phone settings, synchronized between the primary and secondary controller regardless of whether the users make the changes while they are on their primary or secondary controller. Refer to the System Administration Tool Help for more information.

T1 E1 Trunk Resiliency

The new T1/E1 Combo MMC supports T1/E1 trunk resiliency. If a site's primary controller fails, this feature automatically transfers the T1/E1 trunks to a T1/E1 Module on the secondary controller. See [T1/E1 Trunk Resiliency](#) on page 97.

1.2.11 3300 ICP Release 6.0

CXi Platform

The new CX and CXi controllers provide a cost-effective solution in environments with 64 or fewer users. The CXi platform supports Rapid Spanning Tree Protocol (RSTP). For more information on configuring RSTP, refer to Appendix: "Engineering Guidelines for Resiliency". For additional information on the CX and CXi platforms, please refer to the *3300 ICP Engineering Guidelines*.

For detailed information about the new hardware, refer to the *3300 ICP Hardware Technical Reference Manual* and the *Technician's Handbook*.

1400 users on LX

The 1400-user system configuration increases the capacity and efficiency of the 3300 ICP installations on the LX platform. These installations mainly include resilient and hot desk installations.

The 1400-user system configuration allows the following:

- 1400 active IP devices (an increase from 700)
- 5600 configured IP devices (an increase from 700)

For more information on the 1400-user system configuration, refer to the System Administration Tool Help.

New Supported Devices and Applications

6010 Teleworker Solution

The 6010 Teleworker Solution connects a remote office to the corporate voice network to provide full access to voice mail, conferencing and all the other features of the office phone system.

The Teleworker Solution can be completely configured at the head office using an IP phone. Using a two-click process, the phone is set to operate in teleworker mode. The telephone keypad is used to enter the IP address of the 6010 Teleworker Gateway installed at the head office. The phone can then be taken off-site and plugged into any broadband Internet connection. When the phone is powered up, it automatically establishes a connection with the Teleworker Gateway and is registered as a standard extension of the office phone system. The phone can also be returned to normal (non-teleworker) mode with the touch of a button.

In Release 6.0, the 3300 ICP supports resiliency for the 6010 Teleworker Solution.

Line Interface Module

The Line Interface Module provides 5220 IP Phone (Dual Mode) users with the ability to make and receive calls on an analog line in either of the following modes:

- Line Interface Module Mode: user can access the analog line at any time by pressing the programmed LIM key.
- Fail-over Mode: user can use the analog line when the IP connection has failed.

The 3300 ICP now supports resiliency for the Line Interface Module.

IP DECT OMM Resiliency

In Release 2.0 of the IP DECT Solution, both Open Mobility Manager (OMM) Resiliency and OpenPhone Resiliency are supported. To support OMM Resiliency, two Open Mobility Managers have to be configured in an OMM network; one working as the primary or active, and the other working as the secondary or standby OMM. In the event that the RFP designated as the primary OMM fails, the other RFP, designated as the secondary OMM, automatically assumes the role of Open Mobility Manager.

IP DECT OpenPhone 27

The IP DECT OpenPhone 27 is wireless telephone that is supported by the 3300 ICP. In Release 6.0, the 3300 ICP supports resiliency for this wireless device. The OpenPhone 27 has the same resiliency characteristics as other resilient IP phones.

1.2.12 3300 ICP Release 5.2

SpectraLink Wireless Telephones

SpectraLink NetLink e340 and NetLink i640 wireless phones are mobile IP handsets that are supported by the 3300 ICP. In 3300 Release 5.2, the 3300 ICP supports resiliency for these telephones. The SpectraLink wireless phones have the same resiliency characteristics as other resilient IP phones.

1.2.13 3300 ICP Release 5.0

Resilient Clustered Hot Desking

The 3300 ICP supports resiliency for clustered hot desk users and devices. For information about clustered Hot Desking, refer to the Release 5.0 *3300 ICP System Administration Tool Help*. For information about resilient clustered Hot Desking, refer to [Feature Resiliency](#) on page 48.

Layer 2 Spanning Tree Protocol Support

Release 5.0 introduces support for the IEEE 802.1D version of Spanning Tree Protocol (STP) on the LX 3300 Controller, in resilient as well as non-resilient systems. For network engineering guidelines and behavior pertaining to resilient systems running Spanning Tree, see [Connecting MiVoice Business Systems to L2 Switches](#) on page 302. For more information about support for IEEE 802.1D, refer to the 3300 ICP 5.0 *System Administration Tool Help*.

IP Phone Display Improvements

Release 5.0 IP phone display improvements apply to resilient phones during Fail-over and Fail-back and to resilient clustered Hot Desking phones during login/logout. These improvements include

- Improved progress indication during Fail-over/Fail-back and handoff operations
- New debug display

For information about display improvements, see [IP Phones](#).

5215 and 5220 DUAL MODE IP PHONES

3300 ICP Release 5.0 Resiliency and resilient clustered Hot Desking are supported on the 5215 and 5220 Dual Mode IP Phones. Pre-release 5.0 Resiliency and Hot Desking are not supported on these phones.

To provision resiliency for these phones, the primary and secondary ICPs must be Release 5.0 -- the phones do not fall over to Release 4.x ICPs.

For more information and for clustering requirements for resilient clustered Hot Desking on the dual mode IP phones, see Hot Desking (Resilient Clustered).

1.3 Glossary

ACD Gateway Controller

An ICP/ MiVoice Business system that hosts ACD paths. This system receives calls from the Central Office and distributes the calls to the resilient agent skill groups on the primary and systems.

Alternate ICP/ MiVoice Business System

The ICP/ MiVoice Business system (either primary or secondary) on which a resilient IP DEVICE is not in service.

See also: PRIMARY ICP/ MiVoice Business, SECONDARY ICP/ MiVoice Business, OTHER ICP/ MiVoice Business.

Boundary Node

There is only one boundary node for any given call, and it changes on a call-by-call basis. A boundary node is an element in a RESILIENT CLUSTER that routes calls to the home element of non-resilient devices or to either the primary or secondary ICP/ MiVoice Business system of resilient devices, depending on which element the device is in service on. The boundary node is the first Release 4.0 or later 3300 ICP that is encountered on a call route for calls originating inside or outside the cluster. For a call originating inside a resilient cluster, the call-originating element can also be the boundary node for that call. Boundary nodes route calls according to ARS. See [Basic Resilient Call Routing Scenarios](#) on page 175.

See also: TRANSIT NODE.

Call Resiliency

The ability to maintain existing calls in the presence of ICP/ MiVoice Business or network failures, and cleanly clear them down when they end. Part of the RESILIENCY solution. This combines CALL SURVIVAL and RESILIENT CALL CLEARDOWN.

See also: DEVICE RESILIENCY, FEATURE RESILIENCY.

Call Survival

During Call Survival, a device that has lost its controller retains PSTN access and preserves any active calls but cannot access phone features or dialing functionality. Call survival is the process of keeping voice streams related to active calls from being dropped when a device loses contact with its ICP/ MiVoice Business system (that is, when a device's ICP/ MiVoice Business system has failed during a call and begins failure-handling processes).

See also: CALL RESILIENCY.

Centralized Embedded Voice Mail

A centralized EMBEDDED VOICE MAIL server can be used to provide voice mail service to devices from a separate ICP/MiVoice Business system dedicated as a voice mail server. This server functions like an EXTERNAL VOICE MAIL server, except that if the hosting ICP/ MiVoice Business system fails, all sets lose their voice mail capability until it recovers. (See [Dedicating MiVoice Business Systems to Specific Functions](#)).

See also: CENTRALIZED EXTERNAL VOICE MAIL, RESILIENT VOICE MAIL.

Centralized External Voice Mail

An EXTERNAL VOICE MAIL server that is programmed as a centralized, shared voice mail server for the cluster.

See also: EXTERNAL VOICE MAIL.

Controller

Provides call control for IP DEVICES. Also referred to in this document as ICP or MiVoice Business system.

See also: PRIMARY ICP/ MiVoice Business system , SECONDARY ICP/MiVoice Business system .

Device Resiliency

The ability to move IP phones or other IP devices to a different ICP/ MiVoice Business system for control in case of ICP/ MiVoice Business system or network failures and recovery. Part of the RESILIENCY solution.

See also: CALL RESILIENCY, FEATURE RESILIENCY, IP PHONE, IP DEVICE.

Embedded Resilient Device Support

In a network or cluster that supports Remote Directory Number (RDN) Synchronization, you can configure resilient devices from the MiVoice Business System Administration Tool. MCD Release 4.0 or later software is required on all the element in the cluster to support this functionality. Prior to MCD Release 4.0, you had to use OPS Manager to configure resilient devices. Refer to the following books in the System Administration Tool Help for details:

- Voice Networking -> Manage Network -> Embedded Resilient Device Support,
- Voice Networking -> Manage Network -> Remote Directory Number Synchronization

Embedded Voice Mail

Exists on each ICP/ MiVoice Business system. Embedded voice mail can be configured as a local voice mail resource for the hosting ICP/MiVoice Business system or as a centralized, shared resource for the cluster. When embedded voice mail is configured to be resilient for a resilient device, the device receives voice mail service from its SECONDARY ICP/MiVoice Business system if its PRIMARY ICP/MiVoice Business system fails.

See also: CENTRALIZED EMBEDDED VOICE MAIL, CENTRALIZED EXTERNAL VOICE MAIL, RESILIENT VOICE MAIL.

External Voice Mail

Hosted on a dedicated server that is connected to a network element. In most cases the server is connected to an ICP/MiVoice Business system, but it does not have to be connected to a Mitel product. If the PRIMARY ICP/ MiVoice Business system of a device fails, the device fails over to its SECONDARY ICP/ MiVoice Business system and continues to be provided with voice mail service from its external voice mail server. However, if the external voice mail server or the element to which it is connected fails, all devices that obtain voice mail service from it also lose their voice mail capability until the server recovers. Can also be a centralized voice mail server.

See also: USER, RESILIENT USER, RESILIENT VOICE MAIL, EMBEDDED VOICE MAIL, CENTRALIZED EXTERNAL VOICE MAIL, CENTRALIZED EMBEDDED VOICE MAIL.

Fail / Failure / Failure Mode/ Fail State

Either a network failure or an ICP/ MiVoice Business system failure has occurred in the communication between an IP device and its current ICP/ MiVoice Business system. The device is currently not in contact or registered with any ICP/ MiVoice Business system; it has no call controller. A failure is considered primarily from the perspective of the IP device and may be the result of

- ICP/ MiVoice Business system failure
- Network failure between the device and the ICP/ MiVoice Business system

- A planned outage of the ICP/ MiVoice Business system (for example, for a scheduled re-boot or upgrade)

The device does not differentiate between the causes of failure; it only recognizes that a failure has occurred.

See *also*: Fail-back, Fail-over.

Fail-back

A special case of HANDOFF. The process of REHOMING to and re-registering with a primary ICP/ MiVoice Business system as a result of intentional handoff from an IP DEVICE's SECONDARY ICP/ MiVoice Business system.

See *also*: FAIL, Fail-over, HANDOFF/Fail-back INITIATOR, HANDOFF/Fail-back DESTINATION.

Fail-over

The process of a resilient device HOMING to and registering with a secondary ICP/ MiVoice Business system for control, after FAIL is detected. This may be:

- From PRIMARY ICP/ MiVoice Business system to SECONDARY ICP/ MiVoice Business system (Fail on primary ICP/ MiVoice Business system), or
- From secondary ICP/ MiVoice Business system back to primary ICP/ MiVoice Business system (fail on secondary ICP/ MiVoice Business system). Not an intentional HANDOFF.

See *also*: FAIL, Fail-back.

Feature Resiliency

The ability to retain access to features from a device that is in service on its secondary ICP/ MiVoice Business system. Part of the RESILIENCY solution.

See *also*: CALL RESILIENCY, DEVICE RESILIENCY.

Foreign Exchange Office (FXO)

An interface (normally an RJ-11 connector) that is used to connect to the PSTN (or Central Office).

Foreign Exchange Station (FXS)

An interface (normally an RJ-11 connector) that is used to connect to a standard telephone set (or FAX) and supplies ring, voltage, and dial tone.

Glare

Occurs when both ends of a telephone line or trunk are seized at the same time, usually resulting in deadlock of the line or trunk.

Handoff

The case of intentionally moving an IP DEVICE from its SECONDARY ICP/ MiVoice Business system to its PRIMARY ICP/ MiVoice Business system.

See also: Fail-back, HANDOFF/Fail-back INITIATOR, HANDOFF/Fail-back DESTINATION.

Handoff/Fail-back Initiator

The ICP/MiVoice Business system that initiates HANDOFF of an IP DEVICE to a different (destination) ICP/MiVoice Business system .

See also: HANDOFF/Fail-back DESTINATION.

Handoff/Fail-back Destination

The ICP/ MiVoice Business system that is the target of a HANDOFF from a different (Initiator) ICP/ MiVoice Business system . Upon successful Handoff, this ICP/MiVoice Business system becomes the new primary ICP/MiVoice Business system for the device.

See also: HANDOFF/Fail-back INITIATOR.

Health Check

ICP/MiVoice Business system -to-ICP/MiVoice Business system messaging mechanism that detects ICP/ MiVoice Business system recovery after FAIL and causes resilient devices to Fail-back to their healthy PRIMARY ICP/MiVoice Business system .

See also: HEALTHY STATUS.

Healthy Status

For the purpose of ICP/ MiVoice Business system-to-ICP/MiVoice Business system health checking, leading to HANDOFF or Fail-back, an ICP/ MiVoice Business system declares itself “healthy” (begins to respond to health check requests from another ICP/ MiVoice Business system) only after all its internal components are started and fully configured, and it is 100% ready to accept the requested handoffs. At the requesting side, the destination ICP/MiVoice Business system is declared “healthy” (meaning that phones can be handed off to the destination) only after multiple consecutive affirmative responses to the health check requests it sends to the destination MiVoice Business systems.

See also: HEALTH CHECK, HANDOFF/Fail-back DESTINATION, HANDOFF/Fail-back INITIATOR, IDLE STATE.

Home / Homing

The process by which an IP DEVICE seeks an alternate ICP/ MiVoice Business system (PRIMARY ICP/MiVoice Business system or SECONDARY ICP/MiVoice Business system) to register with. The MiVoice Business systems that the IP device searches for are defined by the ICP/ MiVoice Business system list processing, as a part of Fail-over or Fail-back and HANDOFF operations.

See also: REHOME / REHOMING.

Host ICP/MiVoice Business System

The ICP/ MiVoice Business system (either primary or secondary) that is supporting the resilient device.

Hot Desk ACD Agent

A hot desk ACD agent can log into any ACD set and the system will apply the agent's personal phone profile to the phone. After the agent logs in to the ACD set, the agent has access to his or her own feature keys and phone settings.

Idle State

A phone is idle when it is not hosting an active call stream. The idle state corresponds to a phone's readiness to Fail-over or Fail-back, when its corresponding call control state moves to CC Idle. Supplementary call control information such as held calls and pending callbacks is not a defining factor of idle state.

See also: NON-IDLE STATE.

IP Device

Refers to IP phones, IP appliances, and other desktop IP devices such as consoles and programmable key modules (PKMs).

See also: IP PHONE.

IP Phone

Refers to MiVoice IP Phones and appliances.

See also: IP DEVICE.

IP User License

A resilient user requires an IP User License at their PRIMARY ICP/MiVoice Business system but not at their SECONDARY ICP/MiVoice Business system . Validation is done upon initial configuration.

Local Hunt Group

A hunt group that contains only local members. The hunt group and all of its members are hosted on the same controller.

Network Hunt Group

A hunt group that contains remote members. Remote members are hosted on other controllers in the network. Remote hunt group members must be accessed through IP trunking, T1/E1, DPNSS, or XNET trunks.

Mass Fail-over/Fail-back/Handoff Event

A large scale Fail-over, Fail-back or other HANDOFF activity, involving large numbers of IP DEVICES at the same time (or close in time). Such events can create large flooding loads at the Destination ICP/ MiVoice Business system(s) involved and/or the underlying IP network.

See also: HANDOFF/Fail-back DESTINATION, HANDOFF/Fail-back INITIATOR.

Mitel Mezzanine Card (MMC)

A hardware module that you can install in a 3300 ICP to provide additional functionality such as trunk support (examples include the T1/E1 Combo MMC, Dual T1/E1 Framer, or Quad BRI Framer).

Mixed Cluster

A mixed resilient cluster contains at least two Release 4.0 or later ICP/MiVoice Business system plus any combination of pre-Release 4.0 ICP systems.

Network Resiliency

The ability to recover from underlying network-level failures.

Non-Idle State

A device is in a non-idle state when it is hosting an active call stream. Supplementary call control information such as held calls and pending callbacks is not a defining factor of non-idle state.

See also: IDLE STATE.

Other ICP/ MiVoice Business system

In a resilient cluster, an “other” ICP/ MiVoice Business system is any ICP/ MiVoice Business system in a cluster that is neither the primary nor secondary ICP/ MiVoice Business system of a given device. In the role of “other” ICP/ MiVoice Business system, controllers can route calls to a resilient device in service on either of its MiVoice Business systems, but it does not host the device.

See also: PRIMARY ICP/ MiVoice Business System, SECONDARY ICP/ MiVoice Business System, ALTERNATE ICP/ MiVoice Business System.

Primary ICP/ MiVoice Business System

The designated master ICP/ MiVoice Business system of an IP DEVICE. An IP device has only one primary ICP/ MiVoice Business system.

See also: SECONDARY ICP/ MiVoice Business system, OTHER ICP/ MiVoice Business system, ALTERNATE ICP/ MiVoice Business system.

Priority Stream

For the purpose of call survival operation, a particular stream is marked as a “priority stream” when it is to be preserved by an IP DEVICE in the event of an ICP/ MiVoice Business system or network failure. This priority stream status is flagged when the related call moves into a stable two-way talk state.

See also: STREAMING.

Point of Failure

A physical point (hardware) in a cluster or network that can fail, causing the failure or partial failure of telephone services.

See also: SINGLE POINT OF FAILURE.

Rehome / Rehomings

The process by which an IP DEVICE that is in service on its SECONDARY ICP/MiVoice Business system seeks out and re-registers on its PRIMARY ICP/MiVoice Business system. Also the act of seeking out and re-registering on the primary ICP/ MiVoice Business system.

See also: HOME / HOMING, Fail-back.

Remote Directory Number Synchronization

You can migrate a network or cluster of ICP/ MiVoice Business systems to support the synchronization of remote telephone directory entries across all the element databases. If you migrate a network or cluster to support Remote Directory Number (RDN) Synchronization, any telephone directory entries that you add, modify, or delete at an element are automatically distributed as shared data updates to the other elements in the network or cluster via System Data Synchronization (SDS).

RDN Synchronization is a pre-requisite for Embedded Resilient Device Support.

Resiliency

The term used to describe the ability of a network to adjust to and recover from system failure. Characterized by CALL RESILIENCY, IP DEVICE RESILIENCY, and feature resiliency, resiliency differs from redundancy in that there is no attempt to fully recover dynamic states such as call state. The Mitel Resiliency solution enables users to preserve phone functionality and features in the event of an ICP/ MiVoice Business system or network failure by allowing IP DEVICES to Fail-over to a backup controller, a SECONDARY ICP/ MiVoice Business system.

See also: RESILIENT CALL, USER, PRIMARY ICP/ MiVoice Business system .

Resilient Hot Desk ACD Agent

A hot desk ACD agent that has been configured as resilient. If the primary controller of a resilient hot desk ACD agent fails, the agent and the agent's group fail over to the secondary controller. After failover, the agent can answer ACD calls that are directed to the agent's group on the secondary controller.

Resilient Call

An active call that is preserved from the USER perspective through a CALL SURVIVAL operation. In a resilient system, all IP devices experience resilient calls, regardless of whether the devices themselves are resilient.

See also: RESILIENT CALL CLEARDOWN.

Resilient Call Cleardown

The process of clearing down a RESILIENT CALL as a result of USER action at the controller-less endpoint and/or trunk-side activity detected at the controlled end.

Resilient IP Console

A 5540 IP Console or MiVoice Business Console that is configured to be resilient. The console is considered to be an IP DEVICE.

See also: RESILIENT DEVICE.

Resilient Cluster

A cluster consisting of at least two Release 4.0 or later ICP/ MiVoice Business systems (plus any number or combination of other ICP/ MiVoice Business systems and associated IP Phones (and/or other IP devices), which are able to hand off phones, route calls, and maintain calls to each other.

See also: BOUNDARY NODE, TRANSIT NODE, IP DEVICE, IP PHONE.

Resilient Device

An IP DEVICE that is configured to be resilient, that is, the device is programmed to Fail-over or HOME to a SECONDARY ICP/ MiVoice Business system if it's PRIMARY ICP/ MiVoice Business system or link to the ICP/ MiVoice Business system fails.

Resilient Multi-device User Groups

Multi-device User Groups are an association of devices that are programmed as Multi-device users. The devices ring simultaneously (Ring All) when called.

Resilient Pair

A primary and its secondary controller. The System Data Synchronization feature shares user and device data at the "Resilient Pair" scope between a primary and secondary controller.

Resilient Peer

The other controller in a resilient pair (see above).

Resilient Personal Ring Groups

Personal Ring Groups (PRGs) are an association of single-user devices under a single Directory Number (DN). The devices ring simultaneously (Ring All) when called.

Resilient T1/E1 Trunk

A T1/E1 trunk that is configured as resilient. The trunk is connected to the main port on a T1/E1 Combo MMC in a primary controller. A second redundant port on the T1/E1 MMC Combo in the primary controller connects to a T1/E1 MMC in a secondary controller. If the primary controller fails, a bypass relay in the T1/E1 MMC connects the trunk to the T1/E1 MMC in the secondary controller. Any calls that are in progress on the trunk are dropped at the time of failure. After the relay connects the trunk to the module in the secondary controller, users can place incoming and outgoing calls on the trunk.

Resilient User

A USER who has a device programmed to Fail-over to a secondary ICP/ MiVoice Business system in the event of a PRIMARY ICP/ MiVoice Business system or network failure. This means that the user's directory number (DN) is associated with both a primary and secondary ICP/ MiVoice Business system, and the user's IP DEVICE can HOME to and register for call control on both. The primary and secondary call routing data for this user is propagated across all Release 4.0 or later ICP/ MiVoice Business systems in the cluster.

See also: IP USER LICENSE.

Resilient Voice Mail

(CENTRALIZED) EXTERNAL VOICE MAIL that continues to be provided to a device that fails over to its secondary ICP/ MiVoice Business system, or a (CENTRALIZED) EMBEDDED VOICE MAIL account that is configured on both the PRIMARY and SECONDARY ICP/ MiVoice Business system of a device. In either case, during a failure, voice mail service continues to be provided to a resilient device while it is in service on its secondary ICP/ MiVoice Business system.

- For external and centralized external voice mail, the voice mail server provides service to a device on its secondary ICP/ MiVoice Business system.
- For embedded voice mail, a device obtains embedded voice mail service from its secondary ICP/ MiVoice Business system. For centralized embedded voice mail, a device continues to obtain voice mail service from the voice mail server but loses voice mail service if that server fails.

Secondary ICP/MiVoice Business system

The designated backup or alternate controller of a resilient IP DEVICE. Non-resilient devices do not have a secondary ICP/ MiVoice Business system.

See also: PRIMARY ICP/ MiVoice Business system, OTHER ICP/ MiVoice Business system, ALTERNATE ICP/ MiVoice Business system.

Single Point of Failure

A single physical point (hardware) in a cluster or network that can cause the complete failure of telephone services to its associated devices if it fails. Non-resilient devices have a single point of failure because they can only be registered on one controller.

See also: POINT OF FAILURE.

Single Point of Provisioning

As a system management tool, the MiVoice Business System Administration Tool is the single point of provisioning for RESILIENT USERS and RESILIENT DEVICES. The System Administration Tool (via System Data Synchronization) also propagates Resiliency user data to other elements in the RESILIENT CLUSTER, and synchronizes the telephone directory information for all cluster elements.

Streaming

A term used to describe the transformation of telephone audio into IP packets and the subsequent transmission of these packets across the network between IP phones and other IP endpoints.

See also: PRIORITY STREAM.

System Data Synchronization

A feature that shares system programming data among a network or cluster of elements (ICP/ MiVoice Business systems) and synchronizes the system data of those elements with the data of a master element. This feature reduces the amount of time required to maintain consistent data across the elements. SDS also allows you to share user and device data between a primary and secondary controller.

Transit Node

A network element that allows calls to be routed through it, to a resilient ICP/ MiVoice Business system call destination. There can be any number of transit nodes on any given call route; however, there is only one boundary node for any given call. Any network element can participate in resilient call routing as a transit node.

See also: BOUNDARY NODE.

Trunk Resiliency

See RESILIENT T1/E1 TRUNK.

User

Desktop user of an IP DEVICE.

User license

See IP USER LICENSE.

1.4 About Resiliency

Maintaining telecommunications services in the event of system outage or failure is a primary concern for businesses, institutions, organizations, and communities (critical and emergency services), a concern which becomes no less important as we migrate our solutions to the world of IP. While failures and outages in any complex network are unavoidable, the MiVoice Business Resiliency solution provides the ability to preserve system functionality in the event of network difficulties by distributing network intelligence throughout “resilient” clusters that anticipate and pro-actively mitigate system failure and, thus, minimize the need for human intervention.

1.4.1 The Advantages of Resiliency

...in Maintaining Essential Service

Resiliency gives your network the ability to maintain calls in progress, handle new incoming and outgoing calls, and continue to provide voice mail service in the event of an MiVoice Business system failure or a network-level failure.

The resiliency solution is ideal for mission-critical environments in which it is essential that a voice path be maintained if an MiVoice Business system fails. Resiliency is achieved through setting up a network of MiVoice Business systems in a resilient cluster, which is a specially configured network of MiVoice Business systems that can direct IP phones and route and maintain calls.

Over Redundancy

Using self-correction techniques that take advantage of the IP-network characteristics of location independence and network element distribution, resiliency stands out from other more costly and less flexible solutions such as redundancy. While the redundancy model is highly effective and reliable, especially for traditional TDM solutions, it can be an unnecessarily costly solution for IP networks.

In Mitel's resiliency solution, IP phones can effect a virtual "move" from a failed primary call controller to a secondary controller without the need for wiring changes. Distributed resilient networks offer the advantage of being able to route around failed or otherwise inaccessible portions of an IP network, providing the following distinct advantages over the centralized, 1 + 1 hardware requirements of a redundant solution:

- There is no single point of failure in a resilient network.
- IP devices from one MiVoice Business system can be programmed to fail over to any number of different MiVoice Business systems.
- Network resiliency is reliable down to each IP device, not just to a shelf.
- Efficient use of existing hardware means lower hardware costs.

Rather than dedicating expensive, robust hardware to solving temporary and often infrequent system failures, resiliency makes efficient use of a system's existing capacity. In resilient networks, a secondary MiVoice Business system is not limited to acting as a dedicated backup call-control host. In fact, in most cases, the secondary controller can also function as a

- Primary controller for other devices in the network
- Voice mail server
- PSTN gateway
- IP network gateway
- Video conference controller
- Call center controller
- Wireless access controller
- Group controller, or
- All of the above.

See [Dedicating MiVoice Business Systems to Specific Functions](#).


1.5 Mitel Resiliency Solution Overview

The MiVoice Business Resiliency solution provides

- **Device resiliency** — A secondary ICP/ MiVoice Business system provides call control service to IP phones, SIP Phones and IP Consoles in the event of a failure on their primary MiVoice Business system or in the event of a networking failure between phones and their primary system . For more information, see [Device Resiliency](#) on page 24.
- **Call resiliency**—Calls in progress are maintained if a failure occurs. For more information, see [Call Resiliency](#) on page 47.
- **Feature resiliency**—Many features support resiliency and continue to be available to a device while it is in service on its secondary ICP/MiVoice Business system . Some resilient features may behave differently on a secondary system . For information about resilient features and feature behavior, see [Feature Resiliency](#) on page 48 .
- **ACD resiliency**— Resiliency is supported hot desk ACD agents. Resiliency is not supported for traditional ACD agents. Resiliency is also not supported for ACD Express Groups, but a certain level of resiliency can be programmed. For more information, see [ACD Resiliency Solution](#) on page 106.
- **Hunt group resiliency**—A hunt group and its members can be resilient. If the primary ICP/MiVoice Business system fails, support for the hunt group is automatically handled by the secondary system . For more information, see [Hunt Group Resiliency](#) on page 91.
- **Personal ring group and Multi-device User Group resiliency**—Personal ring groups (PRGs) and Multi-device User Groups (MDUGs) are resilient. Support for the PRG and MDUG passes to the secondary ICP/MiVoice Business system if the primary fails. For more information, see [Personal Ring Group Resiliency](#).
- **Ring group resiliency**—Ring groups are resilient. Support for the ring group passes to the secondary ICP/MiVoice Business system if the primary fails. For more information, see [Ring Group Resiliency](#) on page 96.
- **T1/E1 trunk resiliency**—If a primary ICP/MiVoice Business system fails, the T1/ E1 trunks connected to the T1/E1 Combo MMC in that controller are automatically transferred to a T1/E1 module in the secondary controller. For more information, see [T1/E1 Trunk Resiliency](#) on page 97.
- **Single point of provisioning**—The Mitel system management tool, the MiVoice Business System Administration Tool, is used to provision resilient users and devices on primary and secondary systems.

The following figure depicts the topology of a basic resilient cluster:

Figure 1: Basic Resilient Cluster



The transfer of phone service between primary and secondary MiVoice Business systems as well as the maintenance of calls in progress ensures that most system failures are not noticed by desktop users.

For detailed information about the Resiliency solution, see [Mitel Resiliency Solution](#).

Resiliency is primarily an IP-enabled capability that builds on existing MiVoice Business Voice Clustering (Portable Directory Number [PDN]), and call-routing principles. Existing clustering techniques are used to set up a cluster, which is then made resilient through the programming of boundary nodes, transit nodes, and call routing ARS (see [About this Chapter](#) on page 167).

For details about MiVoice Business clustering, see the *Mitel Clustering Design and Implementation* document.

Building on the MiVoice Business Platform

While only 3300 ICP Release 4.0 or later controllers can be resilient controllers, resilient clusters can contain pre-4.0 systems. These other devices cannot function as secondary controllers, but they can be part of a resilient solution as boundary nodes and transit nodes:

- **Boundary node**— 3300 ICP Release 4.0 or later systems can function as boundary nodes, actively routing calls between resilient systems and non-resilient (pre-Release 4.0) systems.
- **Transit node**—All cluster elements can function as transit nodes, passively allowing calls to be routed through them, on their way to a resilient ICP/MiVoice Business system.

Migration from a Pre 3300 ICP 4 0 Release

At Pre- 3300 ICP Release 4.0 customer sites where a site-wide upgrade to 3300 ICP Release 4.0 or later is not being undertaken, it may be necessary to have Pre-4.0 and 4.0 or later ICPs co-existing for some period of time.

Pre-4.0 ICPs cannot provide resilient functionality; however, they can initiate calls to devices that exist within resilient clusters, as long as

- At least one node in the route to the called device is resiliency-capable (Release 4.0+). The first such node in the route becomes the boundary node.
- The boundary node must be identified to the Pre-4.0 ICP.

For an illustration of this call scenario, see [Basic Resilient Call Routing Scenarios](#) on page 175.

For information about upgrading and migration to resiliency, see [Considering Upgrading and Migration Issues](#) on page 225.

This chapter contains the following sections:

- [Introduction](#)
- [Device Resiliency](#)
- [Call Resiliency](#)
- [Feature Resiliency](#)
- [Line Appearance Resiliency](#)
- [Hunt Group Resiliency](#)
- [Personal Ring Group and Multi-device User Group Resiliency](#)
- [Ring Group Resiliency](#)
- [T1/E1 Trunk Resiliency](#)
- [SIP Resiliency Support](#)
- [Resilient System Elements and Interactions](#)

2.1 Introduction

This chapter describes the main elements of the Mitel Resiliency Solution:

- Device resiliency
- Call resiliency
- Feature resiliency
- Line Appearance resiliency
- Hunt group resiliency
- Multi-device user group resiliency
- Personal ring group resiliency
- Ring group resiliency
- T1/E1 trunk resiliency
- Resilient system elements and Interactions

2.2 Device Resiliency

You can decide which users/devices in your system to make resilient. You may want to make all devices resilient or only those required for critical services. You can also select which MiVoice Business system each device will fail over to.

2.2.1 Resilient Devices

For a list of Mitel IP devices that support resiliency, see [Mitel Devices that Support Resiliency](#).

Table 2: Mitel Devices that Support Resiliency

Devices		Comments
IP Phones & IP Appliances	5224 IP Phone	You provision these devices with resiliency using SDS. For provisioning information, refer to the <i>System Administration Tool Help</i> .
	5304 IP Phone	
	5312 IP Phone	
	5324 IP Phone	
	5330 IP Phone (3300 Rel 7.0 UR2 and later)	
	5340 IP Phone (3300 Rel 7.0 UR2 and later)	
	Cordless Handset Module	
	5560 IPT	
	6920 IP Phone (MiVoice Business Rel 8.0 and later)	
	6930 IP Phone (MiVoice Business Rel 8.0 and later)	

Devices		Comments
	6940 IP Phone (MiVoice Business Rel 8.0 and later)	
	6905 IP Phone (MiVoice Business Rel 9.0 SP3 and later)	
	6910 IP Phone (MiVoice Business Rel 9.0 SP3 and later)	
Wireless SIP Devices	5603/5613, 5604/5614, 5617/ 5619, 5610, 5624/5634, 612 SIP-DECT, 622 SIP-DECT, 632 SIP-DECT, 650 SIP DECT.	
SIP Devices	5320, 5330, 5340, 5505, 6865, 6867, 6869, Mitel UC360 Collaboration Point, MiVoice Conference Unit, MiVoice Video Unit, Single Line UC Endpoint, MiVoice Conference, MiVoice Video, MiCollab Clients (Windows PC, Android, iOS).	You provision these devices with resiliency using SDS. For provisioning information, refer to the <i>System Administration Tool Help</i> .
Other Devices	5305 and 5310 IP Conference Units	A conference unit is automatically resilient by SDS if the IP device to which it is attached is resilient. No additional provisioning is required.

Devices		Comments
	5410/15 and 5412/48 PKMs	You provision a resilient Programmable Key Module, using SDS.
	5540 IP and MiVoice Business Consoles	You provision a resilient console with the System Administration Tool. For provisioning information, see IP Console Resiliency .
	Teleworker Solution	Resiliency is supported by SDS for Teleworker phones, on the LAN side of the MBG Teleworker Solution only (see Resilient Hot Desk Devices and E911 Service). Teleworker resiliency is backwards compatible with Release 3300 ICP Release 4.0. Requires minimum Teleworker Release 2.0.

For information about MiVoice IP Phones, refer to the appropriate User Guides at Document Center.

2.2.2 States of Resilient Operation

During Fail-over and Fail-back/handoff, resilient devices progress through several states of resilient operation.

For information about device user interfaces during the states of resilient operation, see [Resilient IP Device Behavior and User Interfaces](#) on page 28.

Fail over States

Fail-over is initiated when a resilient device loses connectivity to its current MiVoice Business system (primary or secondary). The device progresses through the following states of resilient operation:

- **Early warn state**—The early warn state occurs when a phone detects the possible loss of its link to its current MiVoice Business system . A phone can recover from this

state and resume normal operation on the MiVoice Business system if a failure is averted.

- **Fail state**—Fail state follows early warn state if a MiVoice Business system or network failure is not averted. Once in fail state, a device will Fail-over to its alternate MiVoice Business system .
- **Homing/Registration**—Once a failure has occurred, a resilient device homes to and registers on its alternate MiVoice Business system. If it is unable to do so, the device goes out of service. This state occurs during all Fail-over, Fail-back, and hand-off events, and it occurs only when a phone is idle (see [“Idle Devices”](#)). A phone does not enter this state until any resilient calls on it are ended.
- **Operation on alternate MiVoice Business system**—Once a device finishes registering on its alternate MiVoice Business system (primary or secondary), it is in service on that MiVoice Business system .



Note:

Registration is transparent to an end user.

Fail back Handoff States

When a resilient device is in service on its secondary MiVoice Business system and the system determines that the primary system of the device has regained healthy status, the secondary system initiates a Fail-back or handoff of the device to its primary system. When a device fails back or is handed off to its primary system, it progresses through three main states of resilient operation:

- **Fail-back/Handoff Initiation**— The secondary (and current) MiVoice Business system of a resilient device initiates the Fail-back/handoff process.
- **Homing/Registration**— The device rehomes to and registers on its primary MiVoice Business system .
- **Normal operation on primary MiVoice Business system**— Once the device finishes registering on its primary MiVoice Business system , it is in service on that system.

2.2.3 Resilient IP Device Behavior and User Interfaces

IP devices exhibit slight variations in resilient behavior and resilient user interfaces (UIs). This section describes IP phone, programmable key module (PKM), and conference unit behavior and UI for non-idle and idle phones during link failure detection, call survival, Fail-over to the secondary MiVoice Business system , and Fail-back to the primary

MiVoice Business system . For information about idle and non-idle states, see [Idle Devices](#) , and [Non-idle Devices](#).

For information about the IP Console behavior and resilient UI, see [IP Console Resiliency](#) .

Note:

SIP device resiliency behavior differs from that of MIP phones. For information about those devices, see [Resilient IP Device Behavior and User Interfaces](#) on page 28.

IP Phones

Behavior

Resilient IP phones are able to detect MiVoice Business system or network failures and fail over (home) to a secondary system . When a Fail-over event occurs, the secondary system monitors the health of the primary system through a health check process (see [MiVoice Business Health Check](#) on page 284). Phones fail back (rehome) to the primary system once it regains healthy status.

Phone functionality and features are available while a phone is in service on an MiVoice Business system but temporarily suspended during Fail-over and Fail-back/handoff processes, when the phone is not in service on an MiVoice Business system . IP calls in progress on an IP phone during the detection of a failure, experience call survival.

Note:

TDM calls (analog, PRI, or T1/D4 trunk) involving an IP phone do not experience call survival.

Display

Resilient phones use audio and visual signals to indicate their progress through the various states of Fail-over and Fail-back/handoff (see [States of Resilient Operation](#) on page 27).

The phone display varies, depending on whether a failure is detected while a phone is hosting a call or while it is idle and also whether a phone is failing over to a secondary system or failing back (or being handed off) to a primary MiVoice Business system . The displays for these states are presented in [Failure Detection and Fail-over Display for Non-idle Phones](#) , [Failure detection and Fail-over display for idle phones](#) , and [Fail-back/Handoff Display](#) .

Note:

The 69xx series IP phones do not indicate a fail-over state using audio or visual signals.

1. Early Warn State

- **Frozen display**—In early warn state the phone display is frozen, indicating that the phone has detected a potential MiVoice Business system or network failure. In this state, soft keys, speed call keys, and other user-programmable keys do not function. Once the display is frozen, if a user interacts with the phone, either by pressing keys or by taking the handset off-hook, a PLEASE WAIT message is displayed (see [Figure 4: Fail-back/Handoff Display](#) on page 33). The user must wait until the phone is registered again on an MiVoice Business system before placing a call or using features.

2. Fail State

- **Frozen display**—Once a MiVoice Business system or network failure is imminent, a resilient phone enters fail state operation. The display remains frozen as in early warn state.
- **Call survival tone**—A user on an active call from a phone in fail state hears a call survival tone consisting of two quick beeps that interrupt the voice stream every twenty seconds. The tone indicates that the phone has lost connection to its MiVoice Business system but also that the current voice stream will be preserved until the call is ended by either party (the resilient device user can only use the hookswitch or Cancel key to end the call). For the remaining duration of the call, the user cannot access features or use the dialpad. The phone fails over to its secondary system only once the call is ended (phone must be idle to fail over or back). See [Call Resiliency](#) on page 47 .

3. Homing/Registration

- **Progress Bar**—A progress bar appears in the top right-hand corner of the phone display to indicate that the phone is homing to and registering on a MiVoice

Business system . The progress bar appears during Fail-over and Fail-back/handoff.

Note:

Login and Logout for resilient clustered Hot Desking users is also displayed during homing/registration. See [Resilient Feature Support](#).

4. In Service


Once a phone registers on a MiVoice Business system , the display indicates whether the phone is in service on its primary or secondary system .

- **On secondary MiVoice Business system**—A flashing square appears in the top, right-hand corner of the display (when the phone is idle) to indicate that a phone is in service on its secondary system . Basic phone functionality and resilient features are available on the secondary system (see [Feature Resiliency](#) on page 48).
- **On primary MiVoice Business system**—When a phone is in service on its primary MiVoice Business system , there is no flashing square in the display.

Display for Failure Detected During a Call

For an illustration of the display that you see if your phone detects a MiVoice Business system or network failure during a call, see [Failure Detection and Fail-over Display for Non-idle Phones](#).


Figure 2: Failure Detection and Fail-over Display for Non-idle Phones



Display for Failure Detected While Phone is Idle

The display for an idle phone that detects a MiVoice Business system or network failure is the same as for a non-idle phone (see [Display for Failure Detected During a Call](#)), except that the user does not hear the call survival tone (see [Failure detection and Fail-over display for idle phones](#)).


Figure 3: Failure detection and Fail-over display for idle phones



Display During Fail-back/Handoff

A phone that is in service on its secondary system will Fail-back to its primary MiVoice Business system once that system regains healthy status. Fail-back/handoff will begin only once a phone is idle. [Fail-back/Handoff Display](#) illustrates the phone display during the Fail-back process.

Figure 4: Fail-back/Handoff Display



Debug Interface

At any time during the Fail-over and Fail-back/handoff processes, you can press the debug star (*) key to toggle between the default user display and a debug display that provides the IP phone boot sequence and detailed progress information. For information about the debug interface, refer to the *3300 ICP Technician's Handbook* (see "IP Phone Boot Sequence" and "Checking the IP Phone Progress Display").

Programmable Key Modules

Behavior

Resilient PKMs Fail-over and Fail-back with the resilient phones that they are attached to.

User Interface

PKM keys function only while the associated IP phone is registered on a MiVoice Business system. The keys do not function during call survival or while the phone is in the process of homing to and registering with a MiVoice Business system .

Conference Units

Behavior

Conference units function as another type of speaker for IP phones. They fail over to secondary system s and fail back to primary systems along with their resilient host

phones. Like IP phones, conference units experience call survival whether they are resilient or not, and, if they are resilient, conference units continue to function along with their host phones on a secondary system.

User Interface

Conference-unit functionality is available when the associated IP phone is in service on a MiVoice Business system and during call survival. Functionality is not available during the Fail-over and Fail-back processes when the phone is homing to and registering with a MiVoice Business system .

2.2.4 Resilient SIP Device Behavior

A resilient SIP device, on failure of service, automatically obtains service from another call server/proxy. This section contains recommendations on how to ensure effective resiliency on SIP devices.

A SIP device must be programmed to manage the basic scenarios of failing over to a new controller or being redirected to a new controller, in order to be deemed a resilient device on the MiVoice Business system. SIP device resiliency plays an important role in providing an enhanced customer experience. There are several ways to improve resiliency; this guide provides guidelines with regards to MiVoice Business.

Addresses

Resilient Devices with knowledge of two or more addresses

For a SIP device to be resilient in a clustered MiVoice Business network, one of the following conditions must apply:


- It should be programmed with either two or more MiVoice Business (elements) addresses within the cluster.
- It should obtain two or more addresses using DHCP option 125.
- It should be programmed with a Fully Qualified Domain Name (FQDN) that resolves with a Domain Name Server (DNS) to at least two addresses.

In each case, there should be a minimum of two addresses so that if one address is unavailable, the other address is able to provide initial service. Simultaneous failure of both addresses results in some services not being available.

Learning the MiVoice Business Home and Secondary Element Addresses

Each resilient SIP device is provisioned in the MiVoice Business cluster with a home element (HE) and a secondary element (SE). The addresses of the HE and the SE may be different from the addresses programmed directly in the devices, or in the network supporting the devices. The programming of the HE and the SE is done in the same

manner for SIP devices, and MiNET devices. Following is the illustration of a sample scenario:



In the above example, SIP extension 1000 has been configured to learn two addresses (A and C) - both MiVoice Business elements in the cluster. SIP extension 1000 first tries address A (step 1), and then learns (due to the redirection) that the HE programmed for this extension is B (step 2) and the SE programmed for this extension is C. SIP extension 1000 then tries its HE: B (step 3) and is informed that it has now obtained the service (step 4).

It is not necessary (as shown in this example) that the addresses programmed into the device on bootup directly line up with the HE and the SE. The SIP Protocol allows a REGISTER message to be redirected to another server so that the device will eventually reach an element from which it can obtain service.

The advantage of lining up the devices' network address with the HE and the SE is the time saved by not having devices redirected from one element to another. However, forcing the devices' network addresses to line up with the HE and the SE may be a complex process (depending on the number of devices and nodes in the network) that is better to avoid. We recommend that your devices use the P-Alternate-Server header (see [P-Alternate-Server Header](#)) for learning the HE and SE addresses and improving the search process for the MiVoice Business element from which it can obtain service. After a device learns the HE and SE addresses, it does not have to use the original addresses from the bootup/config file. MiVoice Business does not control the behavior of the device, and each device may behave slightly differently.

Redirection

A MiVoice Business element will respond to a Registration request from a device with a **301 Moved Permanently** response when it is required that the device move to a different MiVoice Business element for service (redirection).

The message contains two headers that provide useful information - the Contact header and the P-Alternate-Server header.

Contact Header

The Contact header is an industry-standard header that indicates to which MiVoice Business address the registration process must be moved (see example below).

Example: Contact: <sip:10.42.67.112>;q=0.60,<sip:10.39.68.165>;q=0.40

In this header, available MiVoice Business elements are listed with an assigned priority (or **q** parameter). The element with the higher **q** value is the next MiVoice Business element that the device tries. A redirection response is sent when:

- A device has chosen to register with an element that is not its HE or SE. In the best case, this happens only when the device first boots up and before it learns the HE/SE addresses. If devices choose not to use the P-Alternate-Server header, this can happen more frequently causing unnecessary load on the MiVoice Business cluster. The device learns through this response whether it is resilient or not (whether there are 1 or 2 addresses) and if resilient, the HE and SE addresses.
- A device registered with the SE is informed that it must failback to the HE to receive service. The SE does not redirect the Registration while the device is on a call; it will wait for the device to re-register (when idle) before redirecting.

The proprietary Mitel header called P-Alternate-Server is also included in the Redirection response; this header (described below) contains similar information.

Note:

If a device happens to register (not re-register) with the SE while the HE is working normally, the SE will not redirect the device to the HE immediately. Instead, the SE will provide service until the next re-register; this allows service during period(s) of uncertain network connection.

P-Alternate-Server Header

The P-Alternate-Server header, while included in a Registration redirection response (301 Moved Permanently), is also included in the 200 OK response to a REGISTER message (and other messages such as OPTIONS and SUBSCRIBE); this header provides the most value in the 200 OK response to the REGISTER message.

In the [above example](#), extension 1000 may have learnt the HE and SE addresses from the Redirection message. Suppose an extension 2000 has element A as HE and is directed there immediately on bootup, it would not have learnt the SE through redirection, and when element A failed it must learn the identity of its SE by trying the addresses learnt on bootup. By parsing the P-Alternate-Server header in the REGISTER 200 OK

response, the device would be able to directly contact its SE thereby providing the user with a better failover experience.

The header contains a list of entries with the following parameters:

- **URI parameter-** This mandatory parameter gives a standard description of the URI of a MiVoice Business element; this may be an FQDN instead of an IP address if the option **Allow FQDN for Resiliency** is enabled, and if the remote MiVoice Business element is configured with an FQDN in the **Network Elements** form.
- **State parameter-** This mandatory parameter indicates whether its associated URI should be the URI used by the device; a value of *active* indicates that the device must send requests to this URI. A value of *inactive* indicates that the device must switch to this URI if service from the active URI is interrupted.
- **Ping-interval parameter** - This parameter informs the device of the recommended heartbeat interval to use with a MiVoice Business element. The MiVoice Business element sets the ping-interval parameter to a value in the range 32 seconds through 900 seconds (default: 60 seconds), as defined by the **SIP Endpoint Resiliency – Device Heartbeat Interval** field in the *Controller Registry* form. The use of OPTIONS messages or a TLS heartbeat is optional, but if used, enables faster detection of service issues. Some devices already have a configuration field for the heartbeat interval; this field can be used to override the configuration value when the device is connected to a MiVoice Business element.

Example: P-Alternative-Server: <sip:10.10.10.10:5060;transport=udp>;state=active;ping-interval=60, <sip:10.20.20.20:5060;transport=udp>; state=inactive;ping-interval=60

The above header indicates that the first entry is the active element, and all requests from the device should be routed to this element; the second entry is the element to be used if the active element fails to respond to a request.

The presence of this header allows devices to alter the type of resiliency to be used with the controller. If this header is provided, it means that the controller requires the device to override locally learned addresses with the HE and SE addresses; it also means that the devices should wait for fail-back, and not poll the primary controller to re-obtain service. If neither the HE nor the SE is available, devices may revert to the addresses obtained originally to attempt to learn if changes have been made to the network.

Previous Bronze/Silver/Gold/Platinum levels

In the previous versions of MiVoice Business Resiliency Guidelines documents, the levels: Bronze, Silver, Gold and Platinum are defined; however, the nomenclature of these levels may create the impression that one level is better than another and that there are limited ways to improve resiliency. In reality, the methods used by a device to obtain its initial set of addresses using DNS, DHCP, or device configuration (as mentioned in the earlier sections of this document) does not matter; the notion that the device can switch between these addresses is useful, but it is actually better if the device

switches between the learned MiVoice Business HE and the SE. Using a heartbeat message (for example, an OPTIONS message) is also recommended.

Failover and Failback

Failover

The term *failover* is used to indicate that the device has detected a failure with its current MiVoice Business element, and is switching to another element to obtain service. Normally, devices failover to the SE address when the HE has been shut down, moved, restarted, or is unreachable; it is less common that a device fails over from an SE to the HE, but this can happen depending on re-registration times and how quickly an element is taken out of service after another element is returned to service.

The most common way a failover is detected is when a message request or heartbeat receives no response. For SIP messages (REGISTER, SUBSCRIBE, INVITE, or OPTIONS), the timeout interval is usually 32 seconds. Once a timeout is discovered, a device can check the available addresses and determine the address to try next. Preferably, devices must utilize the information learned about the HE and the SE addresses and go directly to the SE for service. A “Persistent TLS” Keep Alive (common with Mitel SIP devices) is also a good way to detect failovers.

During failover, it is also important to consider whether the user is active on a call; a failure in signaling does not necessarily imply that the media has been interrupted too. See [Advanced Survival Mode](#) for a description of how to improve customer experience by managing media that survives a failover.

The INVITE failure merits a separate discussion; see the next section. When a device fails over, it should Register and then Subscribe for the usual services (hotdesk, message summary and so on). MiVoice Business prefers that the devices do not poll the *dead* primary server but rather wait for Failback.

Failover detected by INVITE

If a call is placed using an INVITE to an HE that is no longer in service, the call typically fails. SIP messages have a timeout of 32 seconds, which is too long for a user to wait for ring-back. This timeout can be reduced to 5 or 6 seconds and a successful “first call after failure” call could be provided.

For a better customer experience, this call must be saved for the customer. After a short timeout (less than 32 seconds), the device can abandon the call, register (and subscribe if a hotdesk) with the SE, and send an INVITE message to that element instead. The user might detect that the setup time is longer than usual, but the call would be successful.

For this type of failover handling, it is recommended that you set a timeout period of 5-6 seconds (as the MiVoice Border Gateway (MBG) has implemented this feature for non-

Hotdesk users). The MBG reattempt is after 4 seconds, and therefore we recommend setting a slightly longer timeout for these applications to work together efficiently.

Failback using Registration Redirection

The Failback from the SE to the HE happens when devices are idle and they re-register with the SE when the HE is in service again. If devices are using a 10-minute Registration period, it means that all the idle phones will move back within a 10 minute period; all the phones do not failback at the same time.

Devices polling the HE to determine whether it is available can lead to all devices trying to register with the primary at the same time, causing high levels of congestion. It also prevents the SE from detecting exactly when a device is leaving, causing a delay in signaling information about this device to the HE. The SE uses the **SIP Endpoint Resiliency – Peer Heartbeat Interval** in the *Controller Registry* form (default 45 seconds) to determine when the HE will be back in service. Once the HE is back in service, it may take the SE this heartbeat interval to learn of the HE's new status, and start allowing devices to failback. When a device fails back, it should Register and then Subscribe for all the usual services.

Failback using Hotdesk Logout – Server Unavailable

Hotdesk users can be failed back before they send a re-register message. Hotdesk subscriptions can be logged out by MiVoice Business using the failure reason *Server Unavailable*, which triggers Hotdesk devices to switch to the alternate element on behalf of the user. In this case, the device should re-login automatically for the user. This is useful for SIP ACD Agents, which should move back shortly after the associated ACD Groups move.

Registration with both the HE and SE

MiVoice Business does not support the concept of being registered on both the HE and SE; instead, the failover/failback model used is described above.

Polling the Primary

Some phone vendors poll the primary server to detect if it is in service so that network phones can be moved back to the primary server. This is not the approach to be used with MiVoice Business because large numbers of devices can be serviced by the MiVoice Business application and moving all devices back at once may present a problem. Additionally, in some cases, the SE triggers information to the third-party services when a device is failing back. If a device abandons the SE in favor of the HE, some issues might occur for these services. Devices should wait for a REGISTER Redirection or a Hotdesk Logout, which indicates that a service must be received elsewhere.

MiVoice Business journaling of UDP/TCP information

When devices register successfully with MiVoice Business, the information about the UDP or TCP connection is journaled between the HE and SE. This does not apply to Hotdesk SIP devices or devices that register with TLS. The purpose of journaling is to allow calls to the device to be successful even when the device has not yet registered with the MiVoice Business element. For instance, the device might be registered with the HE, but the HE might have just gone out of service. In this case, an outbound call from the SE can be made to the device using the journalled registered contact information.

The journaling between network elements does not include Message Waiting subscription information. Subscription information is made available only after the device re-subscribes to the new element.

Resiliency Performance Impact

When a SIP device is resilient between two MiVoice Business elements, there is a heartbeat established between the elements. This heartbeat is used by the SE to determine when to move devices back to the HE. This heartbeat interface is used to journal information about SIP devices between MiVoice Business elements.

Message Waiting Subscriptions

The message waiting subscription (also known as message-summary) must be done after the Registration; MiVoice Business rejects attempts to Subscribe by devices that do not have a valid Registration yet. When a device fails over to the SE or fails back to the HE, the device should shortly thereafter Subscribe for message-summary and thereby allow the message waiting indication to be updated.

Typically Phones are configured with a longer subscription period than registration period; this can mean that if the subscription is terminated, it may take a long time to get restored. If MiVoice Business is shut down due to an upgrade, for example, it may terminate the MWI subscriptions during the shutdown. The termination of the message-summary subscription is a trigger that can be used by the device to re-initiate Registration and Subscription to the alternate MiVoice Business element. To ease the load on the MiVoice Business and other servers, it would be good to have a randomized delay before re-registering and re-subscribing due to this trigger (thousands of devices can quickly learn of this type of event). Devices that use a message-summary subscription should trigger a re-subscription if the connection to a controller has failed or moved; they should not wait the full period before reattempting the subscription.

Device Types

Mitel supports two groups of SIP devices:

- Fixed SIP devices use a fixed extension number and Register and usually Subscribe for message-summary, using UDP/TCP or TLS.
- Shared SIP devices allow a user to hotdesk to that device, and Register. One of their Subscriptions is a Mitel proprietary 'mihotdesk' Subscription.

Fixed UDP/TCP Devices

Fixed UDP/TCP devices use journaling as described in [Resilient IP Device Behavior and User Interfaces](#) on page 28 to allow MiVoice Business to direct calls to the device even when the device has not registered with the alternate MiVoice Business element.

Fixed SIP TLS Devices

For many devices, TLS connections are established from the device to the server. MiVoice Business follows this pattern and does not open TLS connections to devices. This prevents MiVoice Business from placing calls to a device before the device is registered. TLS connections from the device should be permanent, allowing calls to be made from the MiVoice Business element using the same connection. For devices that use TLS, early detection of failures on the HE is necessary to minimize time period during which calls cannot be placed to a device. It is recommended that you use a “Persistent TLS” keep alive timer.

Hotdesking

Some SIP devices provide the option of using the Hotdesking feature, which allows the sharing of a SIP device, and mobility between devices. Hotdesk SIP devices use a Mitel proprietary Subscription package. Similarly to TLS connections, this Subscription is initiated by the device and MiVoice Business does not journal these subscriptions. These devices can learn of failures through subscription termination cause codes and restore the connection for the user, thereby minimizing the time the phone is without service.

SIP Center of Excellence Designations Basic/Advanced

In place of the four levels (Bronze through Platinum) described in earlier versions of the MiVoice Business Resiliency Guidelines documents, we now recommend two levels: Basic and Advanced. Devices must support Basic to be classified as Resilient with a MiVoice Business system. If you want to provide a better customer experience, it is recommended that you implement some or all of the Advanced features. These advanced features, if implemented, are listed in the SIP-CoE interoperability guidebooks.

Basic Resiliency

- Able to obtain two or more addresses on bootup from DNS, DHCP, or local configuration.
- Use the REGISTER-301 Moved Permanently message to redirect registration to an alternate MiVoice Business element. A device must be able to failover from the Home Element (HE) to the Secondary Element (SE), and failback from the SE to the HE.
- At a minimum, a 32-second timeout for the REGISTER, SUBSCRIBE, INVITE or OPTIONS messages should trigger a Failover.
- Hotdesk devices need to support a *Logout-Server Unavailable* message from the element when it is trying to move devices to the other resilient element.

- After Failover/Failback – the device must restart all subscriptions (message-summary, hotdesk)

Advanced Resiliency

- Advanced: P-Alternate-Server

Use the P-Alternate-Server header in the REGISTER-200 OK message to store the HE and SE addresses. See [P-Alternate-Server Header](#) for more information.

- Advanced: Heartbeat

Use a light-weight heartbeat to periodically monitor the health of the MiVoice Business element to which the device is connected. This allows for the device to recover from failures faster without overloading the controlling element. See [Advanced Heartbeat](#) for more information.

- Advanced: Survival Mode

Continue existing conversations when a failure is detected until at least the Session Timer expires or the user takes an action which causes termination. Displaying a message on the device is also recommended. See [Advanced Survival Mode](#) for more information.

- Advanced: First Call after Failure

Implement a policy to time out a new call early if no 18x/2xx message is received. See [Resilient IP Device Behavior and User Interfaces](#) on page 28 for more information.

Advanced Heartbeat

Some sites have attempted to reduce the Registration Refresh Timer to act as a Heartbeat and detect failures of network elements; this puts unnecessary strain on MiVoice Business because the REGISTER message is a heavy-weight keep alive.

In a well-engineered network, the detection of a failure should be done by a light-weight keep alive mechanism that confirms whether the device is connected to the HE or SE. Two mechanisms have been shown to be effective in lowering the load. When using TLS, a Persistent TLS keep alive checks the connection at a rate of usually once a minute. When using UDP or TCP, devices can use an OPTIONS message at a similar rate.

The P-Alternate-Server header (see [P-Alternate-Server Header](#)) includes a ping-interval timer value that devices may wish to use; this value can be used by devices to synchronize keep alive intervals on all devices types. With a ping interval of about a minute, all devices will start the failover process within 90 seconds of a failure. A REGISTER message requires authentication and conducts many checks including querying for any name changes within MiVoice Business. A SIP OPTIONS message does not require authentication and is simply sent a response without any additional checks.

When using a light-weight heartbeat, the time between Register messages can be extended in most cases. Hotels that utilize the Register message to update the display on the Guest Room phone may still prefer to use a maximum 5-minute interval between REGISTER messages so that by the time a guest looks at the phone in their hotel suite, the phone's display has been updated. For other sites, a 10 to 15- minute interval between REGISTER messages should be sufficient. Since most phones refresh their registration at some percent of the registration period, the actual registration timer configured can be longer.

For instance, a 10 minute registration timer that fires at 50% of the interval will re-register every 5 minutes. So at 50%, a 10 to 15 minute interval between REGISTER messages means the user should configure a registration time of 20 to 30 minutes.

To check the status of the SIP Signaling on MiVoice Business, administrators can run the SIP Traffic Reports. These reports indicate problems with Congestion and Retransmissions in intervals of 15 minutes to 1 hour. If the number of these messages does not increase from interval to interval, it means that the SIP signaling is not being overworked.

Advanced Survival Mode

An important consideration in a resilient environment is the status of calls in progress at the time of failover. If a piece of equipment involved in the streaming fails, the audio/ video will be lost, and the user will have to retry the call. If streaming is not lost, and the device has detected a failure with the element to which it was connected, devices should maintain the streaming for a good customer experience. When some devices register with a new system, they drop any active calls. One could delay registering if there is a call in progress, but this would mean that the user might miss out on a new incoming call.

The preferred mechanism is to register and subscribe to the other MiVoice Business element and present everything as normal to the new controller; this state is termed **Survival Mode**. It is not the same as MiNET Resilient Talk. In this state, a message could be displayed on the device indicating that the device is in this state. If the user puts the current call on hold or tries to answer a new incoming call, the current call will be terminated because the signaling connection to the old element no longer exists.

Another reason for which the call might drop after entering **Survival Mode** is that the SIP Session Timer has expired. This timer is periodically signaling keepalive for a particular call between the device and the host element, and if the host element fails this timer will eventually be triggered. The SIP Session Timer must always be enabled as it prevents resources from being lost. It is recommended that this timer is set to 3 hours in resilient connections. Normally this is refreshed at a 50% interval, so it needs to be successful every 90 minutes for the call to survive. This mode is important for sites that need to perform upgrades while some calls may be active. Usually an administrator can busy out external connections before attempting an upgrade, but usually the whole network is not busied out. This allows the survival of internal device-to-device calls while the upgrade is in progress.

Invite Ringing Response Timer

For devices that are mobile, MiVoice Business provides an option called **Invite Ringing Response Timer**. This timer can be set to between 1 second and 30 seconds and defines the time interval MiVoice Business waits to receive a ringing response before the call is forwarded to another destination. In this way, MiVoice Business attempts to redirect the call to Voicemail or a cell phone as programmed for the end-user. This option can be set to 5 seconds, but if it is relied upon heavily, it can be reduced to 3-4 seconds depending on the device; it needs to be set longer than the time it normally takes for the device to begin ringing so that the device gets all its usual calls.

Check Signaling for Excessive messaging

Tuning a site for SIP is also useful and makes a positive difference in the responsiveness of the MiVoice Business system. The SIP Traffic Reports are a useful tool that can be used; these reports are generated when **Traffic Reports** are enabled on the MiVoice Business System Administration Tool (through the SIP TRAFFIC Maintenance Commands). These reports give statistics on calls made and received and also provide important information about Discarded or Retransmitted Messages and Error responses received. When an element is overworked or the network is having issues, there may be attempts to retransmit messages. To reduce the load, you can analyze the messaging being exchanged and accordingly change the appropriate timers to reduce the frequency of messages. Two common error responses that should be addressed are: 422 and 423.

- 423 Interval Too Brief responses (RFC 3261)
 - When the MiVoice Business **Registration Period Minimum** is higher than the *expires* time in a Register message, the MiVoice Business system is forced to delay the registration while the Registration time is negotiated. Users should ensure that the *expires* time programmed on SIP devices is equal to or greater than the minimum in the MiVoice Business system.
 - The same applies for Subscriptions (message-summary, hotdesk, and device user data); the *expires* times should be greater than or equal to the **Subscription Period Minimum**.
- 422 Session Interval Too Small (RFC 4028)
 - As with the 423 message, the Session Timer value is checked; the value set in the *SIP Device Capabilities* form is the minimum value that can be accepted. For resilient solutions, a time of 3 hours (10800 seconds) is recommended.

Most SIP devices do not negotiate early Media, and the PRACK message can be disabled to reduce the overhead of signaling on incoming and outgoing calls. This is done by setting the SIP Device Capabilities option **Disable Reliable Provisional Responses** to **Yes**.

Summary of Recommended Timer settings

1. Time between REGISTER messages – 10 minutes
2. Device heartbeat check – 60 seconds (see [Advanced Heartbeat](#))
3. Time between SUBSCRIBE messages – 1 hour or more
4. Session Timer – 3 hours
5. Peer heartbeat check – 45 seconds
6. First Call After Failure Timer – 5 seconds (see [Resilient IP Device Behavior and User Interfaces](#) on page 28)
7. Invite Ringing Response Timer – 5 seconds (see [Invite Ringing Response Timer](#))

For installations that do not utilize a *Device heartbeat check*(2), the *Time between REGISTER messages* (1) setting can be lowered to 2.5 minutes between messages. This might mean that devices are configured with a 5-minute Registration period since most phones refresh at 50% of the timer period.

Time between SUBSCRIBE messages (3) can be much longer than the Re-Registration time. The main issue for subscriptions is that they should be restarted when a device moves and is restarted (if stopped) when a Registration is successful.

Session Timer(4) is set to a large value, and allows for calls to enter into Survival Mode (see [Advanced Survival Mode](#)). If a site is prone to losing resources, you may set this timer to a smaller value for faster recovery of resources.

Peer heartbeat check (5) is the frequency with which the status of other elements in the cluster is checked to determine whether they are in service.

2.2.5 System Support for Resilient Devices

Call Routing

Resilient call routing handles cases in which a phone is in service on its secondary system or in transition between MiVoice Business system s. All Release 4.0 and later systems in a cluster are aware of the primary and secondary systems associated with each IP phone and are able to route calls to these phones when they are in service on either their primary or secondary systems.

Single Point of Provisioning

The MiVoice Business System Administration Tool, together with SDS, is used to program primary and secondary system s, enabling you to consistently and efficiently configure users and devices on their primary and secondary systems.

Data Distribution

SDS distributes and synchronizes telephone directory data throughout the cluster.

Features

While in service on a secondary system, IP phones retain basic call service. Most call features are available during Fail-over, some with possible behavior differences (see [Feature Resiliency](#) on page 48).

2.3 Call Resiliency

Call resiliency consists of two system capabilities:

- Call survival
- Resilient call clear-down

All IP phones are capable of call resiliency, regardless of whether they are provisioned as resilient. (That is, if an MiVoice Business system or the link between the system and its associated IP phones should fail, any active calls on IP phones survive whether the phones are programmed as resilient or not.)

Constraints

- IP phone connections to trunk or TDM devices through a failed MiVoice Business system are not maintained.
- Active calls routed through the embedded trunk module on a failed controller are not maintained.

2.3.1 Call Survival

Call survival is the process of keeping active calls alive when a device involved in an active call loses contact with its MiVoice Business system. A device does not have to be resilient to experience call survival. Once a resilient call is ended, a non-resilient device goes out of service, but a resilient device fails over to its secondary system. During a MiVoice Business system or network failure, both non-resilient and resilient devices experience call survival, but because they have lost the link to their system (call controller) they cannot access phone features or dialing functionality.

Only calls in talk state are deemed to be priority streams and are given call survival treatment by IP devices, in a failure situation. Calls that are in the process of being set up or are in feature transition, for example, do not experience call survival.

2.3.2 Phone User Interface

If a failure affects a device's current MiVoice Business system during a call, the user hears two quick beeps indicating that

- The phone's primary system has failed, or the link between the phone and the system has failed.
- The current call will be maintained, but for the remaining duration of the call, the user is unable to use the keypad or any phone features.

Once the user ends the call, the IP phone will home to its secondary system, at which time basic call and feature functionality is restored (see [Feature Resiliency](#) on page 48).

If a failure occurs while a phone is idle, the phone immediately homes to its secondary system and retains basic call functionality and most feature functionality.

2.3.3 Resilient Call Cleardown

Resilient call cleardown is the process of cleanly clearing down resilient calls when they are ended by either party.

Once a resilient call ends

- A resilient phone rehomes to its secondary system and retains service.
- A non-resilient phone goes out of service.

During a failure, MiVoice Business system s do not clear down calls in progress, and the Fail-over/Fail-back of any phone hosting a call is deferred until the call ends and the phone is idle.

If a resilient call is ended at the controllerless endpoint, the controlled endpoint automatically goes into idle state (dialtone); that is, no user intervention is required to clear down the call at the controlled endpoint. The end of stream detection delay (refer to the Controller Registry form in the MiVoice Business *System Administration Help*) handles the call cleardown. This behavior applies to all Mitel endpoint types, including resilient calls terminating on

- MiVoice IP Phones (all 50xx, 51xx, 52xx)
- Mitel IP Consoles (5540 and MiVoice Business Console)
- 3300 ICP E2T-trunking connections (all types)
- 3300 ICP E2T-DNIC connections
- 3300 ICP E2T-ONS connections

2.4 Feature Resiliency

This section provides information about feature availability and interaction during resilient operation, including

- Feature handling within call processing during resilient operation

- Accessing features locally on the secondary system
- Resilient feature support
 - unsupported features
 - special behavior, constraints, special programming
 - interactions of features with resilient routing
- Administrative feature interactions with resilient routing

2.4.1 Feature Handling

When you invoke a feature, call processing performs a local extension status check on the local device. If it is successful, the feature is given the local destination. In the event of a failure, the feature is invoked if it supports a remote directory number (RDN). The following features check the local device status:

- Direct page
- Remote Retrieve
- Directed Pickup
- TAFAS (only if the RDN record points to a device in the local resource software ID)

Feature handling invoked on the secondary system results in remote directory number (RDN) translation and virtual call routing to the primary system. Feature handling will only be performed on the secondary system if the primary system is out of service or the link to it is down.

Note:

You can use superkey sessions, feature access codes, and feature access keys to access and handle features locally, on the secondary system.

For information about using Speed Dial keys to access voice mail, see [Accessing Voice Mail](#).

2.4.2 Resilient Feature Support

When a device fails over to its secondary controller, the secondary controller supports most MiVoice Business features for the device, although some features have constraints and others require special programming or routing.

Summary of Feature Support on Secondary

The below table summarizes the features that are available on a resilient device that has failed over to its secondary controller.

Table 3: Features Supported on Secondary

Feature	Feature Supported on Secondary	Conditions
Account Code – Non-Verified	Yes	See Account code (Non-Verified) for details.
Account Code – Verified	Yes	See Account Code - Verified for details.
ACD Express Groups and Agents	Partial	See ACD Express Groups and Agents for details.
ACD Agent Request Help	Yes	Supervisor and agent must have same secondary controller. If the primary controller fails, any active Request Help calls, Silent Monitor sessions, or Record-a-Calls that are in progress are dropped. Both the initial call and the Request Help call, Silent Monitor session, or Record-a-Call are dropped.
ACD Silent Monitor	Yes	
Advisory Messages	Partially	See Advisory Messages for details.
Alpha Tagging	Yes	See Alpha Tagging for details.

Feature	Feature Supported on Secondary	Conditions
Attendant Cancel Call Forwarding	Yes	See Attendant Cancel Call Forwarding for details.
Auto Answer	Yes	Auto Answer feature status is shared between primary and secondary controllers in 3300 Release 7.0 and later.
Background Music	Partially	See Background Music for details.
Bandwidth Management	No	Not supported
Busy Lamp Field (BLF)	Partially	Monitored and monitoring devices can be resilient. See Busy Lamp Field (BLF) for details. For information about programming BLF keys, see “Programming Clustered Resilient Busy Lamp Field Keys” .
Call Logs (Call History)	Partially	See Call Logs (Call History) for details.
Call Forward	Yes	See Call Forward for details.
Call Forward - End Chaining	Yes	None

Feature	Feature Supported on Secondary	Conditions
Call Forward Override	Yes	Feature is selectable and functions correctly on the secondary controller.
Call Hold	Partially	See Call Hold for details.
Call Park – Remote Retrieve	Yes	See Call Park – Remote Retrieve for details.
Call Pickup – Clustered	Yes	See Call Pickup – Clustered for details.
Call Pickup – Dialed	Yes	Must be configured as clustered call pickup. See Call Pickup – Dialed for details.
Call Pickup – Directed	Yes	Functions correctly on secondary controller.
Call Privacy	Yes	Functions correctly on secondary controller.
Call Transfer to Ringing	Yes	When a controlling party transfers its held party to its other ringing party, any or all of the devices involved can be on their secondary controller without affecting the transfer.
Call Transfer to Talk	Yes	When a controlling party transfers its held party to the other party it is in conversation with, any or all of the devices involved

Feature	Feature Supported on Secondary	Conditions
		can be on their secondary controller without affecting the transfer.
Callback	No	Not supported.
Camp on	Partially	See Camp on for details.
Cancel Call Forwarding - End Chaining	Yes	None
Cancel Callback and Cancel Individual Callback	Yes	None.
Clear All Features; Remote Clear All Features	Yes	None.
Clustered Hospitality	No	Resiliency support in a hospitality application is limited to devices only; guest services (wake-up calls, room status information, suite services etc.) are not resilient.
Conference Call	Yes	See Conference Call for details.
Conference Call Split	Yes	Feature can be invoked on the secondary controller for conferences that were set up from a device on that controller.
Dialed Day/Night Service	Yes	Feature can be invoked on the secondary controller.

Feature	Feature Supported on Secondary	Conditions
		When you invoke this feature on the secondary controller, it does not affect the Day/Night Service status of the primary controller.
Direct Page	Yes	See Direct Page for details.
Do Not Disturb	Yes	Do Not Disturb status is shared between the primary and secondary controller by SDS.
Do Not Disturb – Remote	Yes	See Do Not Disturb – Remote for details.
Direct Station Select (DSS)	Yes	DSS can be invoked from a device on its primary or secondary MiVoice Business system. The information obtained by this feature is dependent upon the data it receives from the BLF feature (see Programming Clustered Resilient Busy Lamp Field Keys).
DPNSS Feature Support	Yes	See DPNSS Feature Support for details.
Feature Access Codes	Yes	See Feature Access Codes for details.
Feature Access Keys	Yes	See Feature Access Keys for details.

Feature	Feature Supported on Secondary	Conditions
Embedded Unified Messaging	Yes	Users continue to receive new voice messages but they are not synchronized until the user's device is back to the primary controller. A synchronization can still occur with a resilient controller as long as the Embedded Unified Messaging settings are also configured (including user IMAP login) on the controller. In this case, both the primary and resilient controllers are logging into the users inbox but they are interested in a different set of VM messages because the message-ID formatting includes the unique system ID of the controller.
Emergency Call Notification	Yes	Support for Emergency Services is provided on the secondary controller.
External Hot Desk User and External Hot Desk Agent	Yes	External Hot Desk Users and Agents support resiliency in the same way as other resilient devices. The resiliency mechanism used is the IP trunk heartbeat mechanism that other features, such as ACD and ring groups, use.

Feature	Feature Supported on Secondary	Conditions
Flexible Answer Point	Partially	See Flexible Answer Point for details.
Float Keys	Yes	Primary and secondary must both support the Float Key feature; otherwise, the Float Key will not appear if the two sets fail over and rehome at different times. To illustrate: Set (A) has a float key to answer calls to Set B. The sets fail over but Set A rehomes before Set B. While Set B is on its secondary (which does not support the Float Key feature), calls to it won't ring Set A nor will Set A's float key for Set B appear.
Force and Release Party Release	Yes	Users can invoke Release and Force Party Release against the other party in a two party call. These features apply on a per call basis.
Group Listen	Yes	See Group Listen for details.
Group Page and Meet Me Answer	Partially	See Group Page and Meet Me Answer for details.
Group Park	Yes	A resilient device programmed as a backup on the current switch but not hosted by that switch can be used as the park

Feature	Feature Supported on Secondary	Conditions
		destination DN. However, when the Group Park FAK is pressed, the call is parked remotely on the resilient device's hosting switch and the Group Park FAK on the current switch does not flash. If there is a Group Park FAK programmed on the resilient device's hosting switch, then the Group Park FAK will flash.
Group Presence	Yes	See page Group Presence for details.
Handoff	Yes	See Handoff for details.
HCI/CTI Application	Yes	HCI/CTI application is single system oriented. See HCI/CTI Application for details
Headset Operation	Yes	The on/off status of the headset feature is shared between the primary and secondary controller by SDS.
Hold on Hold	Partially	Both parties of a two-party call can put the call on Hold. The call appears on Hold to both parties. The line indicator associated with the original host of the call or the Hold/Retrieve LED blinks. Hold on Hold can take place over an

Feature	Feature Supported on Secondary	Conditions
		MSDN/DPNSS network or cluster. Resiliency support for Hold on Hold is the same as Call Hold. See Hold on Hold for details.
Hot Desk ACD Agent Login	Yes	Agent login status is shared between the primary and secondary controllers in 3300 Release 7.0 and later. If a hot desk ACD agent fails over to the secondary controller and then logs out of the agent session, the phone can take up to 45 seconds before it will return to service. After the phone returns to service, LOGGED OUT appears in the phone display.
Hot Desk ACD Agent Logout	Yes	
Hot Desking (Resilient Clustered)	Yes	You can provision hot desk directory numbers to be resilient. However the hot desk user must log into a set that has hot desking and resiliency provisioned. See Hot Desking (Resilient Clustered) for details.
Hotel/Motel (Hospitality) Features	No	Resilient Clustered Hospitality is not supported.
Hunt Groups	Yes	See Hunt Groups for details.

Feature	Feature Supported on Secondary	Conditions
Individual Trunk Access	Yes	Users can invoke this feature on the secondary controller; however, the trunk number specified refers to, and selects, a trunk on the secondary controller. Note: On an EX platform, individual trunk access is not supported because trunk numbers are not available when they are being used with the feature access code.
Language Change	Partially	A device's language setting is not shared to the resilient peer. See Language Change for details.
Last Group Member Routing	No	See Ring Group Resiliency
Last Number Redial	Partially	The number stored by the last number redial feature is not shared between resilient controllers. Last

Feature	Feature Supported on Secondary	Conditions
		numbers dialed are lost after a system reset.
Loudspeaker Paging	Partially	See Loudspeaker Paging for details .
Make Busy	Yes	SDS shares the status of the Make Busy feature between the primary and secondary controllers.
Message Waiting Indication (MWI) (on personal keys)	Yes	While a phone is being supported on its secondary controller, callers can leave messages on personal keys that have been programmed as message waiting indicators. However, after the primary controller returns to service, the phone will not rehome to its primary controller until the user clears all the message waiting indications from the phone.

Feature	Feature Supported on Secondary	Conditions
		i Note: However, that hot desk phones will rehome to the primary controller while there are flashing messaging waiting indications on the hot desk phone's personal keys.
Meet-Me Conference	No	If the MiVoice Business system hosting the conference fails, the conference terminates. If a conference party is on another system and it fails, the party is disconnected from the conference and must dial in again after the set has failed over to its secondary.
Multi-device User Groups	Yes	The non-prime member of an External Twin group must be primary on the same node as the prime member. No additional user or other licensing is required to support resiliency.
Multi-Level Precedence and Preemption (MLPP)	Yes	None

Feature	Feature Supported on Secondary	Conditions
Music On Hold	No	Music On Hold is not resilient. If an IP device is receiving MOH and the controller (primary, secondary, or other) that is supplying the MOH fails, the MOH audio will be disconnected.
Name Suppression on Outgoing Trunk Call	Yes	This feature works on a call by call basis. Users of a resilient device can invoke this feature while on the secondary controller.
Override (Intrude)	Yes	Users of a resilient device can invoke this feature while on the secondary controller.
Personal Keys	Yes	See Personal Keys for details.
Personal Ring Groups and Multi-device User Groups	Yes	See Personal Ring Group and Multi-device User Group Resiliency on page 93 for details.
Private Line Automatic Ringdown	Yes	The digital link upon which PLAR trunking runs can be resilient.
CDE Speed Calls	Yes	CDE Speed Calls are programmed in the Multiline Set Keys form. The data in this form is

Feature	Feature Supported on Secondary	Conditions
		shared by SDS with the resilient peer.
User Speed Calls	Yes	Users assign speed calls to keys (buttons) on their sets. Speed calls programmed into keys are available on both the primary and secondary controller.
Personal Speed Calls	Partially	See Personal Speed Calls for details.
System Speed Calls	Yes	System Speed Calls are available for use from all phones on the system.
Phonebook	Yes	The telephone directories of the elements in the cluster must be kept in synchronization through SDS to ensure that the Phonebook directory on the primary and secondary controllers will contain all the same entries.
Phone Lock	Yes	See Phone Lock for details.
Queue Status	Yes	Queue status information obtained on a controller pertains only to the wait queue status local to that controller: Resiliency support for ACD II supports the re-queuing of queued calls on the secondary

Feature	Feature Supported on Secondary	Conditions
		controller in case of a primary controller failure. In this case, the feature reports the status of the queue including the queued calls.
Record-A-Call	Yes	Record-A-Call uses the user's voice mail box. It is therefore subject to all the constraints of voice mail use on the backup controller. See Voice Mail for details . Record-A-Call fails if the controller hosting any of the three conference members fails.
Remote Clear All Features	No	This feature is not allowed on the secondary controller.
Ring Groups	Yes	See Ring Group Resiliency
Swap / Trade	Yes	This feature operates on a call-by-call basis.
Tag Call	Yes	This feature operates on a call-by-call basis.
Timed Reminder	Partially	See Timed Reminder for details.
Tone Demonstration	Yes	None.

Feature	Feature Supported on Secondary	Conditions
Travelling Class Marks	Yes	None.
Triple Ring Callback – Setup	Partially	See Triple Ring Callback – Setup for details.
Trunk Answer From Any Station	Yes	None.
Trunk Single/Double Flash	Yes	None.
Two B-Channel Transfer	No	Not applicable

Feature Conditions and Constraints

The following list identifies any conditions, constraints, or special programming requirements that apply to feature operation when a resilient device is on its secondary controller.

- **Account Codes (Non-Verified):** This feature is supported on the secondary controller. Users can activate this feature from their phones by using a feature access code, a feature access key, or the Superkey. The following conditions apply:
 - If non-verified account codes and SMDR printing are configured for the system, it generates an SMDR record for every non-verified account code that users enter.
 - Depending on the type of system failure, an SMDR may or may not be generated for the non-verified account code. If the controller loses communication with the set(s) due to a network failure, then the SMDR record is still generated on that controller. If the controller that would normally generate the SMDR record fails, then the record is lost.

- **Account Codes (Verified):** This feature is supported on the secondary controller. Users can activate this feature from their phones by using a feature access code, a feature access key, or the Superkey. The following configuration dependencies apply:
 - Verified Account Codes allow users to dynamically select a new COS/COR on a call-per-call basis to gain access to features that are not normally available under the user's configured COS/COR.
 - In resilient scenarios, physical configurations may limit the transparency of dynamically selected service profiles. These issues arise mostly in the context of dialing network destinations. Consider, for instance, the following example:




Figure 5: Dynamically Selected Service Profile

- Extensions 107100 and 107200 are on their primary iPBX 107. Both extensions can make local PSTN calls using their default COR 5. When either extension needs to make a long distance call, it dials a verified account code to acquire COR 6.
- If Extension 107200 fails over to iPBX 106, it can still make local PSTN calls using COR 5 and long distance calls using COR 6. The same may not be true for Extension 107100 if it fails over to iPBX 225 that has no direct PSTN access. The IP trunk from iPBX 225 iPBX 106 may not allow calls to PSTN via DPNSS (for example, COS 3). Even if this problem is corrected by proper configuration of the trunk service parameters (for example to COS 5), when the user selects COR 6 to make a long distance call, the new COR cannot be transmitted to the

IP trunk remote end point. Thus, iPBX 106 may deny the call based on the trunk's programmed COR.

- You can solve the particular problem in this example by setting up a route list to allow the selection of an appropriately configured IP trunk based on the caller's service profile (see [Dynamically Selected Service Profile](#)).




Figure 6: Dynamically Selected Service Profile (Solution)

- Advisory Messages:** This feature is partially supported on the secondary controller.
 - Users can access this feature by using Superkey on the secondary controller.
 - The set's Advisory Message status (on/off) and the selected message are not shared by SDS. Therefore, after the set fails over from the primary to the secondary, the user must select and activate the message again.
 - If the user selects an Advisory Message while the set is on the primary controller, the message will not be automatically selected when the set fails over to the secondary controller. However, when the set re-homes to the primary, the Advisory Message selection that was originally set on the primary controller prior to the failure will be recovered. The reverse occurs if the selection is made on the secondary controller but not on the primary controller.
 - Advisory Messages are language dependent. Users select the language of the set display by using the Superkey on the set. Ensure that you program the primary controller and secondary controllers to default the set displays to the same language. You program the default set display language with the LANGUAGE SELECT maintenance command.
- Alpha Tagging:** Alpha Tagging associates names with external numbers entered in the system telephone directory. When a call from an "alpha tagged" number is received on a display device, the user sees the associated name and number. Alpha tagging is supported on display sets that have failed over to their secondary controller.
 - SDS is responsible for synchronizing the alpha tagging entries in the telephone directories of the controllers.
 - The alpha tag lookup is performed on the terminating switch. If alpha tagging is enabled on the secondary controller and the database entry is programmed on that

controller, then a display set which has failed over will display the alpha tag name display.

- **Background Music:** This feature is only partially supported on the secondary controller. The following support is provided:
 - Users can access this feature when their set is on the secondary controller by pressing a feature access key or from the Superkey menu.
 - SDS shares the data that is programmed in the System Access Points form for the music source between the resilient controller pairs. However, the Embedded Media Sources form is not shared, nor are the embedded media files.
 - The set's music on/off status is not shared; a user who is listening to background music from a set on the primary controller will not hear background music when the set fails over to its secondary controller. Users must manually turn on background music at each controller.
 - If a controller resets, background music is turned off on all sets. If a fail-over is caused by a failure of the primary controller, the background music of the set will be off when the set re-homes to the primary controller. However, if the fail-over is caused by a network failure not involving a reset of the primary controller, the background music of the set will be on when it re-homes back to its primary controller.
 - Embedded music sources are digital files and are not shared by SDS. If you want resilient sets to play the same music whether they are on their primary or secondary controller, you must install the same embedded background music files on both the primary and secondary controllers.
 - External music sources are E&M trunks; these are physical trunks and cannot be shared. A resilient pair configured with E&M music will play the same music only if the trunks are tuned to the same audio sources.
 - Sharing of the music source configuration in the System Access Points form is only possible if the appropriate physical configuration exists on both controllers. If, for instance, PLID 2/1/4/1 refers to an E&M circuit on the primary controller but PLID 2/1/4/1 refers to an LS circuit, or is not configured on the secondary controller, SDS data sharing fails.
- **Busy Lamp Field:** When a monitored device is configured on the primary or secondary system of a resilient device with a related busy lamp, the busy lamp will be updated only when monitored device and the busy lamp host device are in service on the same MiVoice Business system, with the exception of DND status changes that are initiated by third party sets.

When a monitored device is configured on a MiVoice Business system other than a related busy lamp host device's primary or secondary system, regardless of

Resiliency, the lamp will only be updated when the busy lamp host device is registered on its primary system

While a device is on the secondary controller, busy lamps will intermittently appear to miss events for random periods of time and will intermittently be out of synchronization with the monitored device.

- **Call Forward:** In 3300 Release 8.0 and later, call forwarding is fully supported for resilient devices.
 - Call forward programming and on/off status is maintained for resilient devices that fail over to the secondary or fail back to the primary.
 - Call forwarding feature access codes and call forward feature keys are supported on both the primary and secondary.
 - Call forward Superkey session is available to users of resilient devices on both the primary and secondary controller.

Ensure that you set the following options in the System Options form to the same values on both the primary and secondary controller:

- Number of Forward Hops
- Call Forwarding Always -- Line Status Indicator ON
- Feature Active Dial Tone -- Call Forwarding

These options affect call forwarding behavior, but they are not shared between the primary and secondary controller by SDS. If you do not set them the same, users may experience differences in the call forwarding behavior when their phones are hosted on the secondary controller.

In 3300 Release 7.0 and earlier, call forwarding is not supported on the secondary controller. Call forwarding feature access codes and the Superkey call forwarding session do not function while a device is hosted on its secondary controller. Any call forward settings that users have programmed on their phones is lost when the device fails over to the secondary. However, the call forward programming is restored after the device fails back to the primary.

- **Call Forward - Cancellation of Call Forwarding for Extensions:** In a non-resilient configuration, an attendant can cancel the call forwarding settings of devices that are registered on the same controller as the IP console.

In a resilient configuration with 3300 Release 8.0 software or later an attendant console can cancel call forwarding for any device that is programmed as a primary or secondary on the controller that the console is registered on. Note that the following behaviors apply:

- The attendant can cancel call forwarding from the device regardless of which controller the device is registered on and regardless of whether or not the console

is resilient. If Call Forward Profile sharing is enabled, SDS cancels the call forwarding on the resilient peer controller.

- The console can cancel call forwarding on the device even if it isn't registered with the same controller as the console. The device just has to be configured as primary or secondary on that controller.
- If the console fails over to its secondary it can cancel forwarding from ALL devices that are configured as either primary or secondary on that controller (in the Multiline Sets form). The devices don't have to be registered with the secondary controller, they just have to be configured as either primary or secondary on that controller.
- If the console cancels forwarding for a device that is hosted on another controller (that is, the console controller is the device's secondary, and the device is registered with its primary controller), the change is only propagated to the primary controller if Call Forward Profile sharing is enabled through SDS.
- **Call History Logs:** This feature logs up to 20 missed calls (external incoming and external outgoing) for the phone. This feature is partially supported on the secondary controller.
 - Users can access this feature when their set is on the secondary controller by pressing a feature access key or from the Superkey menu.
 - Call History record only reflects the activities of the set on the current controller. If the set is on the primary, Call History only shows activities that occurred on the primary controller. If the set is on the secondary, Call History only shows activities that occurred on the secondary controller.
 - Call History records survive a system reset. After a device re-homes to its primary controller, records for activities that occurred before the failure are available. However, activities that occurred on the secondary will not be available. Similarly, if a device fails over to its secondary, the device's Call History records will only include previous activities that occurred on the controller.
- **Call Hold:** The following conditions and constraints apply to this feature:
 - Call hold on the primary controller does not survive a resilient fail-over to the secondary.
 - A held party over MSDN/DPNSS is dropped if a loss of communication occurs with the holding party.
 - A holding party over MSDN/DPNSS is dropped if a loss of communication occurs with the held party.
 - Users can use the "Call Hold", "Hold Retrieve" and "Hold Remote Retrieve" features while their phones are on the secondary controller.
 - A resilient set on the secondary controller will not re-home to the primary controller while it is in a call that is on hold or while it has a call on hold.
- **Call Park:** The Call Park feature allows you to place a call in a special hold state. You, or someone else, can then retrieve the call from any extension in the system. After

parking the call, the system can automatically connect you to paging equipment so that you can announce the call to the requested party. The following constraints apply:

- Calls are parked at the destination directory number's primary system whenever possible. If the primary system is unreachable because it has failed or because of network issues, the call will be parked on the secondary system if it is available. When the primary system returns to service, calls parked on the secondary will remain on the secondary. This means that calls may be parked on both the primary and secondary for a short period of time.
- Call Park - Retrieve LEDs on the park destination device will only reflect the park queue status of the MiVoice Business system to which the device is connected. Thus, if a call is parked on the primary while the device is connected to the secondary, the Call Park - Retrieve LED will not be lit. If the call is parked on the secondary when the device is homed to the secondary, the LED will be lit.
- Call Park retrievals will follow resilient routing. Therefore, a call that is not retrieved on the primary will result in a retrieval attempt on the secondary without any additional actions by the user.
- While the call is parked, if the parker rehomes to another MiVoice Business system, the recall to the parker will follow resilient routing.
- Parked calls may be retrieved from anywhere in the cluster.
- If a call is parked at the parkee's secondary controller when the primary controller is down, after the primary comes up, all calls parked on the secondary cannot be retrieved. The parked calls time out and follow the timeout handling for the parkee.
- Parked calls are dropped if the controller that hosts the parked call fails.
- Parked calls are dropped if the transit node or the parker's controller fails.
- In order for the system to support resilient call park from an IP console, both the primary and secondary controllers for the IP console must have 3300 Release 7.0 or later software.
- Call Park - Retrieve LEDs on the park destination device will only reflect the park queue status of the MiVoice Business system to which the device is connected. Thus, if a call is parked on the primary while the device is connected to the secondary, the Call Park - Retrieve LED will not be lit. If the call is parked on the secondary when the device is homed to the secondary, the LED will be lit.
- **Call Pickup - Clustered:** This feature allows phones on other elements in a cluster to be members of a pickup group on a local controller. Resiliency is supported for the resilient members of a clustered pickup group. If resilient members in a clustered pickup group fail over to their secondary controller, they are able to pick up calls that ring the other resilient, non-resilient, and remote phones in their pickup group. The other resilient, non-resilient, and remote members in the clustered pickup group can also pick up the calls that ring resilient members who have failed over to their secondary controller.


You must enable System Data Synchronization (SDS) for the cluster. The SDS feature maintains the group information and group resiliency membership on the cluster elements. However, SDS does not share the group and member information with

all the cluster elements. SDS shares group and member information only to cluster elements that have at least one local member programmed in the pickup group.

If you program a directory number on a primary or secondary controller and you then add that directory number to a clustered call pickup group, SDS adds the number to the call pickup group on the resilient peer. If the group does not exist on the resilient peer, SDS will create the group and add the member.

[Clustered Pickup Group](#) shows an example an example of a clustered pickup group.

Figure 7: Clustered Pickup Group



Refer to the Clustered Call Pickup feature description in the System Tool Online help for programming instructions.

- **Camp-on:** Users can activate Camp-on from their sets on the primary controller by using a feature access code, feature access key, or softkey. This feature is only partially supported on the secondary controller:
 - A caller can camp on to a resilient device while the device is on its secondary controller.
 - A Camp-on set from a resilient device on its primary controller will be lost if the device fails over to its secondary controller:

If a busy device with call waiting enters resilient talk, it causes all queued Camp-on calls to clear down. The camped on caller receives “Out of Service” on the display and re-order tone.

If a resilient device is camped on a busy extension and then loses communication with its primary controller, the call clears down and the resilient device fails over to its secondary controller.

- A caller is camped-on to a busy resilient set that is on its secondary controller. If the resilient device rehomes to its primary, the waiting call clears down and the caller receives re-order tone.
- If a resilient device on its secondary controller is camped on a busy phone, the resilient device does not re-home to the primary controller until the camp-on is serviced.
- **Cancel Individual Callback and Cancel Callback:** Cancel Individual Callback allows users to cancel a specific callback message that they set from their phone. Cancel Callback allows users to cancel all triple ring callback messages that they set from their phone. The following constraints apply:
 - Users can only access the Cancel Callback and Cancel Individual Callback features by using a feature access code on the primary controller.
 - Users cannot set Triple Ring Callback while their phone is on its secondary controller. Therefore, Cancel Callback and Cancel Individual Callback are not applicable on the secondary controller.

- **Conference Call:** Users can set up a conference from a resilient device that is on its secondary controller by using a feature access code, hardkey, or softkey. The following conditions apply:
 - A resilient device that is on its secondary controller will not re-home to its primary controller while the device is in a conference call.
 - If a Conference Bridge disappears, the
 - conference clears down
 - conferenced resilient devices on the failed switch fail over to their secondary
 - conferenced non-resilient devices on the failed switch wait for communication
 - conferenced devices over MSDN/DPNSS receive an error display and reorder tone
 - conferenced PSTN trunks receive reorder tone from the CO.
 - If the controller of a resilient device that is a party in a conference fails, the resilient device is dropped from the conference. However, the other parties in the conference continue to be connected.
- **Day/Night Service:** Users can invoke the following features while their sets are on the secondary controller by entering a feature access code or by using the Superkey:
 - “Dialed Day/Night Service - Setup” to setup the Day/Night Status of the secondary controller.
 - “Dialed Day/Night Service – Inquire” to obtain the Day/Night Status of the secondary controller.
- **Direct Page:** The following conditions apply to this feature:
 - Users can invoke Direct Page while their phone is on the secondary controller or direct page a resilient device that is on its secondary controller.
 - If a controller fails while a Direct Page is in progress, the page clears down. If a resilient paging or paged device has the failed controller as Primary, it fails over to its backup. If the resilient paging device or paged device is not on the failed switch, it receives an error display and re-order tone.
 - A paging or a paged device on the secondary controller cannot re-home to the primary while engaged in the page.
- **Do Not Disturb - Remote:** Typically, this feature does not lend itself to resilient applications.
 - The Flexible Answer Point must be a local extension number. As such, resilient devices cannot be designated as Flexible Answer Points.
 - A resilient device on the secondary controller is able to invoke this feature.
 - The scope of the feature is limited to the controller on which it was invoked; if invoked on the primary controller, it defines Flexible Answer Points for calls processed by the primary controller; if invoked on the secondary controller, it defines Flexible Answer Points for calls processed by the secondary controller.

- **DPNSS Features:** The following DPNSS features are supported on resilient devices in service on their secondary system

Resilient DPNSS Feature

- Callback Setup
- Callback Cancel
- Callback from Callee
- Callback Message Setup
- Callback Message Cancellation
- Diversion Validation
- Message Waiting Notification Activation
- Message Waiting Notification Deactivation
- **Feature Access Codes:** Features that interact with the call control database in a device's primary controller are not available to users while their phone is in service on its secondary controller, with the exception of class of service (COS) which follows normal exception handling. If intercept handling is programmed, users are presented to the intercept destination; if no intercept destination is programmed, users receive reorder tone and see NOT ALLOWED on the display. The following feature access codes will be rejected at the secondary controller (feature access codes that do not appear in this list will function on the secondary):

Feature Access Codes Rejected at the Secondary Controller

- Hotel Room Monitor Listen
- Hotel Room Monitor Setup
- Hotel Room Personal Wakeup Call – Set
- Hotel Room Personal Wakeup Call – Cancel
- Hotel Room Status
- Hotel Room Wakeup Call From Guest Extension
- Hotel Room Wakeup Call From Guest Extension – Cancel
- **Feature Access Keys (Multiline Phones):** All feature access keys are available on a device while it is in service on its secondary controller.
- **Flexible Answer Points:** Flexible Answer Point allows station and console users to program a night answer point for their incoming trunk calls or intercept handling calls. Telephone calls to an extension number listed in either the Trunk Attributes form or the Intercept Handling form will respect Flexible Answer Point programming. Typically, this feature is not used in resilient configurations. The following conditions apply:
 - The Flexible Answer Point must be a local extension number, so resilient devices cannot be designated as Flexible Answer Points.
 - A resilient device on the secondary controller is able to invoke this feature.
 - The scope of the feature is limited to the controller on which it was invoked. When invoked on the primary controller, it defines Flexible Answer Points for calls

processed by the primary controller. When invoked on the secondary controller, it defines Flexible Answer Points for calls processed by the secondary controller.

- **Group Listen:** The current state of the Group Listen feature will be maintained when in resilient call state. Thus,
 - if the Group Listen feature is active during system failover, the call will remain in Group Listen mode until the call is terminated (pressing the feature access key assigned to Group Listen will not turn off the feature or change the state of the associated LED).
 - If the Group Listen feature is inactive when failover occurs, the Group Listen feature will remain off until the call is terminated (pressing the feature access key assigned to Group Listen will not enable the feature or change the state of the associated LED).
- **Group Page and Meet Me Answer:** Group Page is not a resilient feature and thus you should not try to configure page groups as resilient. SDS does not share the data in the Page Groups form. Users of resilient IP phones can invoke these features while their phones are on the secondary controller providing that
 - the controller that hosts the Page Group is operating, and
 - the phone can communicate with the controller that hosts the Page Group (a device on the secondary controller cannot Meet Me Answer a page that it received while the it was on the primary controller).
- **Group Presence:** This feature allows users to take themselves out of the ACD, hunt, ring, personal ring groups, or multi-device user groups in which they are members. While out of (or *Absent* from) the group, calls to it bypass the user's phone until he or she returns to the group. Group Presence is resilient but with the following conditions and constraints:
 - for agent group, ring group, personal ring group and multi-device ring group members to support the group's resiliency, their primary and secondary host controllers must align with the primary and secondary of the pilot number. If controllers are misaligned—that is, when the group member and its group pilot number are hosted by different controllers—the LED associated with the Group Presence feature key will fail to indicate the member's true presence status when the key is toggled. Similarly, when an ACD Agent's host controller is different from its Agent Skill Group(s) controller, Group Presence for the particular Agent Skill Group will not be updated when changed using the Group Presence feature access codes for Join All and Leave All ACD Groups.
 - for hunt groups (networked and non-networked), member host controllers do not have to align with the host controller of the group pilot number. Remote hunt group members can change their Group Presence by feature access codes or feature access keys, however LEDs associated with Group Presence feature access keys hosted on remote members' telephones will remain off, regardless of the members' Group Presence status.
- **HCI/CTI Applications:** The MiVoice Business HCI / CTI application is single-system oriented. A CTI application can receive call and feature status information only from

devices that are locally configured and registered to the connected MiVoice Business system. The CTI application can initiate, answer, and clear calls only through control of locally configured and registered devices. A resilient device would have to be HCI monitored on its primary and backup systems. In order to support resiliency, an HCI / CTI application would have to interface to each MiVoice Business system independently and the application would have to integrate the data and control for the resilient devices.

MiTAI allows applications to transparently monitor resilient IP phones which fail over to a different MiVoice Business system. To provide continued telephone service in the event of a MiVoice Business system failure, MiTAI supports resiliency (a secondary MiVoice Business system assumes task of continuing support for IP phones that lose support from their primary system). If a MiVoice Business system failure occurs, the IP phone and associated MiTAI monitor fail over to secondary system.

A new event (ResilientDeviceEvent) informs the MiTAI monitor on an IP Phone in the following scenarios:

- if the phone's primary system fails,
 - after service is restored by a successful failover to the secondary controller, and
 - when the primary registration becomes available again.
- **Handoff with Personal Ring Groups and Multi-device User Groups:** [Handoff](#)
 - **Hot Desking (Resilient Clustered):** The MiVoice Business clustered hot desking feature supports resiliency. It is recommended that all clustered hot desk users and devices be provisioned to be resilient. (Nodal Hot Desking does not support resiliency because resiliency requires clustering.) For more information about clustered hot desking, refer to the MiVoice Business *System Administration Tool Help*.

As part of the Resiliency solution, resilient clustered hot desking offers device, call, voice mail, and feature resiliency, as well as the following features related specifically to resiliency in hot desking:

- Hot desk user-profile resiliency
- Persistent login
- Hot desk login, logout, and remote logout on the primary and secondary controller.

With resilient clustered hot desking, a hot desk user can be provisioned to be resilient and to support persistent login. Persistent login means that an active hot desk login survives during failover and failback. Resilient hot desk users can log in and log out of a hot desk session as long as they are in service on either their primary or secondary controller and as long as the hot desk device that they are using is in service on either a primary or secondary controller.

Note:

1. If a hot desk agent fails over to the secondary controller and then logs out of the agent session, the phone can take up to 45 seconds before it will return to service. After the phone returns to service, LOGGED OUT appears in the phone display.
2. When a resilient hot desk device rehomes to a switch that has newer device firmware, the user (including hot desk ACD agents) is automatically logged out to allow the firmware upgrade to proceed.

To experience resilient hot desking behavior

- A user must be provisioned in the MiVoice Business System Administration Tool as a resilient user, and
- The user must be using a device that is provisioned in MiVoice Business System Administration Tool, both as a hot desk device and as a resilient device.

Resilient hot desk users are not required to purchase any additional user licenses other than those that are required for external hot desking (refer to the *MiVoice Business System Administration Tool Help*).

Resilient clustered hot desking supports E911 services. For more information, see [IP Phone Behavior During Resilient Mode Operation](#) on page 286, and also refer to the *MiVoice Business System Administration Tool Help*.

- **Language Change:** The following conditions and constraints apply:
 - A device's language selection is not shared by SDS. Consider a system that has the following language configuration:
 - English (default)
 - French (auxiliary 1)
 - Spanish (auxiliary 2)

If a user selects French (auxiliary 1) on the primary controller, the device's language selection will revert to English (default) on the secondary controller.
 - Users can change the display language of their sets while on their secondary controller. The device's language selection on the secondary controller, however, is not shared with the primary controller. For instance, a user can select English on

the primary controller and French on the secondary controller. This selection will persist every time the set fails over to its secondary and re-homes to its primary.

- A controller's language selection (default, auxiliary 1, and auxiliary 2) are specific to that controller and are not shared. When a resilient device fails over it uses, by default, the default language configured on the secondary. In addition, the device's language selection options are limited to one of the three languages configured on the backup switch.
- To support language resiliency in a cluster, all nodes in the cluster must have the same setting for their default, auxiliary 1 and auxiliary 2 languages.
- **Loudspeaker Paging:** Typically, this feature is not configured as resilient because the paging zones configured on the primary controller are unlikely to match the paging zones that configured on the secondary controller.
 - Loud Speaker Paging uses E&M paging systems which are physical, non-resilient, trunks. They pertain only to the controller on which they are configured.
 - The paging zones, defined in the Loud Speaker Paging form, are not shared by SDS.
 - Users can invoke Loud Speaker Paging on the secondary controller using a feature access code or feature access key. The announcement is made over the paging zones on the secondary controller.
- **Phone Lock:** Phone Lock locks a set preventing access to the majority of features, with the following exceptions: unlocking the set via a user PIN, Hot Desk Login and Logout support and Emergency Call Notification support. Phone Lock has no effect on incoming calls but restricts outgoing calls, with the following exceptions: calls to emergency trunk routes and local operators. As well, the majority of keys on the device are disabled, except the dial pad and volume keys.
 - The user's PIN and Locked status are resilient when a switch resets, a fail-over occurs, or software is upgraded.
 - The lock/unlock status, the Phone Lock COS, Phone Lock FAK and FAC must be shared with the secondary system (if applicable) via SDS. These values need to be shared whenever there is a change in them.
 - The Phone Lock FAK and FAC are accessible whether the set is on primary or secondary controller. The Phone Lock FAK and FAC are configurable through Management Interfaces whether the set is on primary or secondary controller.
 - User PINs can be programmed through Superkey both from the primary and secondary controllers
 - A Hot Desk user may lock a set like any other user. It should be noted that Hot Desk users are currently not resilient when a set reboots and thus a set cannot be considered secure by simply locking a Hot Desk user. The previous locked status of Registration DN is applicable. If the previous state (before hot desk user logged in) is unlocked, the set would get unlocked. Only the Hot Desk user's profile is secured by invoking Phone Lock.
 - Hot Desk Users are not persistent across set reboots in terms of staying logged in. This means that a set reboots as the Registration DN instead of the Hot Desk DN. If

a hot desk user is logged in to one set and then logs into a different set, the original set's locked status before the hot desk user logged in is restored.

- **Programmable Keys:** Users can assign features, such as a Do Not Disturb, to programmable keys on their phones. The programmed keys are shared by SDS. If a user programs a key on either the primary or secondary controller, the key will also be present when the phone is hosted on the resilient peer controller.

 Note:

As of 3300 Release 9.0, to ensure proper display of labels on sets that support user-programmable key labels, both resilient peers must have same system software.

- **Personal Speed Calls:** Users can store and dial frequently-used numbers by using access codes and index numbers. This feature is partially supported on the secondary controller:
 - A user can store, remove and invoke personal speedcall numbers using the appropriate FACs.
 - The personal speed call numbers are not shared. The numbers stored on the primary controller are only available on the primary controller, and those stored on the secondary controller are only available on the secondary controller.
 - Deletions of stored numbers on a controller are not distributed by SDS to the other controller.
- **Timed Reminder:** This feature allows users to program their sets to ring and provide a message at a specified time within the 24-hour period. The following conditions apply to this features in resilient configurations:
 - Users can program reminders using the Superkey when their sets are on the secondary controller.
 - Reminders that are programmed while the set is registered on the primary controller are not distributed to the secondary controller. Reminders that are programmed while the set is registered on the secondary controller are not distributed to the primary controller.
 - If a reminder matures on the primary controller while the set is registered with the secondary controller, the set is not notified. The reminder expires on the primary

controller. The functionality is the same if the reminder matures on the secondary while the set is on the primary controller.

- If a resilient device sets up a timed reminder on its primary controller then fails over to its backup switch, the following should be expected should the device re-home back prior to the reminder maturing:
 - If the device fails over because its primary controller reset, the reminder will be lost, and the set will not be notified.
 - If the device fails over without the primary controller resetting, the set will be notified when the reminder matures.
- **Resilient Monitored Device:** If a monitored device is resilient, its primary and secondary controllers can initiate busy lamp updates, if the feature is configured. When a monitored device comes into service, an update is sent, indicating that the device is idle. This occurs each time a device registers. An out of service (OOS) monitored device generates no call-related busy lamp updates. However, a busy lamp update is issued for an OOS monitored device if it is placed in or out of Do Not Disturb (DND), by another telephone.

Busy lamps configured on devices that are programmed on the same MiVoice Business system as a resilient monitoring device's primary or secondary system are only updated when the monitored device is registered on the same system, with the exception of DND status changes that are initiated by third party sets.

Busy lamps configured on devices that are programmed on MiVoice Business systems other than a monitored device's primary or secondary system will be stored as RDNs. These busy lamps are updated by the monitored device's primary and/or secondary system. If the monitored device's system fails, busy lamps on remote systems become unsynchronized with the monitored device's status until it rehomes.

- **Serial Call:** Serial calls route to the attendant console on the primary or secondary system. The system originating the serial recall may require changes to support resiliency.
- **Superkey Sessions:** The following types of superkey sessions are available to a phone that is in service on a secondary controller:
 - Emergency Call
 - Phonebook
 - Messaging
- **Transfer to Busy:** When a controlling party transfers its held party to a busy third party, any or all of the devices involved can be on their secondary controller without affecting the transfer. Transfer to Busy causes the controlling party to make a new call to the busy party requesting camp-on. If the new call successfully rings or camps on the busy party, then the controlling party transfers its held party to the ringing or waiting party. If the busy party switches to its alternate controller before the controlling party invokes transfer, the new call will receive basic call resilient routing support to reach the destination on its alternate controller.

- **Transfer Recall:** When a party is transferred to a ringing or busy party, the previously held party sets up a recall timer. If the ringing or busy party does not answer within a configurable time period, the previously held party releases its current call and recalls the controlling party. During the time period between transfer invocation and attempted recall, the controlling party may have switched to its alternate (primary or secondary) controller. In this case, when the originating party attempts to recall the controlling party, basic call resiliency support causes the recall to route to the controller party on its alternate controller.

If a resilient device performs an unsupervised transfer and then moves to its alternate controller and the previously held party recalls, the device will not receive indication that the call is a recall. To generate a recall indication, the waiting party's controller would have to signal DPNSS redirection in the ISRM.

- **Triple Ring Callback Setup:** Triple Ring Callback allows a user to request notification from the system when a busy extension or a busy external line becomes idle, or when an unanswered station goes off-hook and on-hook. This feature is only partially supported on the secondary controller:
 - Users can activate this feature when their set is on the secondary controller by pressing a feature access key or from the Superkey menu.
 - [Triple Ring Callback Setup](#) Support identifies the scenarios in which the triple ring callback feature is supported:

Table 4: Triple Ring Callback Setup Support

Calling Extension	Busy Extension		
	Non-Resilient Device	Device on Primary	Device on Secondary
Non-Resilient Device	Allowed	Allowed	Not Allowed
Device on Primary	Allowed	Allowed	Not Allowed
Device on Secondary	Not Allowed	Not Allowed	Not Allowed

2.4.3 Administrative Feature Interactions with Resilient Routing

This section describes how administrative features behave when they interact with resilient routing. For information about resilient call routing/ARS and resilient call forwarding, see [Resilient Call Routing](#).

Voice Mail

During a failure, resilient voice mail continues to be provided to a resilient phone while it is in service on its secondary system .

Two forms of voice mail can be resilient:

- Embedded voice mail (can be centralized)
- External voice mail (can be centralized)

Embedded Voice Mail

Embedded voice mail capabilities exist on each MiVoice Business system. Embedded voice mail can be configured as a local voice mail resource for the hosting MiVoice Business system or as a centralized, shared resource for the cluster. As a centralized embedded voice mail server, the MiVoice Business system functions like an external voice mail server, except that if the hosting system fails, all sets lose their voice mail capability until it recovers. When embedded voice mail is configured to be resilient for a resilient device, the device receives voice mail service from its secondary system if its primary system fails.

- **Message Notification and Retrieval:** When voice mail service is provided by a device's secondary system , messages are recorded on that secondary system, and there is no change to the message notification and retrieval process. Resilient devices are notified of messages in the usual way, by the Message Waiting Indicator. Users retrieve messages in the usual way, unless using speed-dial keys for this purpose. Upon failing back to the primary system, a device is notified of any messages waiting on the secondary voice mailbox. Since users may have messages waiting on both mailboxes (primary and secondary), their phone is notified of messages waiting on both.
- **Centralized Embedded Voice Mail:** Centralized embedded voice mail is hosted on a MiVoice Business system dedicated as a central voice mail server for the cluster. This server provides voice mail service to devices while on their primary or secondary system s. If this server fails, however, all associated devices lose their voice mail capability until it recovers. There is no change to the message notification and retrieval process for centralized embedded voice mail.

External Voice Mail

An external voice mail server provides voice messaging to devices on their primary and secondary systems. If the external voice mail server or the element to which it is attached fails, all associated devices also lose their voice mail capability until the server recovers.

External voice mail is hosted on a dedicated server that is connected to a network element. In most cases this element is a MiVoice Business system, but this element does not have to be a Mitel product. A cluster may contain several external voice mail servers or one as a centralized, shared resource for the cluster. The voice mail server continues to provide voice messaging to devices that are in service on their secondary system. However, if the external voice mail server or the element to which it is attached fails, all the devices that obtain voice mail service from the server also lose their voice mail capability until the server recovers.

- **Message Notification and Retrieval:** While in service on their secondary system, resilient devices can continue to obtain voice mail service from an embedded (or centralized embedded) voice mail server or from an external (or centralized external) voice mail server. Regardless of whether phones are in service on the primary or secondary system, users are notified of waiting messages in the same way, by the Message Waiting Indicator. Notification is attempted at both of a device's systems, and is received by the device through the system the device is in service on. Users retrieve messages in the usual way, unless using speed-dial keys for this purpose (see [Accessing Voice Mail](#)).
- **Centralized External Voice Mail:** Centralized external voice mail functions in the same way as external voice mail, with no difference to the message notification and retrieval process.

Accessing Voice Mail

You access voice mail messages left on the secondary system in the same way as on the primary. Voice mail services such as message waiting indication and message retrieval are available while an IP phone is in service on its secondary system. Any messages received while on the secondary continue to be indicated and retrievable when the IP phone returns to service on the primary system (see [Superkey Sessions](#), and [Voice Mail](#)).

If a user is using speed dial keys to access voice mail, two speed dial keys must be programmed: one for the primary system and one for the secondary system. The Message Waiting Indicator stops flashing when all messages have been retrieved from both systems.

Saving Voice Mail Messages

When messages have been saved, no indication is given about which voice mailbox they have been saved on. If using speed dial keys to access voice mail, a user may need to access both voice mailboxes to retrieve all saved messages.

Station Message Detail Recording

SMDR generates logs in a resilient system.

Internal SMDR

If both telephones in a call are resilient and both are registered on the same MiVoice Business system , then an internal SMDR record will be generated.

On a destination's primary or backup system , if a local call is resiliency re-routed to the alternate MiVoice Business system , it is not expected that an internal SMDR record will be generated. From an internal SMDR perspective the call was external.

With the appropriate SMDR configuration, a call from one local party to another local party that is OOS generates an internal SMDR record. If a party is forwarded on OOS, internal SMDR generates one record indicating the calling party, controller party and the nominated party and a forwarded call indicator.

External SMDR

On a destination's primary or backup system , if a local call is resiliency re-routed to the alternate MiVoice Business system , it is expected that an external SMDR record will be generated for the re-routed call.

If a call is originated from a system that is the neither primary nor backup system for the destination and resilient routing is attempted, it is expected two external SMDR records are generated, one for each call attempt. If the re-routed call is initiated by a boundary MiVoice Business system , it is expected the boundary system will generate two SMDR records and the originating system only one record.

Emergency Services

CESID numbers (emergency location reference data) are programmed against extensions. For emergency calls this CESID number supplements the ANI information sent to a CO. If an extension is resilient, this CESID number must be programmed the same on the extension's primary and backup MiVoice Business systems. The MiVoice Business System Administration Tool will support propagating this data between MiVoice Business system s.

In addition to the CESID number, the system can be configured to provide local notification of an emergency call. When configured, attendant consoles and emergency hunt group members are informed of the emergency call. A system alarm and maintenance log is generated.

If 911 ARS routing, attendant consoles and emergency hunt groups are configured on the primary and backup system , a resilient device will trigger the emergency services on which ever system the device is registered. The emergency service features of the primary and backup systems will operate independently. Each system must have its own

access to the CO, otherwise the resilient device may be unable to place the emergency call.

If the devices to be notified of an emergency call are resilient they will be notified only if they are registered on the same system as the trunk hosting the emergency call.

HCI CTI Application

The MiVoice Business HCI[®] / CTI application is single-system oriented. A CTI application can receive call and feature status information only from devices locally configured and registered to the connected MiVoice Business system. The CTI application can initiate, answer and clear calls only through control of locally configured and registered devices. A resilient device would have to be HCI monitored on its primary and backup systems. A HCI / CTI application would have to interface to each MiVoice Business system independently and the onus would be on the application to integrate the data and control for a resilient device.

Without additional support it is not expected that HCI make call functions would receive resilient routing treatment when the HCI interface resides on the resilient device's primary or backup system. If the HCI interface was on an originating MiVoice Business system, and an HCI initiated call transited a resilient boundary MiVoice Business system, then make call type functions could benefit resilient routing.

Hotel Motel Application

Resilient routing is applicable only to basic calls involving hotel stations.

The Hotel / Motel features of the MiVoice Business system are currently single-system oriented. The attendant console provides Hotel / Motel specific menus to enable the following types of features.

Most Hotel room features are also applicable to suites and linked suites.

Note:

Most Hotel room features are also applicable to Suites and Linked Suites features.

- **Display Guest Room Information:** Provides information about the guest room, the guest, and the room extension. Existing information can be changed or new information added to a room.
- **Check In and Check Out Guests:** Allows you to keep track of arriving and departing guests.

- **Change the Status of a Room:** Allows you to set the condition and occupancy status of a room.
- **Find Rooms:** Allows you to search for rooms by using the room condition and occupancy status as search parameters.
- **Monitor Rooms:** Allows an attendant console, line, or trunk to listen to a room monitor extension.
- **Set Automatic Wake-up Calls:** Allows you to set an automatic wake-up call for a room extension.
- **Enable Call Blocking:** Prevents calls from being made between guest rooms.
- **Apply Call Restrictions:** Restricts the type of calls that a guest can make from a room extension.
- **Use Message Registration:** Calculates the total cost of calls made from a room extension.
- **Print Hotel Reports:** Provides Automatic Wake-up, Room Status, and Message Registration reports.
- **Access System Logs:** Provides access to the logs generated by the system during operation of the Hotel/Motel feature package.

Most of these features involve configuring and querying data related to sets configured as hotel stations (COS Option). The ability to query and change hotel station data is limited to the data contained in a the CC database of single MiVoice Business system. The ability to restrict a station's ARS access and configure a wakeup calls is limited to local devices. The application does not manipulate data on other MiVoice Business systems. Hotel station data for resilient phones would be managed independently between primary and backup systems. The Hotel data, station restrictions and wakeups would only affect hotel stations when the devices are registered on the MiVoice Business system with the hotel data.

2.5 Line Appearance Resiliency

2.5.1 Description

In general, line appearances are resilient, meaning they follow the host device when it fails over and are fully functional while on the secondary controller.

After you assign a line appearance to a resilient device on the primary controller it too is made resilient by SDS which distributes the line appearance data (DN, host device, and resiliency pairing) to the secondary controller. The line appearance DN and its resiliency pairing (normally inherited from the host device) are distributed within the cluster via Remote Directory Number Synchronization. The directory numbers of the line appearances are then listed in the Remote Directory Numbers (RDN) form at each cluster element.

Observe the following rules and best practices when programming the line appearances to ensure predictable behaviour:

Best Practices

- Assign line appearances to sets with matching resiliency pairings—that is, to sets that have the same primary and secondary controller.
- Only share line appearance between systems running the same version of MiVoice Business software. Distribution by SDS of line key data between systems running different software versions will produce unexpected results.
- Use DSS/BLF keys not line appearances to monitor and answer calls on other prime DN.

Note:

- Upgrading to 3300 ICP Release 5.0 does NOT fix incorrectly programmed line appearances on resilient devices. Manual intervention is required. For example, before Release 5.0, enabling resiliency on a phone that had a multicall line appearance did not trigger SDS to distribute the appearance to the secondary controller. Upgrading to 5.0 and synching with SDS will not resolve the problem either; you must go to the key and execute a Change/Save (without actually changing the data) to get the key distributed to the secondary and the key's RDN updated.
- Phantom extensions (line appearances that are not configured) must not be used in resilient environment; resilient environments support only extensions that are programmed as line appearances or are on registered physical devices.

General Rules

Rule #1: The System Administration Tool allows programming of keys on resilient devices from the primary controller only.

Rule #2 : Line appearances shared among multiple devices must be hosted by devices on the same primary network element. (The shared lines constitute a multicall group which cannot be spread over multiple MiVoice Business nodes.)

Rule #3: Resiliency configuration and behaviour of key line and multicall line appearances are identical.

Rule #4: Resiliency configuration and behaviours are the same whether line appearances are on multi-line IP Phones or Attendant Console (as softkeys).

Rules for Adding Line Appearances

Rule #5: Resilient host data takes precedence over non-resilient host data.

When adding a line appearance to a resilient device, if the appearance is already assigned to non-resilient device, the line appearance data will be distributed to the resilient device's secondary controller.

Example:

- Two sets, DN 1000 and DN 2000 hosted on the same controller (A). Set 1000 is resilient (secondary, B), the other is not.
- Set 1000 is assigned a line appearance (DN 1111), which Set 2000 has already.
- The line appearance is distributed to Set 1000's secondary controller and in the process inherits its resiliency pairing (A,B).

Rule #6: Data distribution respects existing hosting when resiliency pairings conflict.

When adding a line appearance to a resilient set (Set A), if the line appearance is already assigned to another resilient set (Set B) with a different secondary than the first set (Set A), then the line appearance's existing hosting is preserved.

Example:

- Two resilient sets, DN 1000 and DN 2000, each with different resiliency pairings: A|B for Set 1000 and A|C for Set 2000.
- Set 1000 is assigned a line appearance (DN 3333), which Set 2000 has already.
- The line appearance's existing hosting (A|C) prevails and results in an SDS distribution error on controller A as the key cannot be programmed for DN 1000 on controller B. Controller C, meanwhile, generates the following error maintenance log :

"Cannot update record # 0 for <Multiline Set Keys - Programmable Keys> (The directory number cannot be configured as a key appearance because it has already been programmed in a remote node."

Rules for Clearing Line Appearances

Rule #7: The first set that is assigned a shared line appearance drives the resiliency pairing of that line, which matters when clearing line appearances on sets with different resilient pairings.

So, when an appearance is cleared, its hosting information will be changed to the first resilient member of the line group as listed in the Multiline Set Group form. All resilient members matching the new hosting will have the line appearance applied to their secondary controller.

If no resilient members are found, the first non-resilient member will determine the line appearance's hosting.

Example:

- Two resilient sets, DN 1000 and DN 2000.
- Set 1000's resilient pair is A|B; Set 2000's is A|C.
- Both sets share an appearance of DN 1111 with Set 2000 getting it first.
- When the line appearance is cleared from Set 2000, since there are no other sets with A|C resiliency pairing, the first resilient device in the line appearance's group now drives the hosting (in this case A|B).
- The line appearances are distributed to set 1000's secondary controller.
- The line appearance's hosting data is updated throughout the cluster to A|B.

Rules for Changing the Line Type or Directory Number of a Resilient Line Appearance

Rule #8: Changing a line appearance type from one that supports resiliency to one that does not is disallowed.

The User and Services Configuration and Multiline Set Keys form do not allow changing the line type field from a type that has hosting (i.e., an RDN; these include single line, key system, and multicall line types) to a type that does not have hosting (i.e., no RND; which is all other line types). The administrator is prompted to clear the key programming before assigning it to another type.

Rule #9 : Changing the directory number of a resilient line appearance is allowed except when it is seen as a single line on the secondary controller.

This rule is a subset of a more basic rule that applies to all systems, whether or not they are resilient: an appearance of a single line can only be cleared not changed to another DN. To illustrate this rule at work in a resilient application, consider a phone with a shared (hence multicall) line appearance of DN 1111. The phone's resiliency pair is A|B while it is A|C for the other phones with an appearance of DN 1111. From the perspective on first phone's secondary (B), DN 1111 is a single line since no other phones with an appearance DN 1111 have B as its secondary.

Rules for Enabling and Disabling Resiliency on Devices that Share Line Appearances

Rule #10: DO NOT enable resiliency on a device that shares a line appearance with another resilient device unless both devices also share the same resiliency pairing.

Otherwise, the line appearance won't function in fail-over mode. So Rule #10 is really about following Best Practice # 1.

Rule #11: Same logic as Rule #7 on clearing line appearances, which states that the first set that is assigned a shared line appearance drives the resiliency pairing of that line.

So when disabling resiliency on a set, the first resilient member in the line appearance's group (as listed in the Multiline Set Group form) determines the line appearance's hosting.

2.6 Hunt Group Resiliency

Description

In 3300 Release 7.0 software and later, you can configure voice hunt groups and voice mail hunt groups with resiliency. In a resilient configuration, if the primary controller is inaccessible, the calls to the hunt group on the primary are automatically redirected to the voice mail hunt group pilot number on the secondary controller. Calls are routed to the secondary controller until failback occurs.

After you define a resilient hunt group on the primary controller, the System Data Synchronization feature distributes the hunt group information to the secondary controller. The hunt group number and the directory numbers of the resilient local devices are distributed within the cluster either by Remote Directory Number Synchronization. The directory numbers or the resilient devices are then listed in the Remote Directory Numbers (RDN) form at each element. [Resilient Hunt Group Configuration](#) shows a resilient hunt group configuration:




Figure 8: Resilient Hunt Group Configuration

If the primary controller for a resilient hunt group fails, support for the hunt group is transferred to the secondary controller. The failure of the IP trunking link between the

primary and secondary controller initiates the failover to the secondary controller. Once the primary controller returns to service, the IP trunking link to the secondary controller is restored. After the secondary controller detects that the link is restored, support for the resilient hunt group fails back to the primary controller.

Group members fail over and fail back independently of the group and the group fails back to the primary before the members. Some members on the secondary controller may not fail back as quickly as other members.

When a resilient hunt group fails over or fails back:

- The hunt group member device responds the same as a normal resilient device.
- Hunting resumes at the start of the member list
- Calls that are queued to the hunt group are dropped if the controller fails
- Campons and Callbacks to the hunt group are lost.

2.6.1 Conditions

- Hunt group resiliency is supported on systems that have 3300 Release 7.0 or later software.
- Direct IP trunk connections are required between the primary and secondary controller. In a cluster that contains resilient hunt groups, all the controllers should be connected to each other by IP trunks (fully meshed cluster).
- System Data Synchronization (SDS) must be enabled. By default, SDS shares the “Hunt Group Assignment” and “Hunt Group Members” data between the primary and secondary controllers at the “Resilient Pair” scope.
- Hunt group resiliency is only supported for Voice, Voice Mail, Recorder, or NameTag hunt groups (that is, the hunt group must be assigned with type “Voice”, “VoiceMail”, “Recorder”, or “NameTag” in the Hunt Groups form). Currently, only NuPoint Unified Messaging Release 10.0 and later supports voice mail hunt group resiliency.
- It is recommended that you assign a resilient NameTag hunt group and its resilient group members to the same secondary controller. Failure to do so can result in incorrect Name Tag display information from being displayed on group member displays.
- NameTag hunt group resiliency should follow the same resiliency guidelines as a Voice hunt group.
- Resilient hunt groups assigned with type “Voice” can contain local, resilient, and network (remote) hunt group members.
- Resilient hunt groups assigned with type “VoiceMail” can contain local and resilient hunt group members. They cannot contain network (remote) hunt group members.
- To support Record-A-Call (RAC) resiliency, you must program the RAC ports as resilient. Only NuPoint Unified Messaging, Release 10.0 and later supports resiliency for RAC ports. In addition, the resilient voice mail hunt group, the resilient hunt

group that the RAC ports belong to, and the resilient sets that invoke RAC must be programmed with the same secondary controller.

- You cannot add members to a resilient hunt group from the Group Administration Tool. The Group Administration Tool only lists non-resilient hunt groups in the Add Extension and Edit Extension pages.
- You can place Campons or Callbacks to resilient hunt groups provided that the resilient hunt group is not also a network hunt group. The Group Campon and Callback features are not supported for network hunt groups.
- It is recommended that you assign a resilient voice mail hunt group and its resilient group members to the same secondary controller. You do not have to assign the members of a voice hunt group to the same secondary as the resilient voice hunt group. However, if resilient group members have the same secondary controller as the hunt group, the time required to identify an available group member is minimized.
- Resilient voice hunt groups can contain local and network members. You program network members as RDNs or system speed call numbers.
- If you delete an RDN from the Remote Directory Numbers form, the system also deletes it from any resilient hunt groups to which it belongs.
- If you delete a system speed call from the System Speed Calls form, the system also deletes it from any hunt groups to which it belongs.
- Remote Directory Number Synchronization must be used to support the provisioning of users with resiliency and to resolve any inconsistencies that occur across the cluster in the CEIDs of remote directory numbers.

2.6.2 Programming

Refer to [Configure Resilient Hunt Groups](#) on page 237 for instructions.

2.7 Personal Ring Group and Multi-device User Group Resiliency

2.7.1 Description

Both the Personal Ring Group (PRG) feature introduced in MCD Release 4.0 and the Multi-device User Group (MDUG) feature introduced in MCD Release 5.0 are resilient, allowing configuration data for the group to pass to the secondary controller should the primary controller fail.

Distribution of the PRG group pilot number and fail over and fail back behavior are as described in [Personal Ring Group Resiliency](#) for resilient Hunt Groups.

The following group-specific PRG and MDUG information is shared across the resilient pair if resiliency is enabled on the prime member:

- Personal Ring Group and Multi-device User Group number
- One Busy All Busy group option (PRGs only)
- Home Element (read only)
- Secondary Element (read only)

The prime member name is a read-only field. It is editable from the Telephone Directory form. This field is not directly shared by SDS when the Personal Ring Groups and Multi-device User Groups forms are distributed. The data is distributed using other forms.

The PRG and MDUG member data consists of the list of members of the group and their presence. This data is shared with the resilient secondary controller for the group. The PRG and MDUG member data is only shared with the co-hosted member.

Handoff

The Handoff feature can be invoked through Feature Access Keys (FAK), Feature Access Code (FAC) or Softkeys (on some sets).

The Handoff feature access key is resilient for each member in a resilient PRG/MDUG and shared with the peer controller.

User device resiliency in SDS propagates this key in the same manner as other keys configured in the Multiline Set Key form of the System Administration Tool.

The Handoff FAK will appear on the appropriate key while the device is hosted by the secondary controller in the same manner as the key appears when hosted on the primary controller. Pressing the Handoff key while hosted by the secondary controller will either Push a call to the group or Pull an in-progress call from another group member. Changes to a Handoff key while the device is hosted on either controller will be propagated to the resilient peer controller.

Handoff Push and Pull Conditions

The Personal Ring Group (PRG) and Multi-device User Group (MDUG) Handoff feature allows calls to be “Pushed” or “Pulled” from one member of a PRG/MDUG to another member of the same group. A Push passes the call to the PRG/MDUG for pickup by any other available member. A Pull joins an in-progress call with another PRG/MDUG member and any party can then release. The following conditions apply to resiliency for Push and Pull:

- A Pull attempt from another device during a failover will fail and return out of service treatment. The call for which the Pull was attempted will remain connected to the original caller.
- A Push attempt to the PRG/MDUG during a failover will also fail and the caller will remain connected to the group member that attempted the Push.
- A Pull attempt from another device during a failback will succeed only if the PRG/MDUG member being Pulled from is still on the secondary controller. If the member

is already back on the primary controller, then the Pull will fail and the caller will remain connected to the PRG/MDUG member that attempted the Pull. A Pull attempt succeeds if both PRG/MDUG members are hosted by the same secondary controller.

- A Push attempt to the PRG/MDUG during a failback will succeed. The Handoff to the PRG/MDUG could ring members on the secondary controller if the failover has not yet completed, or ring members on the primary controller if the failover has completed.

2.7.2 Conditions

- PRG resiliency is supported on systems that have MCD Release 4.0 or later software.
- MDUG resiliency is supported on systems that have MCD Release 5.0 or later software.
- It is strongly recommended that resiliency programming of all MDUG/PRG members is exactly the same to fully support all associated MDUG/PRG features
- A device can be programmed as a member of a PRG/MDUG if:
 - the device and group are both resilient and primary on the same system. If the device and group have the same secondary system, the member can be programmed on either the primary or the secondary system. If the device and group have different secondary systems, the member can be programmed on the primary system only.
 - for non-resilient devices, the group is resilient and primary on the hosting system for the device.
 - for resilient devices, the group is non-resilient and the group is hosted on the device's primary system.
- If PRG/MDUG members have been configured in such a way that they are on different controllers, and if a particular PRG/MDUG member has network issues with the primary server, then it will failover to the secondary server while the other PRG/MDUG members remain on the primary server.
- If PRG/MDUG members have been configured in such a way that they are on different controllers, and the primary server is rebooted - then all the associated PRG/MDUG devices fail over to the secondary server. When the primary server comes back into service, not all MDUG members failback to the primary server at the same time; some members may stay on the secondary server even after the primary server is fully active. The rate of failback depends on the number of MDUG members on the secondary server.
- If PRG/MDUG members have been configured in such a way that they are on different controllers, then an incoming call will cause devices on the corresponding MiVoice Business server only to ring.

2.7.3 Programming

Refer to [Configure Resilient Personal Ring Groups](#) on page 241.

2.8 Ring Group Resiliency

Description

The Ring Group feature introduced in 3300 Release 8.0 is resilient, allowing support for the ring group to pass to the secondary controller should the primary fail.

Distribution of ring group information and the ring group pilot number is as described on [Hunt Group Resiliency](#) for resilient hunt groups; likewise, the group member fail over and fail back behavior described on [Hunt Group Resiliency](#).

When a resilient Ring Group fails over or fails back:

- After a failover, all calls ringing or queued at a Ring Group are dropped by the failing controller.
- After a failback, all calls ringing a Ring Group on the secondary controller continue to ring with normal overflow handling applying for unanswered calls. If the overflow destination is resilient and available, then the call should overflow successfully.
- After a failback, all calls queued to a Ring Group remain queued on the secondary controller. Since all resilient IP phones will re-home to their primary controllers, no members will be available to answer the calls queued on the now-inactive backup controller. Thus, the call will be subject to overflow handling after the timer expires.

2.8.1 Conditions

- Ring Group resiliency is supported on systems that have 3300 Release 8.0 or later software.
- Direct IP trunk connections are required between the primary and secondary controller. In a cluster that contains resilient Ring Groups, all the controllers should be connected to each other by IP trunks (fully meshed cluster).
- System Data Synchronization (SDS) must be enabled. By default, SDS shares the “Ring Group Assignment” and “Ring Group Members” data between the primary and secondary controllers at the “Resilient Pair” scope.
- You must assign a resilient Ring Group pilot number and its resilient group members to the same primary and secondary controllers.
- Resilient Ring Groups can contain local, resilient and remote DN members.
- Remote Directory Number Synchronization must be used to support the provisioning of users with resiliency and to resolve any inconsistencies that occur across the cluster in the CEIDs of remote directory numbers.

2.8.2 Programming

Refer to [Configure Resilient Ring Groups](#) on page 242 for instructions.

2.9 T1/E1 Trunk Resiliency

Description

If a site's primary controller fails, this feature automatically transfers the support for a T1/E1 trunk from a T1/E1 Combo MMC (PN 50005160) in the primary controller to a T1/E1 MMC in the secondary controller.


On small resilient sites where a single T1/E1 trunk interface is connected only to the primary controller, the trunk is only available to users while their phones are supported by the primary controller. If the primary controller fails, users who fail over to the secondary controller cannot access the T1/E1 trunk. With T1/E1 trunk resiliency, the trunk is connected to both the primary and secondary controller. Then, if the primary controller fails, support for the T1/E1 trunk is transferred to the secondary controller allowing users on the secondary controller access to the external trunk.

The incoming PSTN trunk connects to the Input port on the T1/E1 Combo MMC in the primary controller. The Failover port on the T1/E1 Combo MMC connects to a T1/E1 MMC in the secondary controller. If the primary controller fails, a bypass relay on the T1/E1 Combo card automatically transfers support for the trunk to the T1/E1 MMC in the secondary controller. Any calls that are in progress on the trunk are dropped. After the primary controller returns to service, support for trunk is transferred back to the primary controller.

You can schedule a 1-hour time slot during which the trunk can fail back from the secondary by entering the “Programmed Failback” maintenance command on the primary controller. The system waits until all system circuits are idle before it transfers the trunk back to the primary controller, so that calls in progress are not affected. You also have the option to force the trunk to failback by entering the “EDT Forced Failback” maintenance command. When you force the trunk to failback, any calls that are in progress on the trunk are dropped. See “Controlling the Failover and Failback of Resilient Trunks” in the *3300 ICP Troubleshooting Guide* for details.


[T1/E1 Trunk Resiliency - Normal Operation](#) shows an example of a site configured with T1/E1 Trunk Resiliency in normal operation on the primary controller.

Figure 9: T1/E1 Trunk Resiliency - Normal Operation



T1/E1 Trunk Resiliency - Failover Operation shows an example of a site in which the primary controller has failed and the T1/E1 link is supported by the T1/E1 Dual Framers in the secondary controller.

Figure 10: T1/E1 Trunk Resiliency - Failover Operation



2.9.1 Conditions

Warning:


THE FAILOVER PORT ON THE T1/E1 COMBO MMC DOES NOT PROVIDE SECONDARY PROTECTION. ONLY CONNECT THE FAILOVER PORT TO A T1/E1 COMBO MMC OR DUAL T1/E1 COMBO MMC IN A SECONDARY CONTROLLER.

The following conditions apply to T1/E1 trunk resiliency:

- Both the primary and secondary controller require 3300 Release 7.0 or later software.

- Only the T1/E1 Combo MMC (PN 55005160) supports T1/E1 trunk resiliency. The T1/E1 Combo MMC supports several protocol variants. This resiliency feature is available for all the protocol variants.
- The T1/E1 Combo MMC is not supported in earlier versions of the controller (for example, 100-user controller).
- You cannot divide the T1/E1 channels between a primary and secondary controller (that is, you cannot configure half the channels resilient on the primary controller and the other half resilient on the secondary controller).
- To support trunk resiliency, you must configure the primary and secondary controllers in a resilient cluster.
- You must have IP networking connectivity between the primary and secondary controllers.
- The primary and secondary controllers must be connected by a dedicated trunk. To meet this requirement, use IP trunking between the primary and secondary controllers.
- If you install the T1/E1 Combo MMC in a CX that has pre- 3300 Release 7.0 software, trunk resiliency is not supported. The T1/E1 Combo MMC trunk will function like a trunk on the T1/E1 Dual Framer. MX and LX system with pre- 3300 Release 7.0 software, do not support the T1/E1 Combo MMC.
- By default, T1/E1 trunk resiliency is disabled. You must enable this feature through the System Administration Tool.
- You must configure the resilient T1/E1 link in the secondary controller with the same protocol and physical settings.
- The primary and secondary controllers must be located within 10 meters of each other. The cable that connects the resilient link port in the T1/E1 Combo MMC on the primary controller to the redundant port in the T1/E1 MMC on the secondary controller must not exceed 10 meters in length.
- You must configure the system to initiate trunk failback. You must enter the Programmed Failback maintenance command to set a failback schedule on the primary controller. You can also force the trunk to failback with the “EDT Forced Failback” maintenance command. Note that the trunk will only fail back after all the circuits on the secondary controller are idle.
- Do not daisy-chain T1/E1 Combo MMCs together (see [T1/E1 Trunk Resiliency - Unsupported Configuration](#)). This configuration is not supported.

Figure 11: T1/E1 Trunk Resiliency - Unsupported Configuration



2.9.2 Installation and Configuration

See [Configuring T1 E1 Trunk Resiliency](#) on page 232.

2.10 SIP Resiliency Support

To support Session Internet Protocol (SIP) resiliency in a satellite office, the MiVoice Business systems in the head office must be configured to support resiliency. See also [Resilient IP Device Behavior and User Interfaces](#) on page 28.

2.10.1 Head Office

Configure the MiVoice Business systems in the head office in a resilient cluster. Configure the phones in the satellite office as resilient devices with their primary and secondary controllers in the head office. In the event that their primary controller fails, the phones in the satellite office fail over to the secondary controller.

The primary Domain Name Server (DNS) server located in the head office provides a list of alternate IP addresses to the satellite office gateways when a gateway makes its first request to resolve a Fully Qualified Domain Name (FQDN). The first choice is the primary controller, the next choice is the secondary controller. If neither of the MiVoice Business systems are available, the third choice is the co-located gateway at the satellite office.

For a PRI connection through a digital trunk gateway or SIP gateway, you can configure gateway with two ports (two T1/E1 modules). Each port is connected to a different MiVoice Business system. The systems must be co-located in the same building, no further than 30 meters from the digital gateway. Program the digital gateway to use its hunt group feature to locate the backup MiVoice Business system in the event that the primary system cannot be reached. However, a simpler solution is to configure T1/E1 Trunk Resiliency (see [T1/E1 Trunk Resiliency](#) on page 97).

During a Wide Area Network (WAN) failure, the IP phones in the satellite office cannot make or receive calls. The gateways will remain open if they have a cache of alternate IP addresses. A secondary DNS server could be located at the satellite office to provide the functions of the primary DNS server, although the DNS configurations will be different.

You configure the controller in the head office to direct the calls destined for the satellite office by selecting trunk options through the Public Switched Telephone Network (PSTN).

During a local power failure, the IP phones in the satellite office will be unable to make or receive calls unless you have connected the LAN infrastructure, MiVoice Business system, and IP phones, to redundant power sources. It is recommended that the LAN, MiVoice Business systems, and IP phones be connected to an Uninterruptible Power Supply (UPS). You should also ensure that your Internet Service Provider (ISP) is configured with power redundancy.

Satellite Office MiVoice Business System

During a WAN failure, IP phones do not operate. Access to the MiVoice Business system through the Foreign Exchange Office (FXO) gateway is dependent on the gateway and its configuration. Similarly, access between the Foreign Exchange Station (FXS) gateway and FXO gateway depends on the type of gateways and their configurations.


During a LAN failure or power failure, the IP phones do not operate and the FXO gateway and FXS gateway will be unable to receive and make calls. It is recommended that

- the satellite office LAN, gateways, and IP phones be backed up with an UPS.
- the FXS gateway provide a bypass port to the PSTN to allow outgoing calls during a LAN or power failure.

2.11 Resilient System Elements and Interactions

To achieve device, call, and voice mail resiliency, you must set up two or more MiVoice Business systems in a resilient cluster, which is a specially configured network of MiVoice Business systems and associated IP devices that can direct phones, route calls, and maintain calls to each other. [Resilient-cluster Topology](#) illustrates the topology of a basic resilient cluster, in which IP phones have a primary and secondary system. For information about engineering a resilient network, see Appendix: [Engineering Guidelines for Resiliency](#).

Figure 12: Resilient-cluster Topology



2.11.1 About Resilient Clusters

A resilient cluster

- can be managed from the MiVoice Business System Administration Tool. To manage a cluster, each element in the cluster must be running MCD Release 4.0 or later software, and each must be migrated to support Remote Directory Number (RDN) Synchronization. In addition, all elements in the managed cluster must be MiVoice Business systems.
- must contain
 - A cluster of MiVoice Business systems
 - 5000-series IP Devices that support resiliency
 - Interfaces to PSTN trunks that are connected to MiVoice Business systems inside the cluster
- may contain
 - 3300 ICPs or MiVoice Business systems from a variety of Release levels (3300 ICP Release 3.x, 4.0, or later) with varying user capacities (100, 250, and 700 users). however, only 3300 ICP Release 4.0¹ or later systems can function as resilient (primary and secondary) call controllers.
 - Applications that behave like IP phones such as Mitel Unified Communicator (note that this application cannot be made resilient but can be part of a resilient cluster)
 - 5540 IP Console or MiVoice Business Console
 - Interfaces to 3300s outside of the cluster

The following restrictions apply to Resiliency in mixed clusters:

¹ Not to be confused with MCD Release 4.0 which came after 3300 ICP Release 4.0. See the note on page 5 for an explanation of product branding and release numbering.

- Only 3300 ICPs, Release 4.0 or later can function as primary or secondary ICPs.
- Only 3300 ICPs, Release 4.0 or later can function as resilient-call-routing nodes, or boundary nodes.

Note:

SX-2000 NT and 3200 ICP are not supported in a resilient environment.

2.11.2 Managing Resilient Devices from the MiVoice Business System Administration tool

In a network or cluster that supports RDN Synchronization, you can configure resilient devices from the MiVoice Business System Administration Tool. Prior to MCD Release 4.0, you must use OPS Manager to configure resilient devices.

From the User and Services Configuration form, you can

- Create or delete a resilient IP Phone
- Enable or disable resiliency for an existing IP Phone

From the IP Consoles form, you can

- Create or delete a resilient IP Console
- Enable or disable resiliency for an existing IP Console.

Figure 13: Embedded Resilient Device Support



Refer to the Embedded Resilient Device Support book in the System Administration Tool help for conditions and programming instructions,


2.11.3 Multiple Cluster Interaction

Resiliency introduces no change in cluster interaction. As depicted in [IP-trunked Resilient Cluster](#), all interactions between a resilient cluster and external MiVoice Business systems or other clusters take place through IP trunking or traditional TDM trunking.

Automatic route selection (ARS) handles call routing between clusters, and resilient call routing is invoked only when a call originates in or reaches a resilient cluster. For more information, see the following sections of this document:

- [About this Chapter](#) on page 167
- [Planning ARS Routes](#) on page 221
- [Planning ARS Routes for Boundary Nodes](#)
- [Programming the Cluster and ARS](#) on page 230

Figure 14: IP-trunked Resilient Cluster



2.11.4 Interactions Between System and External Applications

The following Mitel applications are external to the MiVoice Business system, but they can still be used in a resilient environment:

- Mitel Teleworker Solution
- 5810 PDA Application

For information about whether the applications support Resiliency, refer to the application-specific documentation.

All interactions between external applications and the Resiliency solution are restricted to use of the IP phone interface, that is, insofar as an application behaves like a resilient phone, it is treated as one by MiVoice Business system s in the network. From both a user and administrator perspective, the Mitel Teleworker Solution is somewhat different from Mitel IP devices, although both use IP Phone MiNET and streaming to interface with other system elements.

 Note:

The IP Console is considered to be an integral part of the resiliency solution and is discussed in this document.

This chapter contains the following sections:

- [ACD Resiliency Solution](#)
- [Resilient ACD Configurations](#)
- [Converting a Traditional ACD Site: An Example](#)
- [Resilient Virtual Contact Center Configurations](#)

3.1 ACD Resiliency Solution

3.1.1 Description

A resilient ACD configuration typically consists of a trunking gateway or gateways, a primary controller, and a secondary controller programmed within a cluster. The ACD paths are configured on the trunking gateways to direct calls to the resilient ACD agent skill groups that you have configured on both the primary and secondary controllers. During normal operation, the trunking gateways direct the ACD calls to the primary controller. If the primary controller fails, the calls are redirected to the secondary controller. The resilient agents and resilient agent skill groups fail over to the secondary controller and are able to process the incoming ACD calls.

Note:

Resiliency is not supported for ACD Express Groups, but a certain level of resiliency can be programmed. ACD Express Agents can be resilient if the Hot Desk Agent they are based on is resilient.

ACD paths are not resilient. However, you can achieve a level of ACD path resiliency by programming two paths with the same configuration information on separate controllers. You configure each path with a unique directory number within the cluster (see [Resilient ACD Configurations](#) for configuration examples).




Figure 15: Resilient ACD Configuration

Failover Behavior

If the primary controller fails, the trunking gateway is notified of the failure and any calls that are queued to the failed primary, but are not ringing an agent, are re-queued to the corresponding resilient group on the secondary. Any calls that are ringing an agent on the primary are requeued as the oldest call in the queue of the corresponding resilient agent skill group on the secondary controller.

Any ACD agents that are not on calls (that is, ACD agents who are logged in but idle, in Make Busy, or on Work Timer) rehome to their secondary controller. Calls in progress at ACD agents on the primary controller are maintained (call survival). However, the ACD agents are unable to access any features or dialing functions because the primary is out of service. After an agent ends the current call, the agent's ACD set fails over to the secondary controller. As agent members fail over to the secondary, they begin to answer the calls that were requeued by the trunking gateway to the agent skill groups on the secondary controller. Eventually, all ACD agents fail over to the secondary controller and are available to answer calls to the agent skill groups.

While on the secondary controller, ACD agents have access to the same ACD features that they would have on their primary controller. During a failover or failback, the state of an agent -- for example, Make Busy, Do Not Disturb, and so forth -- is maintained. If an agent's set is in Work Timer when the agent's set fails over or fails back, the Work Timer is cancelled.

During a failover, the agent skill group that is configured on the secondary performs the same functions as if it was still on the primary by

- servicing calls that are queued in the path
- presenting calls to the ACD agents
- recording group statistics for the gateway controller as if the failure of the primary controller had not occurred (group statistics are recorded on the secondary controller but the statistics are only collected for calls that arrive on the secondary controller).

The system attempts to maintain the call priority during failover to the secondary. However, depending on the network configuration, the system may initially present calls to the ACD agents out of order.

Failback Behavior

After the agent skill group's primary controller returns to service, the primary controller notifies the trunking gateway. Calls that are currently queued to the agent skill group on the secondary controller, but are not ringing a agent, are re-queued on the primary controller. Calls that are ringing an agent or are in progress are not affected. As ACD agents on the secondary controller become idle, they fail back to the primary and begin answering the re-queued calls on the primary controller. When all the ACD agents are back on the primary controller and all existing calls are queued on the primary, failback is complete.

3.1.2 About ACD Agent Hot Desking

When a traditional ACD agent logs into an ACD set, only the Class of Service (COS) and Class of Restriction (COR) that are associated with the agent's directory number are applied to the ACD set. In order to provide ACD agents with access to user features, such as personal speed calls, some call centers provide ACD agents with a separate non-ACD phone that the agent can then program with personal features. This separate phone can also be programmed with a different COS for Day, Night 1, and Night 2 service than those used by the ACD set.

The ACD Agent Hot Desking feature allows an agent to log into any hot desk enabled set or ACD set and the system re-registers the set with the agent's personal phone profile and ACD functionality. After the agent logs into the set, the agent has access to his or her own personal speed calls, features, and phone settings as well as the ACD agent functions. If you use hot desk ACD agents in a call center, you do not have to provide agents with separate phones for their personal use. Instead, you can make a pool of shared phones available to many agents and any hot desk enabled set or ACD enabled set that a hot desk ACD agent logs into will also function as the agent's personal phone.

After a hot desk ACD agent logs in, the system associates the user's personal phone settings, such as directory number, COS/COR settings, language display, and button programming with the set.

Note:

Only hot desk ACD agents and ACD Express Agents, which are based on hot desk agents, can be resilient. Traditional ACD agents cannot be resilient. Even if ACD agents do not require hot desking capability, they must be configured as hot desk ACD agents in order to be resilient. To convert a traditional agent to an hot desk ACD agent, you set the agent's User Type to "ACD Agent". See [Configure Hot Desk ACD Agents From System Administration Tool](#)

Hot Desk ACD Agent Session

After a hot desk ACD agent logs into a set, the system associates the user's personal phone settings, such as directory number, COS/COR settings, language display, and button programming with the set.

The system also continues to use the CESID (Customer Emergency Services ID) programmed for the set. For example, if an agent makes an emergency call from during an ACD session, the system sends the CESID associated with the set regardless of which profile (hot desk ACD agent or set user) is active on the phone. For local notification (SMDR log and attendant console), the system displays the directory number (DN) and name of the active profile (if available).

Note:

Compression zones are not applicable to hot desk users. Compression is not used when a hot desk user is logged in to a hot desk set because compression zones are associated with the registration DN of a physical device, not the hot desk DN.

After logging in, hot desk ACD agents can use the following features or change their programmed settings:

- Call forwarding (all types)
- Callback messages (message waiting indicator)
- Auto Answer
- Do Not Disturb (DND)
- Last Number Redial
- Timed Reminder
- Triple Ring Callbacks
- Advisory Status Message.

For example, after logging into a hot desk ACD agent session, an agent could press the Redial key and dial the last number that he or she had dialed previously from his or her personal phone profile.

When an agent is logged in, the hot desk agent's DN is in service and the set DN for the hot desk agent follows out of service handling. On logout, the agent's DN follows out of service handling. The ACD set's DN comes back into service and the set's profile, including feature settings, is reapplied. If required, incoming personal calls can be forwarded to the user's voice mail.

There are several options that allow hot desk ACD agents to receive direct (personal) calls on their ACD sets. Refer to ACD Agent Hot Desking in the System Administration Tool Help for details.

Hot Desk ACD Agent Devices

When programmed as hot desk enabled sets or ACD sets, the following devices support hot desk ACD agents:

- 5000 series Multiline IP Phones
- 5000 series Dual Mode IP Phones
- 5000 series IP Appliances
- 5312 IP Phone
- 5324 IP Phone
- 5320 IP Phone
- 5330 IP Phones
- 5340 IP Phones
- 5360 IP Phones
- 6905 IP Phones
- 6910 IP Phones
- 6920 IP Phones
- 6930 IP Phones
- 6940 IP Phones

3.1.3 Features Available to Resilient Hot Desk ACD Agents

ACD Supported Features

The following ACD features are available to resilient hot desk ACD agents who have failed over to their secondary controller:

- ACD Login and Logout
- Make Busy with Reason

- Work Timer
- Queue Status
- Programmable Thresholds
- Request Help
- Silent Monitor

In addition to these resilient ACD features, hot desk ACD agents can also use resilient set features, such as auto answer, account codes, do not disturb, and record a call while they are on their secondary controller.

Note:

Some resilient set features have constraints and may require special programming or routing. See [Resilient Feature Support](#) on page 49 for details.

ACD Feature Settings

When a resilient hot desk ACD agent fails over to the secondary controller or fails back to the primary controller, the states of the following ACD features are maintained:

- ACD Login and Logout status
- Make Busy On/Off status (reason code is not maintained).

In addition, when a resilient hot desk ACD agent fails over to the secondary controller or fails back to the primary controller, the settings of features such as auto answer, do not disturb, and the configured value for the Work Timer are maintained.

Note:

If an agent's set is in Work Timer when the agent's set fails over or fails back, the Work Timer is cancelled.

ACD Feature Access Codes

The following ACD feature access codes are available to resilient hot desk ACD agents who have failed over to their secondary controller:

- ACD Login and Logout
- ACD Make Busy - Setup
- ACD Make Busy Cancel
- ACD Silent Monitor.

In addition to these resilient ACD feature access codes, access codes for resilient features such as auto answer, account codes, and do not disturb are also available to agents while they are on their secondary controller.

3.1.4 Conditions

The following conditions apply to ACD Resiliency:

- 3300 Release 7.0 or later software is required on the gateway controllers, primary controllers, and secondary controllers to support hot desk ACD agents and ACD Resiliency.
- 3300 ICP Release 5.0 or later software is required on the gateway controllers, primary controllers, and secondary controllers to support ACD Path Resiliency.
- You must use either hot desk ACD agents or traditional ACD agents. The use of hot desk ACD agents and traditional ACD agents in the same ACD system, is not supported.
- The primary and secondary controllers must be clustered and connected to each other by IP trunks (fully meshed cluster).
- System Data Synchronization (SDS) must be sharing data between the primary and secondary controllers. See [Configure System Data Synchronization SDS](#).
- You must configure resilient ACD agents on the primary and secondary controllers. You then use the System Administration Tool to configure resilient agent skill groups. After you have configured the agents and agent skill groups, SDS automatically keeps the required user and device data synchronized.
- Ensure that you assign the agent skill group members to the same secondary controller as the agent skill group (that is, both the agent skill group members and the agent skill group to which the members are assigned must have the same secondary controller). If you do not assign the agent skill group members to the same secondary controller as the agent skill group, the agents will not receive ACD calls during a failure of the primary controller.
- You cannot copy a resilient agent skill group from an agent skill group's secondary controller. When you click **Copy**, if the agent skill group's primary element is not local, you will receive an error message. You can only copy a resilient agent skill group from its primary controller.

- If the primary controller fails, any active Request Help calls, Silent Monitor sessions, or Record-a-Calls that are in progress are dropped. Both the initial call and the Request Help call, Silent Monitor session, or Record-a-Call are dropped.
- Only hot desk ACD agents can be resilient. In the User and Services Configuration form, you configure an agent as a hot desk ACD agent by setting or converting the User Type to "ACD Agent". See [Configure Hot Desk ACD Agents From System Administration Tool](#). Traditional agents cannot be resilient.
- Hot desk ACD agents can only log into ACD sets.
- After a system reset, hot desk ACD agents are automatically logged back into the ACD system once the set re-establishes contact with the agent's primary or secondary controller. Hot desk ACD agents are not logged out unless the set loses power or the agent explicitly logs out.
- If a resilient agent logs in and the ACD set rehomes to the primary or secondary controller, there could be a delay of several seconds before login is achieved.
- If a hot desk ACD agent fails over to the secondary controller and then logs out of the agent session, the phone can take up to 45 seconds before it will return to service. After the phone returns to service, LOGGED OUT appears in the phone display.
- If a traditional ACD agent logs out during a call, the system generates the report log immediately (that is, during the call). For resilient agents and hot desk ACD agents, the system does not generate the report until the call ends.
- The Group Administration Tool identifies hot desk ACD agents as "Hot Desk" users.
- You cannot make ACD sets and agent IDs members of hunt groups.
- You can make hot desk ACD agent IDs members of pickup groups. However, traditional ACD agent IDs cannot be members of pickup groups.
- You can park calls against hot desk ACD agents and ACD sets (whether the ACD set is assigned to a traditional ACD agent ID or a hot desk ACD agent). However, you cannot park calls against traditional ACD agent IDs, ACD paths, or ACD agent skill groups.
- In the ACD real time events, the directory number and agent number for a hot desk ACD agent is recorded as the agent's ID.
- When an ACD call rings a remote agent, the caller hears the ringback tone instead of the path RAD or MOH.

 Note:

When an ACD call rings a local agent, the caller hears MOH or a RAD (instead of ringback tone) while the agent is being rung.

3.1.5 Agent Failover and Call Redirection

The **Queue Callers To Group When No Local Agents Are Logged In and Present** option in the ACD Agent Skill Groups form has no bearing on resiliency behavior; calls to a path will queue to its designated group regardless how this option is set if the following conditions are also true:

- If the **Group Contains a Resilient Logged In and Present Agent** option in the ACD Agent Skills Groups form is set to YES, or
- There is at least one local agent logged in and present, or
- There is at least one resilient agent logged in and present

For Releases earlier than MCD 5.0

After the failure of a primary controller,

- agents start to failover from their primary to their secondary controller, and
- the gateway controller(s) initiate the redirection of calls that are queued to the primary controller to queue on the secondary controller.

While the calls are being redirected, a situation can occur where the gateway attempts to queue calls on the secondary controller before the first agent becomes available in the resilient agent skill group on the secondary. In this case, the agents have not arrived on the secondary because they are engaged in existing calls (call survival) or they are in the process of failing over and registering on their secondary controller. Agents don't begin to fail over until their current calls are completed.

If the gateway attempts to queue calls on the secondary controller before the first agent becomes available in the resilient agent skill group, the gateway views the groups on the secondary as having no agents logged in during this transition period. You have two options for handling the redirection of the queued calls.

Note:

The same situation can occur when the primary controller returns to service and calls once again commence queuing at the primary before the first agent has failed back and registered in the agent skill group on the primary controller.

Option 1

In the ACD Agent Skill Groups form leave the option "Queue calls when no local agents are logged in" set to "No" for the resilient agent skill groups. Then, any calls that are redirected to the secondary controller prior to the first agent failing over are not queued to the agent skill groups because there are no agents logged in.

Thus, calls that are queued to the path on the gateway before the primary controller failed:

- first follow Overflow programming when they attempt to queue on the secondary when no agents are logged in, overflow immediately, and
- then follow Interflow handling when the Interflow Timer expires.

All new calls that attempt to queue to the path on the gateway after primary agent controller fails but before the first agent is available on the secondary controller:

- first follow overflow programming and overflow immediately without queueing to the group
- then follow Path Unavailable Answer Point handling.

Option 1 has the following advantages:

- preserves automatic Path Unavailable Answer Point handling when all agents log out at the end of the day.
- ensures previously queued calls will not be left unanswered through Interflow Handling. The calls will eventually route to the interflow point

Option 2

Set the "Queue Callers to Group When No Local Agents are Logged In" option to "Yes" in the ACD Agent Skill Groups form of the resilient agent skill groups. Then, if the gateway redirects calls to the resilient agent skill group on the secondary before the first agent has registered, the calls are queued against the resilient agent skill group.

Be aware that if you enable this option, calls will always queue to the agent skill groups on either the primary or secondary controller when no agents are logged into the group. After all agents log out of an agent skill group at the end of day, the calls will not be routed to the Path Unavailable Answer Point. However, you can still achieve Path Unavailable Answer Point routing by manually putting the Path DN in Do Not Disturb at the end of the day.

You can also use Option 2 if you do not require the Path Unavailable Answer Point to automatically route calls to a "Business is closed for the day" greeting. Instead, you could use the Mitel Intelligent Queue solution to act as a front end with Time of Day routing to the greeting.

3.1.6 Licensing Requirements

The following licensing rules and system maximums apply:

- In order to be resilient, a resilient hot desk ACD agent must log into a resilient device. However, the resilient hot desk ACD agent and the resilient device do not have to be configured with the same secondary controller.
- A resilient hot desk ACD agent requires
 - an ACD license on the primary controller
 - an IP User license on the primary controller
 - a PLID on both the primary and secondary controller
- Hot desk ACD agents can only log into IP phones, so the maximum number of active hot desk ACD agents is limited to the maximum number of IP phones that can be configured as ACD sets on the controller. Refer to the System Administration Tool Help to determine the maximum number of IP phones that can be configured as ACD agents on each type of controller.

[Resilient Hot Desk ACD Agent License Requirements](#) summarizes the license requirements and PLID requirements for each resilient ACD hot desk agent:

Table 5: Resilient Hot Desk ACD Agent License Requirements

License	Primary Controller	Secondary Controller
ACD license	1	0
Device license	0	0
User license	1	0
PLID	1	1

[Resilient Device Licensing Requirements](#) summarizes the license requirements and PLID requirements for each resilient device (either resilient hot desking enabled set or resilient ACD set):

Table 6: Resilient Device Licensing Requirements

License	Primary Controller	Secondary Controller
Device license	1	0
User license	1	0
PLID	1	1

3.1.7 Maximum Number of ACD Agents Agent IDs and Paths

Refer to the ACD documentation in the System Administration Tool Help for the maximum number of agents, agent IDs, and paths supported on standalone and clustered systems.

3.1.8 Maximum Number of Hot Desk User Profiles

The maximum number of hot desk agents and hot desk ACD agents that you can configure on the controllers with 3300 ICP Release 7.0 or later software are listed below:

Table 7: Maximum Number of Hot Desk User Profiles

Platform	Maximum user profiles (includes hot desk users and hot desk ACD agents)	Maximum Active Users (includes hot desk users and hot desk ACD agents)
AX	700	100
CX	700	100
MX	700	200


Platform	Maximum user profiles (includes hot desk users and hot desk ACD agents)	Maximum Active Users (includes hot desk users and hot desk ACD agents)
LX (700 user)	700	700
LX (1400 user)	5600	1400
MXe (base)	700	300
MXe (expanded)	5600	1400
MXe Server/MiVoice Business Virtual/MiVoice Business Multi-instance/ MiVoice Business for Industry Standard Servers	5600	2500

3.1.9 Configuring ACD Resiliency

Overview

The following flowchart presents an overview of the steps required to configure ACD resiliency:

Figure 16: ACD Resiliency Configuration Flowchart



Review Licensing Requirements

Ensure that you have a sufficient number of licenses to support the resilient hot desk ACD agents. Refer to [Licensing Requirements](#) for details.

Plan your ACD Call Center

For basic path and agent skill group planning information, refer to the System Administration Tool Help. For information on planning ACD resilient call center configurations refer to the detailed version of the Resiliency Guide.

- [Resilient ACD Configurations](#)

- [Resilient Virtual Contact Center Configurations](#).

Select Agent Failover Call Redirection Option

Refer to [Agent Failover and Call Redirection](#).

Upgrade to a Cluster with IP Trunking

Upgrade the network to a cluster with IP trunking. Refer to the *Voice Networking -> Configure Network* book in the *System Administration Tool Help* for instructions.

To support ACD resiliency, the gateway controllers, primary controllers, and secondary controllers (all participating controllers) must have 3300 Release 7.0 software or later. To support Embedded Device Resiliency, all participating controllers must have 3300 ICP Release 4.0 software or later and must be migrated to support RDN Synchronization. Embedded Device Resiliency allows you to manage device resiliency from the System Administration Tool without the requirement for OPS Manager in the cluster.

Upgrade the gateway controllers first, followed by the primary and secondary controllers. Refer to the *3300 ICP Technician's Handbook* for software upgrade instructions.

Configure System Data Synchronization SDS

Enable the System Data Synchronization (SDS) feature to share data in the cluster. For ACD resiliency, configure SDS to share the following form data from the master element:

- Class of Service
- Class of Restriction
- Feature Access Codes
- Independent Account Codes
- ACD Agent Skill Group Assignment form (shared by default when SDS is enabled)
- ACD Agent Skill Group Membership (shared by default when SDS is enabled)
- Device and User Data
- User Call Handling Data.

A "master element" is the element that you are logged into when you initiate an synchronization or update a form record (add, edit, or delete). All other elements that belong to the same SDS community (network or cluster) as the master element are referred to as "member elements". After enabling SDS and specifying the shared data, perform a synchronization operation to the member elements. Refer to System Data Synchronization in the *Voice Networking -> Manage Network -> Configure Network* book of the System Administration Tool Help for instructions

Configure User and Device Resiliency

As a prerequisite to ACD Resiliency, you must configure the network to support user and device resiliency. Refer to [“Implementing Resiliency”](#) for instructions.

Configure the Hot Desk Sets and or ACD Sets

Configure the hot desk sets, ACD sets, and Network ACD on the primary controller. Refer to the System Administration Tool Help for instructions.

Program ACD Feature Access Codes

Program the ACD feature access codes on the master element. SDS automatically distributes the ACD feature access codes to the other members elements.

Configure Hot Desk ACD Agents From System Administration Tool

If you have an existing ACD site, you should convert the existing traditional ACD agents to hot desk ACD agents. If you are installing a new site, add the new agent IDs as hot desk ACD agents. This section provides instructions for the following procedures from the System Administration Tool:

- converting a single existing traditional ACD agent into a ACD hot desk agent
- converting multiple existing traditional ACD agents into hot desk ACD agents
- adding a new agent ID as a ACD hot desk agent from the System Administration Tool
- adding multiple agent IDs as hot desk ACD agents from the System Administration Tool.

You manage hot desk ACD agents through the following System Administration Tool forms:

- User Configuration
- Station Service Assignment
- Change Attribute Assignment

Note:

1. The User and Device Attributes form displays a User Type of "Hot Desk" for hot desk users and "ACD Agent" for hot desk ACD agents.
2. You cannot configure hot desk ACD agents from the Multiline IP Sets form. You should configure hot desk ACD agents from the User and Services Configuration form. In the Multiline IP Sets form, the "Hot Desk Type" field and "ACD Set" field are set to "Yes" for hot desk ACD agents.

Converting a Single Existing Traditional ACD Agent into a Hot Desk ACD agent

When you convert an agent, you must enter the ACD agent ID of an existing agent in the Number field in the User and Services Configuration form.

1. Ensure that the agent is logged out. You cannot convert an agent that is logged into the system.
2. Log into the System Administration Tool of the primary controller.
3. Access the User and Services Configuration form. Click **Add**.

Note:

You cannot convert a traditional agent to a hot desk ACD agent by using the **Change** operation. You must perform an **Add** operation.

4. Either leave the Name fields blank or complete the Name fields for the agent:
 - If you leave the Name fields blank, the system does NOT create a new entry in Telephone Directory Assignment, a new user profile, or a new voice mailbox for the agent.
 - If you complete the Name fields, the system creates a new entry in the Telephone Directory form, a new user profile, and a new voice mailbox for the agent.
5. Enter the Agent ID of the existing traditional ACD agent in the Number field.

Note:

An Agent ID could be assigned to more than one name in the Telephone Directory form. The entries in the Telephone Directory form are not changed by the conversion.

6. Leave the Device Type set to **No Device** .
7. Set the User Type field to **ACD Agent**.
8. Configure the hot desk ACD agent's Embedded Voice Mail, Desktop Tool Access, and User profile as required.

Note:

If you have upgraded from a pre- 3300 Release 7.0 software load, the existing traditional ACD agents will not have user profiles and voice mail properties assigned to them. You can copy the user profile and voice mail properties from the ACD enabled set to the new hot desk ACD agent. To copy the properties select the ACD enabled set in the User and Services Configuration form, click **Copy**, and then from the Copy window, enter the agent ID of the traditional agent into the Number field. Leave the device type set to **No Device** , set the User Type field to **ACD Agent** , and then click **Save** .

9. Click **Save**. The system prompts you to confirm the conversion of the ACD agent.
10. Click **Yes** to proceed with the conversion. The system generates a maintenance log for the agent conversion.

Converting Multiple Existing Traditional ACD Agents into Hot Desk ACD Agents

1. Ensure that all agents are logged out. You cannot convert agents that are logged into the system.
2. Log into the System Administration Tool of the primary controller.
3. From the ACD Agent IDs form, export the information for the existing ACD agents to a .csv file.

Note:

For detailed instructions on how to export data, click **Export** and then click the **Help** button located in the bottom left corner of the Export Dialog box.

- Click **Export**.
 - Under Export Range, click **All**.
 - Under File Type, click **Comma Delimited (Spreadsheet)**.
 - Click **Export**. The File Download window opens.
 - Click **Save** and save the .csv file to a folder on your client station.
 - When the download is complete, click **Close**.
4. From the User and Services Configuration form, import the exported agent ID information, such as Agent IDs, COS, and COR.
- Click **Import**.
 - Click **Download a copy of the Import Spreadsheet.xls**. You must use a 3300 Release 7.0 version of the spreadsheet.
 - Click **Open**. Click **Yes**.
 - Extract the file to a folder on your client station.
 - Open the spreadsheet and enable the macros when prompted.
 - Click the **User Configuration** tab at the bottom of the spreadsheet.
5. Next, open the .csv file that you exported from the ACD Agent IDs form. Copy the Agent IDs, COS, and COR settings into the columns in the Import Spreadsheet.xls file. If desired, configure the voice mail, user profile, and Desktop Tool access for the users.
- Select the rows in the column of data (for example, Agent IDs) from the .csv file.
 - Right-click and select **Copy**.
 - Select the first row of the column in the import spreadsheet.
 - Right-click and select **Paste**.
6. In the spreadsheet, set the User Type field for each agent to **ACD Agent**.
7. Click the **Save for Import** button.

8. From the User and Services Configuration form, import the data from the spreadsheet.

- Click **Import**.
- Click **Browse** and navigate to the worksheet (.csv) file in the spreadsheets folder on your client station.
- Select the UserConfiguration_XXXXX.csv file and then click **Open**.
- Click **Next**. The data is uploaded.

 Note:

When you use the import spreadsheet to convert the agents, the system does not present a confirmation prompt. However, if conversions fail, the system displays error messages after the import operation completes.

9. After the import operation is complete, check the ACD Agent IDs form, Telephone Directory form, Multiline IP Sets form, and User and Services Configuration form to ensure that you have imported the data correctly.

In the Multiline IP Sets form, all hot desk ACD agents should have the Hot Desk User field and the ACD Enabled field set to "Yes". In the User and Services Configuration form, the User Type field should be set to **ACD Agent** for each agent.

10. Check the system maintenance logs. The system will have generated a maintenance log for each agent that was successfully converted.

Adding a New Hot Desk ACD Agent

1. Log into the System Administration Tool of the primary controller.

2. In the User and Services Configuration form

- Click **Add**.
- Enter the name of the ACD agent in the Last Name and First Name fields.
- Enter a directory number in the Number field. The number cannot be used anywhere else in the system.
- Set the User Type field to **ACD Agent**.
- Leave the Device Type set to **No Device**.
- Configure the ACD hot desk agent's Embedded Voice Mail, Desktop Tool Access, and User profile as required.
- Click **Save**. The system adds a record to the Multiline IP Set Configuration for the hot desk DN with the ACD enabled device field set to Yes; adds an Agent ID

into the ACD Agent IDs form (the Agent ID will be the directory number that you entered in the Number field); and adds an entry for the ACD hot desk agent to the Telephone Directory form.

3. In the Multiline Set Keys form

- Program the key assignment. ACD agents must have prime lines of the "single line" type.
- Program Specific Group Threshold Alert keys for Supervisors' telephones. Program the directory number for the Specific Group Threshold Alert key at the set, and assign it in the ACD Agent Skill Groups form as the Alert Device.
- Program Generic Group Threshold Alert keys for Agents' telephones. Do not specify directory numbers for Generic Group Threshold Alert keys.

4. In the ACD Agent Skill Groups form

- Add the agent ID as a member of an ACD agent skill group if they are to take path calls
- If required, set the "Group uses Skill Level" field to "Yes".
- Set the "Skill Level" field for the agent to a value between 1 (highest skill level) and 255 (lowest skill level). An agent can appear only once in a single agent skill group. To disable Skill-Based Routing for an agent skill group, clear the "Skill Level" field of all agents, save the form, set the "Group uses Skill Level" field to "No", and save the form again.

Adding Multiple New Hot Desk ACD Agents

1. Log into the System Administration Tool of the primary controller and download the import spreadsheet to your client station.

Note:

For instructions on how to obtain the import spreadsheet, access the User and Services Configuration form, click **Import** and then click the **Help** button located in the bottom left corner of the Import Dialog box.

- Click **Import**.
 - Click download a copy of the Import Spreadsheet.xls. You use a 3300 Release 7.0 version of the spreadsheet.
 - Click **Open**. Click **Yes**.
 - Extract the Spreadsheet.xls file to a folder on your client station.
 - Navigate to the spreadsheets folder on your client station and open the spreadsheet.
 - Enable the macros when prompted.
 - Click the **User Configuration** tab at the bottom of the spreadsheet.
2. Add records for each ACD hot desk agent.
- Enter the Last Name and First Name of the agent.
 - Enter a Department and Location (optional).
 - Enter a Directory Number for the agent.
 - Set the User Type field to **ACD Agent**.
 - Leave the Cab, Shelf, Slot, Circ blank.
 - Complete the COS and COR fields as required.
 - Specify the Personal Speed Calls, Embedded VM, and Desktop Tool Access as required.
 - Click the **Save for Import** button.
3. From the User and Services Configuration form, hot desk ACD agents import the data from spreadsheet.
- Click **Import**.
 - Click **Browse** and navigate to the worksheet (.csv) file in the spreadsheet folder on your client station.
 - Select the UserConfiguration_xxxxxx.csv file and click **Open**.
 - Click **Next**. The data is uploaded.
4. After the import operation is complete, check the ACD Agent ID Assignment form, Telephone Directory form, Multiline IP Sets, and User and Services Configuration to

ensure that you imported the data correctly. In the User and Services Configuration form, all the hot desk ACD agents should have the User Type field set to **ACD Agent**.

Configure the Hot Desk ACD Agents with Resiliency

After you have programmed ACD resiliency, the System Data Synchronization feature keeps the user and device data, such as Make Busy status, and the Do Not Disturb settings, synchronized with the data on the secondary.

Note:

If you have upgraded from a prior software release to 3300 Release 7.0 or later software and have local ACD agents that are programmed with the same agent ID on both the primary and secondary controllers, you must perform the following steps before you can program the hot desk ACD agents with resiliency:

- From the agents' primary controller, convert the ACD agents to hot desk ACD agents (see ["Configure Hot Desk ACD Agents From System Administration Tool"](#)).
- On the secondary controller, either delete the duplicate agents from the groups that they belong to, or delete the agent skill group(s) that the duplicate agents belong to.
- On the secondary controller, delete the corresponding local ACD agents that share the same agent IDs.

To program a Hot Desk ACD agent with resiliency from the System Administration Tool (RDN Synchronization must be enabled for the cluster):

1. Log into the home (primary) element of the ACD hot desk agent.
2. Access the User and Services Configuration form.
3. Select the Directory Number of the ACD hot desk agent.
4. Under "Phone Service Settings", select the name of the secondary element from the drop-down list of cluster members.
5. Click **Save**. RDN Synchronization automatically updates the other cluster elements with the remote directory number of the ACD hot desk agent; System Data Synchronization automatically synchronizes the agent's user and device data and begins sharing it between the primary and secondary controllers.

Program ACD Agent Skill Group Resiliency

After you have programmed the agents with resiliency, program the agent skill groups as resilient. Resilient agent skill groups should only contain hot desk ACD agents. The hot desk ACD agents can be resilient or non-resilient. You assign a secondary controller to a resilient agent skill group through the ACD Agent Skill Groups form in the System Administration form of the primary controller.

For an agent skill group to be fully resilient, all its member agents must be configured as ACD resilient hot desk agents and they must have the same primary and secondary controller as the resilient agent skill group.



CAUTION:

Do not attempt to make ACD agent skill groups resilient by programming the groups directly in the Remote Directory Numbers (RDN) form. SDS will not share updates to the group with the resilient peer if you program an agent skill group directly in the RDN form



Note:

You can use range programming to modify multiple ACD agent skill groups in a single operation. Refer to the System Administration Tool Help for instructions.

1. Log into the System Administration Tool of the primary element.
2. In the Network Elements form, ensure that the primary (local) and secondary elements are sharing data

Note:

After you enable the System Data Synchronization (SDS) feature and start sharing data between the primary and secondary controllers, the "Agent Skill Group Assignment" info and "Agent Skill Group Members" data are shared by default.

3. In the ACD Agent Skill Groups form
 - select an existing agent skill group and click **Change**, or
 - click **Add** to create a new agent skill group.
4. If you are creating a new group, enter the Skill Group ID Number.
5. For new groups, add the Group ID of the group to the Telephone Directory form.
6. Select the host name of secondary resilient peer from the Secondary Element drop-down menu.
7. If you are creating a new group, complete the First Status Threshold, Second Status Threshold, and Alert Device fields.
8. Set the "Queue calls when no local agents are logged in" to Yes or No as required. See ["Agent Failover and Call Redirection"](#) for a description of the two options.
9. Click **Save**.
10. Repeat the above procedure for each ACD agent skill group that you want to configure as resilient.
11. SDS distributes the agent skill group and its remote directory number information to the secondary element. SDS automatically begins sharing any subsequent changes that are made to the agent skill group data between the primary and secondary elements.

Add Resilient Agents to the Resilient Agent Skill Groups

Before you add resilient agents to a resilient agent skill group, ensure that the resilient agents are assigned to fail over to the same secondary controller as the agent skill group.

1. Log into System Administration Tool of the primary controller.
2. In the ACD Agent Skill Groups form, select the resilient agent skill group.
3. Click **Add Member**.
4. Enter the Agent ID of the resilient agent.
5. Set the skill level of the resilient agent.

6. Click **Save**. SDS automatically replicates the resilient agent entry on the secondary controller. Any subsequent changes that you make to the resilient agent membership information are shared between the primary and secondary elements.
7. Repeat the above steps for each agent and each agent skill group.

Make ACD Express Groups and ACD Express Agents Resilient

1. Make the ACD Express Agents that are to form the ACD Express Group resilient. They must have the same primary and secondary system.
2. To form the first instance of the group, program the ACD Express Group on the primary MiVoice Business system of the ACD Express Agents.
3. To form the second instance of the group, program an identical ACD Express Group on the secondary system of the ACD Express Agents but assign a different group pilot DN.
4. Do one of the following:
 - If a trunking gateway hosts the ACD Express Group, use T1/E1 Trunking Resiliency. With this type of resiliency, when the primary MiVoice Business system of the ACD Express Agents is down, the Resilient T1/E1 trunks direct calls for the ACD Express Group to the secondary system of the ACD Express Agents. Digit Modification programming ensures that calls for the ACD Express Group reach the group's appropriate instance, either on the primary or on the secondary system.
 - If a trunking gateway is not the ACD Express Group host, and IP trunks link the gateway to the ACD Express Group host, use ARS Alternate Routing from the trunking gateway to the two ACD Express Group host MiVoice Business systems. If the primary MiVoice Business system of the ACD Express Agents is down, the Alternate Route programming directs calls for the ACD Express Group to the secondary system of the ACD Express Agents. Digit Modification programming ensures that calls for the ACD Express Group reach the group's appropriate instance, either on the primary or on the secondary system.

Complete Path Programming

1. Ensure that any remote cluster agent skill groups are displayed correctly in the Remote Directory Numbers form.
2. On the Answering Point system:
 - Create the ACD Agent Skill Group on the Answering Point system.

3. On the Distributor system:

- Verify that the Agent Skill Group ID that was created on the Answering Point system is present in the Remote Directory Numbers form on the Distributor system.
- Create the ACD Path and assign the Remote Agent Group ID as the Primary Agent Skill Group in the ACD Path.
- Ensure that the ACD Path DN is also a Portable Directory Number (that is, the ACD Path DN is in the Remote Directory Numbers form on the Answering Point system).

3.1.10 Maintaining ACD Resiliency Data

Provisioning a New or Existing ACD Agent as Resilient

1. Add a new hot desk ACD agent to the primary controller. See [Adding a New Hot Desk ACD Agent](#).

or

Convert the existing traditional ACD agent on the primary controller to a Hot Desk ACD agent. See [Converting a Single Existing Traditional ACD Agent into a Hot Desk ACD agent](#).

2. Configure the ACD hot desk agent with resiliency. See [Configure the Hot Desk ACD Agents with Resiliency](#).
3. Add the resilient agent to a resilient agent skill group. See [Add Resilient Agents to the Resilient Agent Skill Groups](#).

Changing the Secondary Controller of a Resilient Agent Skill Group

1. Before you can change the secondary controller of a resilient agent skill group, you must change the secondary controller of the resilient agent members. See [Configure the Hot Desk ACD Agents with Resiliency](#) for instructions.
2. Next, log into System Administration Tool of the primary controller.
3. In the ACD Agent Skill Groups form, select the resilient agent skill group and then click **Change**.
4. Select the name of the new secondary controller from the drop-down menu. This name must correspond to the name of the secondary controller of the resilient agent members.
5. Click **Save**. SDS then deletes the group resilient data from the old secondary and distributes the data to the new secondary controller.

Adding and Deleting Resilient Agent Skill Group Members

You can add and delete resilient agent skill group members through the ACD Agent Skill Groups form of either the primary or secondary controller. SDS distributes the add or delete operation to the peer controller.

Before you can add or delete a resilient agent skill group member,

- the agent device must have been configured as resilient
- the agent device data must have been distributed to the secondary controller.
- the agent must be a member of one or more resilient groups on the agent primary controller

If these conditions are not met, the add or delete update will fail and the SDS feature will generate an alarm. You can correct the problem and retry the data distribution update from the SDS Distribution Errors form.

Note:

To move an agent member, you must delete the agent from the current group and then add the agent to the desired group.

Deleting a Resilient Agent

To delete a resilient agent

1. Delete the resilient agent ID from the agent skill group on the primary controller
2. SDS then updates the group membership on the secondary controller.
3. Delete the resilient agent from the User and Services Configuration form in the System Administration Tool of the primary controller (RDN Synchronization must be enabled for the cluster).

After you delete a resilient ACD agent, the delete operation is propagated to the other elements.

Disabling Agent Skill Group Resiliency

You can only disable agent skill group resiliency from the primary controller. On the secondary controller, the Home Element and Secondary Element fields are read only. To

disable resiliency for an agent skill group, log into the System Administration Tool on the primary controller and set the Secondary Element field to "Not Assigned".

Deleting a Resilient Agent Skill Group

You can delete a resilient agent skill group from either the primary or secondary controller. To delete a resilient agent skill group, access the ACD Agent Skill Groups form, select the group, and click **Delete**. Note that the delete will fail if the group is in used by the local path.

Changing Hot Desk ACD Agents Back to Traditional ACD Agents

To change a Hot Desk ACD agent back to a traditional ACD agent

1. If the ACD hot desk agent is resilient, disable resiliency from the agent, from the User and Services Configuration form of the System Administration Tool on the primary element (RDN Synchronization must be enabled for the cluster).
2. Log into the System Administration Tool of the primary controller.
3. In the User and Services Configuration form, select the ACD hot desk agent, and click **Delete**.

Note:

An agent skill group must have at least one agent. Therefore, you cannot delete the last agent in an agent skill group. You must delete the group.

4. In the ACD Agent ID Assignment form, add the Agent ID for the traditional ACD agent.
5. In the Telephone Directory form, enter the Agent ID as a directory number and assign a name to the directory number. The system writes the name in the ACD Agent IDs form against the Agent ID.

Removing ACD Resiliency

To remove ACD Resiliency complete the following steps.

1. Delete the resilient agent members from the resilient agent skill groups. See [Adding and Deleting Resilient Agent Skill Group Members](#) on page 133
2. Disable resiliency from the agent skill groups. See [Disabling Agent Skill Group Resiliency](#) on page 133.

3. Disable resiliency from the ACD agent from the User and Services Configuration form of the System Administration Tool at the primary element (RDN Synchronization must be enabled for the cluster).

3.2 Resilient ACD Configurations

Overview

The resilient ACD configurations described in this chapter are created by combining the Networked ACD (NETACD), ACD Agent Hot Desking, and the MiVoice Business ACD Resiliency feature. NETACD provides call distribution between the path controllers and the primary and secondary controllers, while ACD Agent Hot Desking and ACD Resiliency supports agent failover and failback between the primary and secondary controllers.

In this document, the configurations have been classed based on the level of resiliency achieved:

- Basic ACD Resiliency
- Advanced ACD Resiliency
- Full ACD Resiliency

The advantages and disadvantages of each configuration are identified.

Note:

Refer to the ACD documentation in the MiVoice Business System Administration Tool Help and to [Configuring ACD Resiliency](#) for programming instructions.

Resiliency --Why It's Important

Call Centers generate profits from call traffic so any downtime in the ACD system will cost the owner. Lost calls results in lost sales. Many call centers operate 24 hours a day, seven days a week and cannot afford to be down. Even call centers that operate as cost centers have service levels that they must meet (for example, answering 80 percent of their calls within 40 seconds). Failure to meet service levels affects customer satisfaction so any system outages have a negative impact on the business.

Also, call center owners do not want to purchase legacy Time Division Multiplexing (TDM) telephone systems to obtain ACD resiliency. They want resilient Voice over IP (VoIP) ACD systems. Prior to this application note, customers may not have considered Mitel's MiVoice Business ACD to be as reliable as its TDM solution.

What is the Solution?

The solution is to configure the MiVoice Business ACD system with NETACD and MiVoice Business ACD Resiliency so that agents can continue answering calls in the unlikely event of a controller outage, ensuring that queued calls are not lost and active calls remain in progress.

Furthermore, even scheduled outages, such as software upgrades, will not take the call center out of service because technicians will be able to upgrade controllers one at a time.

What are the Advantages?

The Mitel MiVoice Business ACD Resiliency solution is highly distributed, spreading trunk density and agents across several nodes. It provides a high level of agent resiliency and ensures that there is no single point of failure. It is also scalable and provides seamless integration with Mitel's existing suite of IP applications such as the Mitel Teleworker Solution.

3.2.1 Basic ACD Resiliency

As shown in [Basic Resiliency - Configuration 1](#) and [Figure 18: Basic Resiliency - Configuration 2](#) on page 137 the ACD paths are programmed on both MiVoice Business ACD agent controllers. In these configurations, the paths are set up with some of the agent skill groups and some of the agents on both controllers and the paths take calls for the same 800 number. In Configuration 1 ([Basic Resiliency - Configuration 1](#)), the Central Office (CO) / Public Exchange calls are split evenly between the ACD primary controller and the ACD secondary controller. In Configuration 2 ([Figure 18: Basic Resiliency - Configuration 2](#) on page 137) exchange calls are split evenly between the ACD controller 1 and ACD controller 2. IP Networking is used between MiVoice Business systems.

Figure 17: Basic Resiliency - Configuration 1





Figure 18: Basic Resiliency - Configuration 2



In Configuration 1:

- Agents on the ACD agent primary controller are programmed as resilient hot desk ACD agents and will fail over to ACD agent secondary controller in the event that primary controller fails.
- Resiliency is enabled for the agent skill groups. The agent skill groups are "primary" on the ACD primary controller and "secondary" on the ACD secondary controller. The agent skill group members are resilient hot desk agents configured with the same primary and secondary as the groups.

- 50% of the calls arrive through the ACD agent primary controller queue to a local path to group A.
- 50% of the calls arriving through the ACD agent secondary controller (backup) queue locally to group A and then remotely to group A on the primary controller via NETACD.
- This configuration provides redundant queuing by allowing each call to queue at both controllers.

In Configuration 2:

- 50 hot desk ACD agents are programmed on ACD controller 1 as resilient devices and will fail over to ACD controller 2 in the event that ACD controller 1 fails. 50 hot desk ACD agents are programmed on ACD controller 2 as resilient devices and will fail over to ACD controller 1 in the event that ACD controller 2 fails.
- Resiliency is enabled for the agent skill groups:
 - On ACD controller 1, resilient agent skill groups are configured with their primary controller as ACD controller 1 and their secondary controller configured as ACD controller 2. The group members are resilient hot desk agents. Their primary controller is ACD controller 1 and their secondary controller is ACD controller 2.
 - On ACD controller 2, resilient agent skill groups are configured with their primary controller as ACD controller 2 and their secondary controller configured as ACD controller 1. The group members are resilient hot desk agents. Their primary controller is ACD controller 2 and their secondary controller is ACD controller 1.
- Calls arriving through ACD controller 1 queue locally on group A and then remotely on ACD controller 2 via NETACD on group B.
- Calls arriving through ACD controller 2 queue locally on group B and then remotely on ACD controller 1 via NETACD on group A.

For both configurations ([Basic Resiliency - Configuration 1](#) and [Figure 18: Basic Resiliency - Configuration 2](#) on page 137), the following conditions apply:

- 3300 ICP Release 7.0 software or later is required on the controllers.
- System Data Synchronization must be enabled on the primary and secondary controllers. By default, SDS shares user and device data at the "Resilient Pair" scope.
- If the elements in the cluster have pre-MCD Release 4.0 software, Mitel OPS Manager is required to program resiliency for the group members. If all the elements in the cluster have MCD Release 4.0 or later software, device resiliency can be programmed through the System Administration Tool (RDN Synchronization must be enabled for the cluster).
- Agent Skill Group resiliency is programmed through the System Administration Tool.
- For call center reporting, Mitel Contact Center Solutions collection points are required for both nodes.
- RADs must be configured on the digital trunking gateway by using the "embedded RAD" functionality that is available through the MiVoice Business embedded voice mail or by using a third-party RAD unit (for example, Mitel Contact Center Intelligent

Queue, Interlalia, or similar unit). Note that the embedded RADs do not use E2T channels; however, the Intelligent Queue is an IP integration that uses E2T channels when RAD recordings are played.

- Up to 256 paths (999 on an MxServer with MCD Release 5.0 and on MiVoice Business ISS, MiVoice Business Virtual , and MiVoice Business Multi-instance) are supported on each controller.
- Up to 1,181 agent IDs (2100 on an MxServer with MCD Release 5.0 and on MiVoice Business ISS, MiVoice Business Virtual, and MiVoice Business Multi-instance) are supported on each controller.
- Up to 100 IP ACD agents are supported on each controller.

Basic Resiliency Configuration 1

Failover Conditions

The following events occur if the primary controller in [Basic Resiliency - Configuration 1](#) fails:

- Calls queued at the primary controller are lost.
- CO / Public Exchange detects that the primary controller is out of service and routes all new calls through the secondary controller (note this is a subscribed service that is provided through the carrier).
- Eventually, 100 hot desk ACD agents fail over from the primary controller to the secondary controller and are available to answer calls that arrived on the secondary and are queued on the secondary. Once they are on the secondary controller, agents have access to the same ACD features that they would have on their primary controller (for example: Make Busy, Do Not Disturb, and so forth). However, an agent's Work Timer is cancelled after failover or fallback.

Return-to-Service Conditions

When the primary controller in [Basic Resiliency - Configuration 1](#) returns to service:

- The CO / Public Exchange detects node is back up and once again splits calls evenly between the primary and secondary controllers. Note that this is a subscribed service that is provided through the carrier.
- 100 hot desk ACD agents home back to their primary controller when idle and health check is complete. After agent devices fail back from the secondary controller to primary controller, the agents are automatically logged into the system and can begin accepting calls.
- Any calls that were queued on the secondary controller while the primary was down will be queued on the primary controller.
- 100 agents now reside on the primary controller (MxServer controller with Ethernet-to-TDM (E2T) card that provides 128 channels) taking calls.

During a secondary controller failure:

- No agents are taken out of service; however, calls queued on the secondary controller are lost.

When the secondary controller returns to service:

- Calls are once again split evenly between the controllers.

Advantages and Disadvantages

The advantages:

- No single point of failure.
- Most economical – two controller solution.
- 100 percent of agents remain in service.
- 100 percent of trunks remain in service if each controller is provisioned with the required hardware to support 100% of the trunk traffic.

The disadvantages:

- Less resilient than Advanced or Full Resiliency Configurations.
- Limited scalability. Maximum of 100 IP agents per controller.
- A controller outage can reduce the number of trunks by 50% if each controller is not provisioned to support 100% of the call traffic.

Basic Resiliency Configuration 2

Failover Conditions

The following events occur if the ACD controller 1 in [Figure 18: Basic Resiliency - Configuration 2](#) on page 137 fails:

- Calls queued at the primary controller are lost.
- CO / Public Exchange detects that the primary controller is out of service and routes all new calls through the secondary controller. (Note this is a subscribed service that is provided through the carrier).
- Eventually, 50 hot desk ACD agents fail over from ACD controller 1 to ACD controller 2. and are available to answer calls. While on the secondary controller, agents have access to the same ACD features that they would have on their primary controller (for example: Make Busy, Do Not Disturb, and so forth). However, an agent's Work Timer is cancelled after failover or failback.
- 100 agents now reside on ACD controller 2 (MXeIII controller with E2T card and 128 channels) taking calls.

Return-to-Service Conditions

After the ACD controller 1 in [Figure 18: Basic Resiliency - Configuration 2](#) on page 137 returns to service:

- CO / Public Exchange detects that the controller is back up and once again it splits calls evenly between ACD controllers 1 and 2 (note that this is a subscribed service provided through the carrier).
- 50 hot desk ACD agents home back to their primary controller when idle and health check is complete. After agent devices fail back from ACD controller 2 to ACD controller 1, the agents are automatically logged into the system and can begin accepting calls.
- Any calls that were queued on ACD controller 2 while ACD controller 1 was down will be requeued on the ACD controller 1.

During an ACD controller 2 failure:

- Same scenario as ACD controller 1.

When the ACD controller 2 returns to service:

- Same scenario as ACD controller 1.

Advantages and Disadvantages

The advantages:

- No single point of failure.
- Most economical two controller solution.
- 100 percent of agents remain in service.
- 100 percent of trunks remain in service if each controller is provisioned with the required hardware to support 100% of the trunk traffic.

The disadvantages:

- Less resilient than Advanced or Full Resiliency Configurations.
- Not scalable -- limited to 100 IP agents per controller.

3.2.2 Advanced ACD Resiliency

In the configuration shown in [Advanced ACD Resiliency - Configuration 3](#), the ACD paths are programmed on the trunking gateway.




Figure 19: Advanced ACD Resiliency - Configuration 3

In Configuration 3 ([Advanced ACD Resiliency - Configuration 3](#)):

- Agents on the ACD primary controller are programmed as resilient hot desk ACD agents and will fail over to secondary controller in the event that the primary controller fails.
- Resiliency is enabled for the agent skill groups. The agent skill groups are "primary" on the ACD primary controller and "secondary" on the ACD secondary controller. The agent skill group members are resilient hot desk agents configured with the same primary and secondary as the group.
- All calls arrive through the trunking gateway and queue to a path on the gateway. The gateway then queues the calls to a resilient agent skill group on the primary ACD agent controller.
- Physical channels are not used between the gateway and primary controller until calls are distributed to agents.
- Under normal operation, calls are only answered from the primary controller.
- 3300 ICP Release 7.0 software or later is required on the controllers.
- All paths are programmed on the trunking gateway.
- Queued callers listen to RADs or MOH at the trunking gateway and are distributed to agents via Networked ACD as agents become available.

- Up to 256 paths are supported on the trunking gateway. With 3300 ICP Release 5.0 software, support is increased to 999 paths on the MXe Server and on MiVoice Business ISS, MiVoice Business Virtual, and MiVoice Business Multi-instance.
- Up to 1181 agent IDs are supported on the agent controllers. With 3300 ICP Release 5.0 software, support is increased to 2100 agents on the MXe Server and on MiVoice Business ISS, MiVoice Business Virtual, and MiVoice Business Multi-instance.
- Scalability - Up to 350 active agents (2100 on an MXe Server running 3300 ICP Release 5.0 and on MiVoice Business ISS, MiVoice Business Virtual, and MiVoice Business Multi-instance) can be supported on the primary or secondary controller depending on the number of gateways. In configuration 3
 - CX and CXi can support up to 65 agents (all IP phones) with one gateway
 - MX can support up to 128 agents (all IP phones) with one gateway
 - MXe (with 128 MHz or higher processor) and LX can support up to 128 agents (all IP phones) with one gateway. Up to 350 IP phone agents can be supported with additional gateways.

Note:

Only IP phones support hot desk ACD agents.

- Recorded Announcement Devices (RADs) can be configured using the “embedded RAD” functionality available through the MiVoice Business onboard voice mail or by using a third-party RAD box (for example, Intelligent Queue, Interallia).
- System Data Synchronization (SDS) is sharing data among the network and elements.
- If the elements in the cluster have pre-MCD Release 4.0 software, Mitel OPS Manager is required to program resiliency for the group members. If all the elements in the cluster have MCD Release 4.0 or later software, device resiliency can be programmed through the System Administration Tool (RDN Synchronization must be enabled for the cluster).
- Agent Skill Group resiliency is programmed through the System Administration Tool.
- For Call Center reporting, Contact Center Solutions collection points are required for all nodes.

Advanced Resiliency Configuration 3

Failover Conditions

If the primary controller in [Advanced ACD Resiliency - Configuration 3](#) fails:

- 100 percent of queued calls are maintained.

- Calls in progress at hot desk ACD agents on the primary controller are maintained (call survival). However, the agents are unable to access any features or dialing functions because the primary is out of service. After an agent ends the current call, the agent's ACD set fails over to the secondary controller. The system automatically logs agents into the secondary controller. As agent members fail over to the resilient agent skill group on the secondary, they begin to answer calls.
- Eventually, 120 hot desk ACD agents fail over from the primary controller to the secondary controller and are available to answer calls. Once they are on the secondary controller, agents have access to the same ACD features that they would have on their primary controller (for example: Make Busy, Do Not Disturb, and so forth). However, an agent's Work Timer is cancelled after failover or fallback.
- Trunking gateway now queues calls on the secondary controller.
- If the primary controller fails, calls that are queued to a resilient agent skill group on the primary are rerouted to its resilient agent skill group on the secondary controller. After the primary returns to service, calls are automatically rerouted from resilient agent skill group on the secondary controller to the resilient agent skill group on the primary controller.
- When the primary controller fails, agents may not fail over immediately to the secondary controller. For example, all agents could be engaged in calls that remain up after the primary fails (call survival). If new calls arrive at the secondary controller before any agents have failed over, the calls will follow path unavailable routing on the secondary controller. However, if you set the "Queue Callers to Group When No Local Agents are Logged In" option to "Yes" for the resilient agent skill groups, the incoming calls will queue to the resilient agent skill group and wait for the agents to home to the secondary. Note that if this option is enabled, then the last agent who logs out of the group at the end of the shift or workday must manually set the path to unavailable.

Return-to-Service Conditions

- Trunking gateway detects that the node is in service and once again resumes queuing on the primary controller.
- 120 hot desk ACD agents home back to their primary controller when idle and health check is complete. After agent devices fail back from the secondary controller to primary controller, the agents are automatically logged into the system and can begin accepting calls.
- Any calls that are active on the secondary controller will be allowed to complete. After agents hang up, their sets fail back.
- Any calls that were queued on the secondary controller while the primary was down will be rerouted to the primary controller.
- 120 agents now reside on the primary controller.

Advantages and Disadvantages

The advantages:

- If the primary controller fails
 - 100 percent of agents remain in service
 - 100 percent of trunks remain in service.
- Agents in talk state can complete active calls before failover.
- Scalability - Up to 350 active agents (2100 on an MXe Server with MCD 5.0 software and on MiVoice Business ISS, MiVoice Business Virtual, and MiVoice Business Multi-instance) can be supported on the primary or secondary controller depending on the number of gateways. In configuration 3
 - Cx and CXi can support up to 65 agents (all IP phones) with one gateway
 - MX can support up to 128 agents (all IP phones) with one gateway
 - MXe (with 128 MHz or higher processor) and LX can support up to 128 agents (all IP phones) with one gateway. Up to 350 IP phone agents can be supported with additional gateways.

Note:

Only IP phones support hot desk ACD agents.

The disadvantages:

- Gateway is a single point of failure.

3.2.3 Full ACD Resiliency

In the configuration shown in [Full ACD Resiliency - Configuration 4](#), the ACD paths are programmed on multiple trunking gateways.




Figure 20: Full ACD Resiliency - Configuration 4

In Configuration 4 ([Full ACD Resiliency - Configuration 4](#)):

- Agents on the ACD primary controller are programmed as resilient hot desk ACD agents and fail over to the ACD secondary controller in the event that the primary controller fails.
- Resiliency is enabled for the agent skill groups. The agent skill groups are "primary" on the ACD primary controller and "secondary" on the ACD secondary controller. The agent skill group members are resilient hot desk agents configured with the same primary and secondary as the group.
- Calls that arrive through trunking gateway 1, queue to a path on the gateway. The gateway then queues the calls on the ACD primary controller.
- Calls that arrive through trunking gateway 2, queue to a path on gateway 2. The gateway then queues the calls on the ACD primary controller.
- Physical channels are not used between the gateway and primary controller until calls are distributed to the agents.
- Under normal operation, calls are only answered from the primary controller.
- 3300 ICP Release 7.0 software is required on the controllers.
- Load sharing (queueing) is provided across two trunking gateways.
- ACD paths are programmed with the same agent skill groups and same agents on both trunking gateways and handle calls for the same 800 number.
- Central Office (CO) calls are split evenly between trunking gateways 1 and 2.

- IP Networking is configured between the controllers.
- Longest idle agent call distribution is supported.
- Queued callers listen to RADs and MOH at the trunking gateways and are distributed to agents via Networked ACD as agents become available
- Up to 256 paths are supported on the trunking gateways, 999 if the gateway is an Mx Server with MCD 5.0 software and on MiVoice Business ISS, MiVoice Business Virtual, and MiVoice Business Multi-instance.
- Up to 1181 agent IDs are supported on the controllers. An Mx Server with MCD 5.0 software and MiVoice Business ISS, MiVoice Business Virtual, and MiVoice Business Multi-instance supports 2100 IDs.
- Scalability - Up to 350 active agents (2100 on an Mx Server with MCD 5.0 software and on MiVoice Business ISS, MiVoice Business Virtual, and MiVoice Business Multi-instance) can be supported on the primary or secondary controller depending on the number of gateways. In configuration 4
 - Cx and CXi can support up to 65 agents (all IP phones)
 - MX can support up to 128 agents (all IP phones)
 - Mx (with 128 MHz or higher processor) and LX can support up to 256 agents (all IP phones) with two gateways. Up to 350 IP phone agents can be supported with additional gateways.

Note:

Only IP phones support hot desk ACD agents.

- RADs can be configured on the trunking gateway using the "embedded RAD" functionality available through the 3300 ICP embedded voice mail or by using a third-party RAD box (Intelligent Queue, Interallia, and so forth).
- System Data Synchronization (SDS) is sharing data among the network and elements.
- If the elements in the cluster have pre-MCD Release 4.0 software, Mitel OPS Manager is required to program resiliency for the group members. If all the elements in the cluster have MCD Release 4.0 or later software, device resiliency can be programmed through the System Administration Tool (RDN Synchronization must be enabled for the cluster).
- Agent Skill Group resiliency is programmed through the System Administration Tool.
- For Call Center reporting, Contact Center Solutions collection points are required for all nodes.

Full Resiliency Configuration 4

Primary Controller Failover Conditions

If a the primary controller in [Full ACD Resiliency - Configuration 4](#) fails:

- 100 percent of queued calls are maintained.
- 150 hot desk ACD agents fail over from the primary controller to the secondary controller. Once on the secondary controller, agents have access to the same ACD features that they would have on their primary controller; for example: Make Busy, Do Not Disturb, and so forth. However, an agent's Work Timer is cancelled after failover or fallback. Their agent IDs on both controllers can be the same.
- Trunking gateways queue new calls on the secondary controller and reroute existing calls to the secondary controller.
- When the primary controller fails, agents may not fail over immediately to the secondary controller. For example, all agents could be engaged in calls that remain up after the primary fails (call survival). If new calls arrive at the secondary controller before any agents have failed over, the calls will follow path unavailable routing on the secondary controller. However, if you set the "Queue Callers to Group When No Local Agents are Logged In" option to "Yes" for the resilient agent skill groups, the incoming calls will queue to the resilient agent skill group and wait for the agents to home to the secondary. Note that if this option is enabled, then the last agent who logs out of the group at the end of the shift or workday must manually set the path to unavailable.
- If the primary controller fails, calls that are queued to a resilient agent skill group on the primary are rerouted to its resilient agent skill group on the secondary controller. After the primary returns to service, calls are automatically rerouted from resilient agent skill group on the secondary controller to the resilient agent skill group on the primary controller.
- Any calls that are active prior to the primary controller failure will be allowed to complete (held calls are lost). After the agents hang up, their sets home to their secondary controller.
- 150 hot desk ACD agents are now available to take calls on the secondary controller.

Return-to-Service Conditions

When the primary controller in [Full ACD Resiliency - Configuration 4](#) returns to service:

- Trunking gateways detect that the controller is in service and once again resumes queuing on primary controller.
- Any calls that were queued on the secondary controller while the primary was down will be rerouted to the primary controller.
- 150 hot desk ACD agents fail back to their primary controller when idle and health checks are complete.
- Active calls on the secondary controller are allowed to complete before the agents' devices fail back (held calls are lost). After an agent hangs up, their set fails back to the primary controller.
- Calls that are queued on the secondary controller while the primary controller was down will be redirected to the primary controller.

- 150 hot desk ACD agents now reside on the primary controller (MXeIII controller with E2T card supporting 128 channels).

Advantages and Disadvantages

The advantages are:

- There is no single point of failure.
- If primary fails, all agents remain in service.
- All trunks remain in service if each trunking gateway is provisioned with the required hardware to support 100% of the trunk traffic.
- All queued calls are maintained when an agent controller fails.
- 50% of queued calls are maintained if a trunking gateway fails.
- If agent controller fails, agents in talk state can complete active calls before failover (held calls are lost).
- Scalability - Up to 350 active agents (2100 on an MxServer with MCD 5.0 software and on MiVoice Business ISS, MiVoice Business Virtual, and MiVoice Business Multi-instance) can be supported on the primary or secondary controller depending on the number of gateways. In configuration 4
 - CX and CXi can support up to 65 agents (all IP phones)
 - MX can support up to 128 agents (all IP phones)
 - MXe (with 128 MHz or higher processor) and LX can support up to 256 agents (all IP phones) with two gateways. Up to 350 IP phone agents can be supported with additional gateways.

Note:

Only IP phones support hot desk ACD agents.

3.2.4 ACD Resiliency Programming Examples

[Programming Example for Advanced Resiliency - Configuration 3](#) depicts an example of ACD programming required on the trunking gateway, primary, and secondary controllers for Advanced Resiliency - Configuration 3 (see [Advanced ACD Resiliency - Configuration 3](#)).

Note:

Refer to the ACD documentation in the MiVoice Business System Administration Tool Help for programming and Chapter [Implementing Resiliency](#) on page 227 of this document for programming instructions




Figure 21: Programming Example for Advanced Resiliency - Configuration 3

In [Programming Example for Advanced Resiliency - Configuration 3](#), the following applies:

- Controllers are clustered. RDN Synchronization keeps the remote telephone directories of the elements synchronized.
- During normal operation, agents reside on primary controller and calls are answered by agents on the primary.
- Agents are programmed as resilient hot desk agents and have the same primary and secondary controller as their agent skill group.
- Resilient agent skill groups are programmed on primary controller. SDS is enabled to share data between the primary and secondary so agent skill group assignment info and membership data is automatically distributed to secondary controller.
- If primary controller goes out of service, agent skill groups and agents fail over to secondary controller.
- Agents are automatically logged into secondary controller on failover and back into primary on failback.

[Programming Example for Full Resiliency - Configuration 4](#) depicts an example of ACD programming required on the trunking gateway, primary, and secondary controllers for Full Resiliency - Configuration 4 (see [Full ACD Resiliency - Configuration 4](#)).




Figure 22: Programming Example for Full Resiliency - Configuration 4

In [Programming Example for Full Resiliency - Configuration 4](#) the following conditions apply:

- Controllers are clustered. RDN Synchronization keeps the remote telephone directories of the elements synchronized.
- During normal operation, all agents reside on the primary controller. All calls through trunking gateway 1 and gateway 2 are queued to primary controller.
- Agents are programmed as resilient hot desk agents have the same primary and secondary controller as their agent skill group.
- SDS is enabled to share data between the primary and secondary controllers. Agent Skill Group assignment info and membership data is automatically distributed from the primary to the secondary controllers.
- If Controller MN53 goes out of service, agent skill group 1 fails over to Controller MN54. If Controller MN54 goes out of service, agent skill group 2 fails over to Controller MN53.
- Agents are automatically logged into their secondary controller on failover and into their primary on failback.

3.3 Converting a Traditional ACD Site: An Example

This section provides an example of how to convert a traditional ACD system on a single controller to a resilient ACD configuration that consists of a gateway controller, a primary agent controller, and a secondary agent controller.

3.3.1 Before the Conversion

In this example, before the conversion the site had a

- single controller (MN45) with traditional ACD configuration
- traditional ACD enabled set with directory number (DN) 5185 that is local to MN45
- traditional ACD agent login ID of 75185 that is local to MN45
- local agent skill group 5155 on MN45 of which ACD agent ID 75185 is a member
- local path 5171 on MN45 with local agent skill group 5155 as the primary agent skill group

[ACD Site Configuration Before Conversion](#) illustrates the site configuration before the conversion.




Figure 23: ACD Site Configuration Before Conversion

3.3.2 After the Conversion

After the site is converted to a resilient ACD configuration:

- The original controller (MN45) is the gateway controller
- Two controllers have been added to the configuration. Primary ACD agent controller (MN46) and secondary ACD agent controller (MN47)
- The ACD agent IDs, ACD enabled sets, and ACD agent skill groups have been moved from the original controller to the primary and secondary ACD agent controllers

- The original ACD agent ID (75185) has been replaced with the ACD enabled set DN (5185)
- The original ACD enabled set, DN 5185 has been converted to a resilient ACD hot desk agent (primary on MN46 and secondary on MN47).
- The original ACD agent ID 75185 has been converted to the resilient ACD enabled DN 75185 (primary on MN46 and secondary on MN47).
- Agent Skill Group 5155 is configured as resilient (primary on MN46 and secondary on MN47)
- Path 5171 is programmed as local on MN46 with remote clustered agent skill group 5155.

Resilient ACD Configuration (After Conversion) illustrates the resilient ACD configuration after the conversion:




Figure 24: Resilient ACD Configuration (After Conversion)

3.3.3 Conversion Procedure

The following procedure outlines the steps that are required to convert the traditional ACD configuration shown in [ACD Site Configuration Before Conversion](#) to the resilient configuration shown in [Resilient ACD Configuration \(After Conversion\)](#). In the following procedure, extension 5185 from [ACD Site Configuration Before Conversion](#) and [Resilient ACD Configuration \(After Conversion\)](#) is used as an example.

1. You must move the agent extensions from MN45 to MN46. Before you move an agent extension from MN45 to MN46:
 - Ensure that the extension does not belong to any group (for example, page group or pickup group).
 - Ensure that the extension is not in the Remote Busy Lamps form. If an extension is in the form, you should export the record before you delete it. Then, you can import the record after you configure the new primary controller with the agents.
 - Check the key assignments of the extensions for single-line keys. If single-line keys are programmed, then delete the associated name from the telephone directory in the MiVoice Business System Administration Tool. Then, perform a synchronization operation.
 - Check the key assignments of the extensions. If any key system keys or multicall keys are programmed, delete them.
2. Move extension 5185 to the primary home element MN46. Assign the secondary element as MN47.
3. Delete the agent name for agent ID 75185 (if an agent name has been programmed) from the telephone directory and perform a synchronization.
4. Log into the primary home element, MN46, through the System Administration Tool and remove agent ID 75185 from all agent skill groups and then delete agent ID 75185.
5. Add 75185 as the physical extension to the primary home element MN46 and assign MN47 as its secondary element. During this step, ensure that you set "ACD Enabled" to "Yes" in the Behavior section of the Add work form.
6. Repeat step 1 to step 5 for any additional sets or ACD agents that you need to move from MN45 to MN46.
7. On MN45 create a temporary agent skill group and make it the primary agent skill group for path 5171.
8. On MN45, record the configuration details of the group 5155 and record all its members. Then, delete agent skill group 5155. Perform an SDS synchronization.
9. On MN46, create agent skill group 5155 and set MN47 as its secondary controller. Add the ACD hot desk agent 5185 to resilient agent skill group 5155. SDS automatically distributes the agent skill group membership information to MN47.
10. Add the name of the resilient agent skill group to the Telephone Directory form. Assign MN46 as the primary controller and MN47 as the secondary controller.
11. Ensure that path 5171 is entered in the telephone directory of MN45 and ensure that it appears in the Remote Directory Assignment forms of MN46 and MN47.
12. On MN46 assign the remote clustered agent skill group 5155 as the primary agent skill group for path 5171.
13. On MN46, reconfigure the Remote Busy Lamps form (if required).

3.4 Resilient Virtual Contact Center Configurations

Introduction

This section describes two configurations for resilient, virtual ACD contact centers. In a virtual ACD contact center, the agents work in different geographical locations, but the contact center functions as if the agents are all located locally in the same office. You create a virtual contact center by combining Networked ACD, MiVoice Business device resiliency, and IP networking.

- Configuration 1 ([Figure 25: Resilient Virtual Contact Center - Configuration 1](#) on page 157) uses only MiVoice Business systems (**recommended** configuration).
- Configuration 2 ([Figure 26: Resilient Virtual Contact Center - Configuration 2](#) on page 161) uses MiVoice Business systems and MiVoice Business Digital Trunking Gateways.

Note:

Refer to Networked ACD in the MiVoice Business System Administration Tool Help and Chapter [Implementing Resiliency](#) on page 227 of this programming instructions.

What is the current issue

Although call distribution is a basic feature of Mitel's ACD functionality on a single system, many contact centers require that calls be evenly distributed (equal load sharing) among a group of agents with a specific skill set regardless of where the agents are located geographically.

Why is it a problem

Prior to IP telephony, multi-site routing was achieved by using Networked ACD which is based upon the ACD “overflow” feature. ACD overflow permits calls originating in a particular geographical location to overflow to agents on a different system in another geographic location in the event of high call volume.

Multi-site routing allows agent resources in another part of the country to reduce the work load of a overloaded primary agent skill group. A shortcoming of this configuration is that the overflow groups only take calls when all agents in the primary group members

are busy on calls. If the primary agents are not all busy, they take all the calls. Multi-site routing is effective for managing call loads and for supporting contact centers that cover different time zones, but it cannot evenly distribute calls among a group of geographically dispersed agents in a “virtual” environment.

What is the solution

The solution is to allow agent devices that are geographically dispersed to be registered to the same MiVoice Business system across an IP network. This solution effectively puts agents that are separated geographically in the same agent skill group (or groups) on the same controller. By implementing this solution, you can have calls evenly distributed among the agents based on which agent has been idle the longest, regardless of where the agent is located. In addition, an ACD supervisor can register on the same controller as the agents regardless of the supervisor’s location, allowing the supervisor to invoke Silent Monitoring on any agent in the agent skill group.

Note:

The functionality described in this section is provided entirely by the MiVoice Business system. An offboard server is not required to manage the call routing in a virtual call center environment.

What is the advantage of this solution


This solution

- is highly distributed, spreading trunk density and agents across several nodes, providing a high level of resiliency, and eliminating single point of failure.
- is scalable and provides seamless integration with our suite of IP applications: Contact Center Solutions Management and Reporting, Teleworker Solution, and so forth.
- allows agents, during a primary controller outage, to function in their “local” ACD environment until their primary controller returns to service and re-establishes the “virtual” environment.

3.4.1 Resilient Virtual Contact Center Configuration 1

Figure 25: Resilient Virtual Contact Center - Configuration 1 on page 157 shows a resilient virtual contact center in which all the ACD paths are programmed on the MiVoice Business systems. **Configuration 1 is the recommended configuration** because this configuration protects against a WAN failure or primary agent controller failure.

Figure 25: Resilient Virtual Contact Center - Configuration 1



The following conditions apply to configuration 1

- All controllers in this configuration are 3300 ICPs (that is, none are 3300 gateways).
- Agents on the primary controller are programmed as resilient hot desk ACD agents and will fail over to their respective ACD agent skill group on the secondary controller in the event that the primary controller fails.
- The controllers that are programmed as the agents' secondary controllers also serve as the digital trunking controllers. All ACD paths are programmed on these controllers (that is, all local Los Angeles (LA) calls come through the LA paths, Dallas calls come through the Dallas paths, and so forth).
- All the 3300 ICPs are clustered
- System Data Synchronization is enabled and sharing data between all the 3300 ICPs.
- All 3300 ICPs are connected by IP networking.
- If the elements in the cluster have pre-MCD Release 4.0 software, Mitel OPS Manager is required to program device resiliency for the group members. If all the elements in the cluster have MCD Release 4.0 or later software, device resiliency can be programmed through the System Administration Tool (RDN Synchronization must be enabled for the cluster).
- Agent Skill Group resiliency is programmed through the System Administration Tool.
- New York agents physically reside in New York, Los Angeles agents physically reside in LA, and Dallas agents physically reside in Dallas.
- All agent phones are registered to the primary agent controller in New York. All agents can therefore be in the same agent skill groups in New York, ensuring that calls are distributed to them based on who has been idle the longest -- even though the agents are geographically dispersed.

- In the event that the primary agent controller goes out of service, the New York agents will fail over to their secondary controller in New York; Los Angeles agents will fail over to their secondary controller in LA; and Dallas agents will fail over to their secondary controller in Dallas. In a virtual configuration, the agents are not in resilient agent skill groups. The agents fail over to a local agent skill group on their secondary controller.
- During the outage, agents work in a “local” ACD environment answering calls that originate from their own geographical region, rather than from a virtual environment.
- The paths are programmed on the primary controller and on each secondary controller. A path’s primary ACD group resides on the secondary agent controller (local trunking gateways); whereas, the overflow ACD groups for the path reside on the primary agent controller in New York.
- All calls arrive through the secondary controller queue first on that controller where the primary agent skill group resides. "Queue caller to Group when No Agents Are Logged In" option is set to Yes for the agent skill group in ACD Agent Skill Groups form.
- During normal conditions, the agent skill group on the secondary controller has no agents logged into it. It is merely used to queue calls locally on that controller in the event that the agent’s primary controller fails.
- After queuing the calls to the local agent skill group, each call will then overflow immediately to the overflow groups on the agents primary controller in New York.
- Under normal operation calls are only answered from the primary controller in a virtual environment.
- Queued callers listen to RADs at the secondary 3300 ICPs, so IP bandwidth is not used unnecessarily.
- Calls are distributed to agents registered to the primary controller in New York as agents become available.
- Up to 256 paths are supported on the secondary 3300 ICPs.
- Up to 1181 agent IDs are supported on the primary and/or secondary agent controllers.
- Up to 350 active agents are supported on the primary and/or secondary agent controllers.
- RADs can be configured on the secondary 3300 ICPs using the "embedded RAD" functionality available through the 3300 embedded voice mail system or through a third-party RAD box (Mitel Intelligent Queue, Intermedia, and so forth).
- For call center reporting, Mitel Contact Center Solutions collection points are required for all nodes.

Primary Controller Outage in New York

In configuration 1, if a primary controller outage occurs in New York:

- 100% of queued calls are maintained.
- LA, NY, and Dallas agent devices fail over from the primary controller in New York to their respective local secondary controller: New York agents failover to their secondary

in New York, LA agents failover to their secondary in LA, and Dallas agents failover to their secondary in Dallas.

- Agents are automatically logged into the agent skill groups on the secondary controllers and commence answering calls. Their agent ID is the same on both controllers.
- Agents have full ACD functionality on their secondary controller. The status of most ACD features is maintained after failover. See [Features Available to Resilient Hot Desk ACD Agents](#) for details.
- Calls continue to queue locally in NY, LA, and Dallas and are serviced by agents now registered at their local controller
- Any calls that are active prior to the primary controller outage will be allowed to complete. After the agent hangs up, their set will register to the secondary. This ensures calls in talk state are not dropped. However, calls on hold are lost.

Primary Controller Returns to Service

- The secondary 3300 ICPs detect that the primary controller in New York has returned to service. Calls resume queuing first locally, and then in New York to re-establish the virtual environment.
- NY, LA, and Dallas agents home back to their primary in NY when idle and health check is complete.
- Any calls that are active on the secondary controller will be allowed to complete. After the agents hang up, their sets fail back ensuring that calls in talk state are not dropped. Held calls are lost.
- When agent devices fail back from the secondary to primary, the agents are automatically logged in on the primary controller.
- Calls that arrived when the primary controller was down will only be queued on the secondary controller. These calls must be redirected back into the path so that they can be queued on the primary. INTERFLOW or Visual Voice Queue capability can be used to redirect calls back into the path to ensure they are requeue on the primary controller.

Secondary Controller Outage

- No agents reside on the secondary controller under normal operating conditions. All calls currently queued via the secondary controller are lost.

Secondary Controller Returns to Service

- No agents reside on the secondary controller under normal operating conditions. Incoming calls will resume queuing after secondary returns to service.

3.4.2 Advantages

Configuration 1 has the following advantages:

- Agents and supervisors in different geographical locations can be represented in the same agent skill groups to create a virtual contact center environment.
- Supported by the 3300 ICP (no off board external application required to support longest idle agent routing).

When the primary ACD agent controller fails:

- 100% of agents remain in service.
- 100% of trunks remain in service.
- Agents can complete 100% of active calls before failover (in talk state only, held calls are lost)
- Up to 350 active agents on primary and up to 100 active agents on secondary controller.
- Local dial '9' and '911' service is available to Los Angeles, New York, and Dallas agents through ARS programming.
- No single point of failure.
- Protects against both WAN failure and primary agent controller failure.

3.4.3 Conditions

The following conditions apply to configuration 1:

- Current route list programming permits only five remote sites to be part of a Virtual Contact Center, that is, the sixth site will not have Local dial '9' and '911' service, rather they must route out of New York. (This can be overcome using loopback IP trunks and the Default Account Code feature).
- Calls queued only to the secondary while the primary was down must be redirected back into the path after all agents fail back to the primary controller. You can use Interflow Timer to redirect the calls.
- Currently, the time displayed on the telephone sets will show New York time zone for all agents.

3.4.4 Resilient Virtual Contact Center Configuration 2

[Resilient Virtual Contact Center - Configuration 2](#) shows a resilient virtual contact center in which the ACD paths are programmed on the 3300 Digital Trunking Gateways.




Figure 26: Resilient Virtual Contact Center - Configuration 2

In this configuration:

- ACD paths are programmed on the 3300 Digital Trunking Gateways (in LA, NY, and Dallas)
- IP-to-TDM gateway functionality (3300 E2T resources) is performed by the Digital Trunking Gateways.
- Agents on the primary controller are programmed as resilient hot desk ACD agents and will fail over to their respective secondary controller in the event that the primary controller fails.

- All the 3300 ICPs are clustered
- System Data Synchronization is enabled and sharing data between all the 3300 ICPs.
- All 3300 ICPs are connected by IP networking.
- If the elements in the cluster have pre-MCD Release 4.0 software, Mitel OPS Manager is required to program device resiliency for the group members. If all the elements in the cluster have MCD Release 4.0 or later software, device resiliency can be programmed through the System Administration Tool (RDN Synchronization must be enabled for the cluster).
- New York agents physically reside in New York, Los Angeles agents physically reside in LA, and Dallas agents physically reside in Dallas.
- All agent phones are normally registered to the primary agent controller in New York.
- All agents can therefore be in the same agent skill groups in New York, ensuring calls are distributed to them based on who has been idle the longest -- even though the agents are geographically dispersed.
- On each gateway, a path's primary ACD group and Overflow groups point to agent skill groups that reside on the primary agent controller.
- Agents on the ACD primary controller are programmed as resilient hot desk ACD agents and fail over to the ACD secondary controller in the event that the primary controller fails.
- Resiliency is enabled for the agent skill groups. The agent skill groups are "primary" on the ACD primary controller and "secondary" on the ACD secondary controller. The agent skill group members are resilient hot desk agents configured with the same primary and secondary as the group.
- Under normal operation calls are only answered from the primary controller.
- Queued callers listen to RADs at the Digital Trunking Gateway (so IP bandwidth is not used unnecessarily).
- Calls are distributed to agents on either the primary or secondary 3300 ICPs in New York using Networked ACD as agents become available.
- Up to 256 Paths are supported on the Digital Trunking Gateway.
- Up to 1181 Agent Ids are supported on the agent controllers.
- Up to 300 active agents are supported on the primary and/or secondary agent controllers.
- RADs can be configured on the trunking gateway using the "embedded RAD" functionality available through the 3300 embedded voice mail application or by using a third-party RAD box (Contact Center Intelligent Queue, Interallia, and so forth).
- For Call Center reporting, Contact Center Solutions collection points are required for all nodes.

Primary Controller Outage

In [Figure 26: Resilient Virtual Contact Center - Configuration 2](#) on page 161, the primary controller is located in the New York office. If the primary controller fails:

- 100% of queued calls are maintained.
- LA, NY, and Dallas Agent devices failover from the primary controller in New York to the secondary controller in New York.
- Agents are automatically logged into the secondary controller and commence answering calls. Their agent IDs are the same on both controllers.
- Agents have full ACD functionality on their secondary controller. The status of most ACD features is maintained after failover. See [Features Available to Resilient Hot Desk ACD Agents](#) for details.
- The Digital Trunking Gateway continues to queue calls on the secondary controller.
- Any calls that are active prior to the primary controller outage are allowed to complete. To ensure that calls in talk state are not dropped, an agent's set does not fail over until after the agent hangs up. Calls on hold are lost.

Primary Controller Returns to Service

When the primary controller in [Figure 26: Resilient Virtual Contact Center - Configuration 2](#) on page 161 returns to service:

- Digital trunking gateways detect that the primary controller is in service and once again resume queuing calls on the primary and secondary controllers.
- NY, LA, and Dallas agents home back to their primary in NY when idle and health check is complete.
- Any calls that are active on the secondary controller will be allowed to complete. After the agent hangs up, their set will failback. This ensures calls in talk state are not dropped. Held calls are lost.
- If the primary controller fails, calls that are queued to a resilient agent skill group on the primary are rerouted to its resilient agent skill group on the secondary controller. After the primary returns to service, calls are automatically rerouted from resilient agent skill group on the secondary controller to the resilient agent skill group on the primary controller.
- When agent devices fail back from the secondary to primary, the agents are automatically logged in on the primary controller.
- Calls that were queued on the secondary controller while the primary was down will be redirected to the primary controller.

Secondary Controller

If the secondary controller in [Figure 26: Resilient Virtual Contact Center - Configuration 2](#) on page 161 fails there is no impact to the call center because no agents reside on the secondary controller under normal operating conditions.

Advantages

Configuration 2 ([Figure 26: Resilient Virtual Contact Center - Configuration 2](#) on page 161) provides the following advantages:

- Agents and supervisors in different geographies can be represented in the same agent skill groups to create a virtual contact center environment.
- Support for this configuration is provided by the 3300 ICP. An external application is not required to support longest idle agent routing.

When either ACD agent controller fails:

- 100% of agents remain in service
- 100% of trunks remain in Service
- Agents can complete 100% of active calls before failover. Held calls are lost.
- Up to 350 active agents are supported on the primary or secondary controller.
- Local dial '9' and '911' service is available to Los Angeles, New York, and Dallas agents through ARS programming.
- No single point of failure.
- A Digital Trunking Gateway failure drops calls that are queued through that gateway but does not take any agents out of service. Local PSTN access (local dial '9' and '911' service) through that gateway is unavailable for the outage period; however, alternate routing via ARS is available to agents.

Conditions

The following conditions apply to Configuration 2:

- Current route list programming permits only five remote sites to be part of a virtual contact center, that is, a sixth remote site would not have Local dial '9' and '911' service. The agents on the sixth remote site must route out of New York. This issue can be overcome using loopback IP trunks and the Default Account Code feature.
- A WAN failure, where LA and Dallas agent devices lose link to New York, takes LA and Dallas agents out of commission for the outage period because they will not be able to home to their secondary. Redundant IP routing paths (NLPS Network Multilayer) may overcome this issue.
- Agents have full ACD functionality on their secondary controller. The status of most ACD features is maintained after failover. See [Features Available to Resilient Hot Desk ACD Agents](#) for details.
- The time displayed on telephone sets will show New York time zone for all agents.

3.4.5 Local Call Breakout in a Virtual Environment including E911

In a virtual contact center configuration, IP phones are registered to a primary 3300 ICP in New York. As a result, ARS (Automatic Route Selection) is programmed in New York to ensure that outgoing external calls (especially E911 calls) for hosted phones are routed through IP trunking to the appropriate gateway (Configuration 1, [Figure 25: Resilient Virtual Contact Center - Configuration 1](#) on page 157) or to the appropriate

gateway controller (Configuration 2, [Figure 26: Resilient Virtual Contact Center - Configuration 2](#) on page 161). ARS route lists are used for this purpose. Since ARS is limited to six ARS routes in a single route list, and one route is assigned for local dialing, then this solution is limited to either

- five 3300 secondary controllers per 3300 ICP head office (in this example New York, Configuration 1), or
- five 3300 Trunking Gateways per 3300 ICP head office (in this example New York, Configuration 2).

This restriction can be overcome using the same techniques (loopback IP trunks and the default account code feature) as described in the Mitel Knowledge Base article 05-5191-00042 “How to route 911 to more than 6 Analog Gateways”.

The following is a general programming guide. Actual implementation may vary from site to site.

Note:

The purpose of these steps is to ensure that the hosted IP phones that reside at the same geographic location as the Gateway will select the proper trunk for local and 911 calls

In the explanation that follows, you can substitute the term “Gateway” for “Dual Purpose 3300 ICP” which is applicable to Configuration 1 described above where the controller serves as both the agent’s secondary controller and a trunking controller.

1. Set up an IP trunk between the primary PBX and all remote gateway(s).
2. On the primary PBX, create a route for each gateway (A maximum of six routes: five routes for remotely situated gateways plus one local route, ensuring that each route uses a unique COR.)

Note:

The selection of route for the appropriate gateway is based on COR.

3. Set up a route list that contains all routes created in step 2 (5 gateway routes plus 1 local route).

4. Program ARS leading digits for both local access and local 911 and assign the route list created in step 3.

 Note:

It is very important that the proper COR is assigned to each route and the corresponding IP phones. The routing mechanism to the proper gateway location is based on the COR of the set and the COR group defined in each route defined in the Head Office.

5. On each gateway, make sure that the ARS leading digits for both local and 911 are created and that the local trunk is selected to process local and 911 calls.

This chapter contains the following sections:

- [About this Chapter](#)
- [Resilient Call Routing ARS](#)
- [Call Forwarding to Resilient Devices](#)

4.1 About this Chapter

This chapter provides general information about resilient call routing , and automatic route selection (ARS) and forwarding to resilient devices. It is recommended that you use call rerouting rather than forwarding in a resilient system.

4.2 Resilient Call Routing ARS

The following points describe how a call is routed to a device in a resilient cluster by resilient call routing, also known as automatic route selection (ARS).

- When a call is initiated ARS will try all non-resilient routes to the primary system, to establish a connection with the primary system.
- If the system determines that a connection to the primary system cannot be completed via any of the defined non-resilient ARS routes, resilient routing is invoked to identify routes to the secondary system.
- Resilient routing uses the RDN table to identify routes to the secondary system and then ARS is used to route the call to the secondary system.
- Once the call is successfully connected with the MiVoice Business system (either the primary or the secondary system) the system checks the status of the called device.
- If the called device is in service, the MiVoice Business system completes the call to the device; however, if the called device is not in service, the system attempts to locate the device on its other MiVoice Business system (either primary or secondary) with a virtual call. This virtual call is an Extension Status Query.

4.2.1 General Guidelines for Resilient Routing ARS

Since resilient call routing will first attempt all ARS routes to the primary system for a device, then will attempt ARS routes to the secondary, ARS configuration of the cluster nodes is critically important for correct and efficient operation of routing at the cluster

level. Keep the following guidelines in mind when planning and implementing resilient routing/ARS:

- It is recommended that you have a fully meshed IP Trunking environment.
- At each node, ARS routes to both the primary and secondary systems serving any resilient device must be in the configuration tables.
- For enhanced resiliency it is desirable to have alternate TDM (non-IP) routes to any high priority primary and secondary systems in the cluster, this will reduce the likelihood of total failure due to IP network failures.
- Resilient call routing is independent of whether the route is over IP trunking/DPNSS or TDM/DPNSS; therefore, when both TDM and IP trunks are available, the system administrator should consider entering both paths into the ARS route.

4.2.2 Remote Directory Numbers Form

Resiliency augments cluster routing information by programming an alternate or secondary MiVoice Business system for each resilient device or directory number (DN). The Remote Directory Number (RDN) table holds a second cluster element identifier (CEID) to support this resilient programming. Each 3300 ICP Release 4.0 or later system in a resilient cluster will have two CEID values for each resilient device. Refer to the *Voice Networking -> Configure Network* book in the *System Administration Tool Help* for details .

Basic non-resilient cluster programming excludes the DNs of local devices from the RDN table; however, for resilient call routing, all resilient devices must have entries in the RDN table on all MiVoice Business systems in the cluster. This information must be on all Release 4.0 3300 systems in order for a call to be redirected to the correct system.

The RDN tables for the cluster elements are displayed in the Remote Directory Number (RDN) Assignment forms. The RDNs are distributed to all the elements in the cluster via Remote Directory Number Synchronization.

The RDN table includes cluster element identifiers (CEIDs) for both the primary and secondary systems of all resilient devices that are visible to the systems in the resilient cluster. The RDN table is used to identify alternate systems, to check for availability of the device, and to determine whether the device is located on the primary or secondary system .

[Remote Directory Numbers Form](#) shows the Cluster Element Index (home or primary system) and Secondary Cluster Element Index (secondary system) numbers in the Remote Directory Numbers form.




Figure 27: Remote Directory Numbers Form

Basic Requirements to Support Resilient Call Routing

The following basic requirements are needed to support a call to a resilient device:

- The Remote DN Table (RDN) must include both the primary and secondary system references (CEIDs) for all resilient-configured devices, this information must be visible to all 3300 ICP Release 4.0 or higher systems in the cluster.

The Remote DN Table is used to identify alternate MiVoice Business systems to check for availability of devices, and to determine if a device is located on the primary or secondary system.

- At least one node in the route to the called device must be resiliency-capable (Release 4.0 or higher) to support boundary node routing which is used for mixed clusters and out-of-cluster calls. The first such node in the route becomes the boundary node.

The boundary node handles selection and checking of primary and secondary system availability in the case of call attempt failure.

For calls originating in mixed clusters and from outside a resilient cluster, at least one MiVoice Business system along a route to a resilient device must be Release 4.0 or higher. This node becomes a “Boundary Node” that handles the selection and checking of primary and secondary system availability in the case of call attempt failure.

Establishing Boundary Nodes

A boundary node is the first 3300 ICP Release 4.0 (or later) system in a resilient cluster that a call encounters as it is routed to its destination. It can also be the ICP that originates the call or that is the call destination. The boundary node routes the call to the ICP that a device is in service on, using normal automatic route selection (ARS).

A boundary node must

- Be placed on routes for calls that are made from an external cluster or external ICP to a destination inside a resilient cluster
- Be used to route intra-cluster calls originating from a pre-Release 4.0 3300 ICP
- Have proper ARS programmed

Call Routing by a Boundary Node

For examples of call routing by a boundary node, see the following sections of this document:

- [Figure 29: Routing to IP Device in Service on Secondary System \(Call Originating Inside Cluster\)](#) on page 175
- [Figure 30: Routing to Device in Service on Secondary System \(Call Originating Outside Cluster\)](#) on page 176
- [Figure 31: Routing for Device Out of Service on Both Primary and Secondary Systems](#) on page 177

Identifying a boundary node MiVoice Business system to a MiVoice Business system *that is external to a resilient cluster*

The boundary node is the first 3300 ICP Release 4.0 system in the route to the destination. A boundary node is required whenever a call is made from an external cluster or external MiVoice Business system to a destination inside a resilient cluster, and whenever an intra-cluster call originates from a pre 3300 ICP Release 4.0.

Network administrators must consider the following guidelines when configuring the system:

- The external MiVoice Business system (or cluster) must be configured such that its ARS call routing will route to at least one 3300 ICP Release 4.0 (or higher) system in the destination resilient cluster. This will cause this Release 4.0 (or higher) system to become the “boundary node” for the resilient cluster.

It is recommended that the external MiVoice Business system be programmed so that the boundary node MiVoice Business system will be the first choice in the external ARS routing table.

- ARS routes originating from any pre- 3300 Release 4.0 ICP in the cluster must contain at least one 4.0+ system in the routes to all other ICPs in the cluster. This will cause this Release 4.0 (or higher) system to become the “boundary node” for the resilient cluster.

It is recommended that the pre-4.0 system be programmed so that the boundary node system will be the first choice in the ARS routing table.

Note:

Since the SX-2000 LIGHT is capable of running the LIGHTWARE software load which is compatible with the 3300 Release 4.0 (LIGHTWARE 32 Release 1.1), such nodes may also participate in resilient call routing as a boundary node, where the boundary is not also a primary or secondary in the call.

Out of Service Handling for Resilient Devices

If a resilient device becomes disconnected or loses power, out of service (OOS) processing is activated on the highest functional and reachable MiVoice Business system

². OOS processing is applied in the following situations:

- Resilient call routing fails if both MiVoice Business systems (primary & secondary) are functional and reachable but the device is not connected to either of them. At this point, the OOS processing is activated on the primary system.
- OOS processing on the secondary system occurs in cases of glare, when the resilient device has homed back to its primary system after being in service on its secondary system, just as the call is being routed to the secondary.
- If the primary MiVoice Business system has failed or is unreachable, and the secondary system is functional and reachable but the device is not registered on it, OOS processing is handled on the secondary system.

Note:

For information about out of service handling scenarios and programming for an IP Console, see ["IP Console Resiliency"](#).

² A MiVoice Business system is considered non-functional if its database is missing RDN data for any device requesting to register with it. Thus, a MiVoice Business system in the process of returning to service (for example, after an offline upgrade) can only accept set registration requests after its database has been completely restored. A set attempting to register with this system before it is functional will be stuck in PIN entry mode and unable to home to its secondary system. Once the primary is functional, PIN entry can resume to complete the registration process.

- For Containers-based MiVoice Business deployments, an Out of Service message is sent when a set is changed from a resilient set to a non-resilient set, or vice-versa using the User and Services Configuration form in the System Administration Tool.

4.2.3 Resilient Call Routing Operation

In a resilient cluster, calls are routed to the MiVoice Business system (primary or secondary) that a device is in service on. [Minimum Resiliency Configuration](#) illustrates the minimal configuration required for resilient call routing.




Figure 28: Minimum Resiliency Configuration

It is recommended that PSTN TDM or analog/digital trunks be connected to both the primary and secondary systems (see [Figure 14: IP-trunked Resilient Cluster](#) on page 104). A customer must subscribe to two CO trunk spans: one span to each MiVoice Business system , with both placed in the same CO hunt group. For example, for a particular listed directory number (LDN) the CO would select from trunks in both spans. Alternatively, if customers require only a few trunks, they can lease fractional T1 or E1.

Call Originating on Primary System

The following two routing sequences describe how a call originating on the primary system of the destination device is routed, depending on whether the destination device is last known to have been in service on its primary or secondary system at the time of the call.

In [Minimum Resiliency Configuration](#), CEID 1 is the primary, and CEID 2 is the secondary system of extension 102. If extension 100 (or a trunk on CEID 1) dials extension 102, CEID 1 attempts to reach the device at the MiVoice Business system (primary or secondary) on which the system is aware that 102 was last registered on.

Call Originates on CEID 1 102 Last in Service on CEID 1 Primary

1. CEID 1 tries to reach 102 locally.
2. If 102 is in service there, CEID 1 handles the call.

3. If the seize fails because the device has gone out of service on CEID 1, CEID 1 attempts to find the device on its secondary system (CEID 2) by querying the status of extension 102 on CEID 2.
4. If 102 is in service on CEID 2, CEID 1 routes the call to CEID 2. (It is possible that 102 goes out of service on CEID 2 and is in the process of rehoming to CEID 1 when the call is routed to CEID 2 by CEID 1. In this unlikely case of glare, when the call is routed to CEID 2, the call fails, and CEID 2 applies out-of-service handling as configured for the extension. At this point, call control on CEID 1 tries to seize 102 as a local device and is either successful or receives a busy signal from 102 if it has not finished re-registering. A busy signal indicates a case of glare.)
5. If 102 is out of service on CEID 2, CEID 1 applies out-of-service handling as configured.

Note:

If extension 100 (or trunk on CEID 1) has Intercept handling configured for 'Directory number out of service', in above scenario, the call will initially be routed to extension 102 on CEID 1, then to the CEID 2. If 102 is out of service on both Primary and Secondary, the Intercept Handling for Directory number out of service will be activated.

Call Originates on CEID 1 102 Last in Service on CEID 2 Secondary

1. CEID 1 tries to route the call to CEID 2.
2. If 102 is in service on CEID 2, CEID 2 handles the call.
3. If this seize fails because the device is out of service on CEID 2, CEID 1 attempts to find the device locally on CEID 1.
4. If 102 is in service locally on CEID 1, CEID 1 handles the call.
5. If 102 is out of service on CEID 1, CEID 1 applies out-of-service handling (configured).

Call Originating on Secondary System

The following two routing sequences describe how a call originating on the secondary MiVoice Business system of the destination device is routed, depending on whether the destination device is last known to have been in service on its primary or secondary system at the time of the call.

In [Minimum Resiliency Configuration](#), CEID 1 is the primary, and CEID 2 is the secondary system of extension 102. If extension 200 (or a trunk on CEID 2) dials

extension 102, CEID 2 attempts to reach the device at the MiVoice Business system (primary or secondary) on which the system is aware that 102 was last registered on.

Call Originates on CEID 2 102 Last in Service on CEID 1 Primary

1. CEID 2 routes the call to CEID 1.
2. If 102 is in service on CEID 1, CEID handles the call.
3. If the seize fails because 102 goes out of service on CEID 1 when the call is routed there (and is in the process of failing over to CEID 2) or because CEID 1 is unreachable or out of service, CEID 2 attempts to find the device locally on CEID 2.
4. If 102 is in service on CEID 2, CEID 2 handles the call.
5. If 102 is out of service on CEID 2, CEID 2 does one of the following two things:
 - If CEID 1 is reachable, CEID 2 routes the call to CEID 1 for configured out-of-service handling for 102.
 - If CEID 1 is unreachable or out of service, CEID 2 applies configured out-of-service handling for 102.

Call Originates on CEID 2 102 Last in Service on CEID 2 Secondary

1. CEID 2 tries to reach 102 locally.
2. If 102 is in service on CEID 2, CEID 2 handles the call.
3. If the seize fails because 102 is out of service on CEID 2, CEID 2 queries the status of the extension on CEID 1.
4. If 102 is in service on CEID 1, CEID 2 routes the call to CEID 1, and CEID 1 handles the call. If the seize fails on CEID 1 because 102 goes out of service on CEID 1 when the call is routed there (and is in the process of failing over to CEID 2), the call fails, and CEID 2 attempts to reach 102 locally.

If 102 is on CEID 2, CEID 2 handles the call.

If 102 is out of service on CEID 2, CEID 2 does one of the following two things:

- If CEID 1 is reachable, CEID 2 routes the call to CEID 1 for out-of-service handling.
 - If CEID 1 is unreachable, CEID 2 applies out-of-service handling, as configured, to the call.
5. If 102 is out of service on CEID 1 and CEID 1 is reachable, CEID 2 routes the call to CEID 1 for out-of-service handling.
 6. If CEID 1 is unreachable or out of service, CEID 2 applies out-of-service handling, as configured, to the call.

4.2.4 Basic Resilient Call Routing Scenarios

There are two basic failure scenarios to consider when planning and programming resilient call routing:

- The IP device is in service on the secondary system.
- The IP device is out of service on both the primary and secondary system.

IP Device in Service on Secondary System

If a call is placed to a resilient device that has failed over to and is in service on its secondary system, the boundary node first tries to route the call to the primary system, using normal automatic route selection (ARS). Upon discovering that the device is out of service, the boundary node invokes resilient routing and routes the call to the destination device's secondary system to see if the device is in service there. If it is, the call is completed.

Routing to IP Device in Service on Secondary System (Call Originating Inside Cluster)

illustrates the steps involved in resilient routing for routing a call placed to a device whose primary system is out of service. In this scenario, the call is originated within the resilient cluster, and the destination device (102) is in service on its secondary system.




Figure 29: Routing to IP Device in Service on Secondary System (Call Originating Inside Cluster)

Routing to Device in Service on Secondary System (Call Originating Outside Cluster)

illustrates the steps involved routing a call to a resilient device whose primary system is out of service. In this scenario, the call is originated from outside the resilient cluster, and the destination device (102) is in service on its secondary system.





Figure 30: Routing to Device in Service on Secondary System (Call Originating Outside Cluster)

Device Out of Service on Both Primary and Secondary System

If a device's primary and secondary systems are functioning properly but the device's links to these systems has failed, the device will be out of service on both systems. When normal ARS and resilient routing are invoked, the systems detect that the device is OOS and inform the boundary node.

[Routing for Device Out of Service on Both Primary and Secondary Systems](#) illustrates the steps involved in routing a call placed to a resilient device that is out of service on both its primary and secondary systems.



IP0534

Figure 31: Routing for Device Out of Service on Both Primary and Secondary Systems

4.3 Call Forwarding to Resilient Devices

Note:

Resilient call rerouting is the recommended configuration for a resilient system.

This section describes how calls are forwarded to resilient devices.

Call rerouting and call forwarding are configured in the same way and exhibit similar behavior. Call forwarding builds on call routing. Note that the basic resilient routing described in [Resilient Call Routing Operation](#) on page 172 also applies to call forwarding to resilient devices.

4.3.1 Call Forwarding Terminology

The following terms are used in the figures and scenarios presented in this section to identify the various devices involved in a call forwarding scenario:

- “Calling device” is the device from which a call is placed.
- “Forwarding device” is the non-resilient device to which a call is placed and that forwards the call to a resilient destination device.
- “Destination device” is the resilient device to which a call is forwarded by a non-resilient forwarding device.

4.3.2 Minimal Configuration

A minimum of two 3300 ICPs (3300 ICP Release 4.0 or higher) or MiVoice Business systems are required for resilient call forwarding. For example, [On System Call Forwarding Minimal Configuration](#) on page 178, shows one resilient device, extension 1002, that is configured with CEID 1 as its primary system and CEID 2 as its secondary system.

There are three kinds of call-forwarding scenarios for calls placed to resilient devices:

- On–MiVoice Business system call forwarding
- Non–home call forwarding
- Off–MiVoice Business system call forwarding

4.3.3 On System Call Forwarding Minimal Configuration

On- system call forwarding refers to cases in which the same MiVoice Business system hosts all three devices involved in a basic call forwarding scenario: the calling, forwarding, and destination devices. On-system forwarding requires that at least the forwarding and destination devices be hosted by the same MiVoice Business system .

This type of call forwarding can handle calls to devices that have the following features programmed:

- Call Forward Always
- Call Forward Busy
- Call Forward No Answer
- Out of Service
- Call Rerouting
- Follow Me
- I Am Here

[On-System Call Forwarding Minimal Configuration](#) illustrates the minimal configuration for resilient call forwarding, which can be used for on- system forwarding.




Figure 32: On-System Call Forwarding Minimal Configuration

Calling Forwarding and Destination Devices on the Primary System

Refer to ["On-System Call Forwarding Minimal Configuration"](#). In this scenario, extension 1002 is in service on its primary system (CEID 1).

1. Extension 1004 dials 1003 and is forwarded to 1002. CEID is the only MiVoice Business system involved in the call between the calling, forwarding, and destination devices.
2. CEID 1 attempts to reach 1002 as a local device. If 1002 is in service on CEID 1, the call is completed.

If 1002 has gone out of service on CEID 1 (and may have homed to CEID 2), the attempt to reach 1002 fails. At this time

- Off-system forwarding is applied, and the call forwarding or call diversion model is sent from CEID 1 to CEID 2 to attempt to reach 1002 on its secondary system (see ["Off System Call Forwarding"](#)).
- Resilient routing to CEID 2 (secondary system) is invoked by CEID 1, which would be acting as the boundary node in such a scenario.

Calling Forwarding and Destination Devices on the Secondary System

Refer to ["On-System Call Forwarding Minimal Configuration"](#). In this scenario, extension 1002 is failed over to and in service on its secondary system, CEID 2.

1. Extension 1000 dials 1001 and is forwarded to 1002. CEID 2 is the only MiVoice Business system involved in the call between calling, forwarding, and destination devices.
2. CEID 2 attempts to reach 1002 as a local device. If 1002 is in service on CEID 2, the call is completed.

If 1002 has gone out of service on CEID 2 (and may have rehomed to CEID 1), the attempt to reach 1002 fails. At this time

- Off-system forwarding is applied, and the call forwarding or call diversion model is sent from CEID 2 to CEID 1 to attempt to reach 1002 on its primary system (see ["Off System Call Forwarding"](#)).
- Resilient routing to CEID 1 (primary system) is invoked by CEID 2, which would be acting as a boundary node in such a scenario.


4.3.4 Non home Call Forwarding

In non-home call forwarding, the calling device is hosted by another MiVoice Business system that does not also host either the forwarding or destination devices. This section discusses resilient call forwarding in which the call-forwarding model is configured on the forwarding device's MiVoice Business system, and the forwarding and destination devices are hosted by the same system. Both call forwarding and call diversion scenarios are described.

This type of call forwarding can handle calls to devices that have the following features programmed:

- Call Forward Always
- Call Forward Busy
- Call Forward No Answer
- Out of Service
- Call Rerouting

Figure 33: Non-home Call Forwarding



Forwarding and Destination Device on Primary System

Refer to [Non home Call Forwarding](#) on page 180. In this scenario, the forwarding and destination device are in service on the primary system of the resilient (destination) device.

1. Extension 1000 on CEID 3 dials 1003 on CEID 1 and is forwarded to 1002, also on CEID 1.
2. In accordance with on-system call forwarding, CEID 1 attempts to reach 1002 locally.

If 1002 has gone out of service on CEID 1 (and may have homed to CEID 2), the attempt to reach 1002 through on-system forwarding fails. In this case, the primary system, CEID 1, acts as a boundary node and sends the call forwarding model to CEID 2 (the secondary system of 1002) to attempt to reach 1002 there. If 1002 is in service on CEID 2, off-system forwarding is invoked (see [Off System Call Forwarding](#) on page 182).

In the call forwarding model:

- The call forwarding model is sent to CEID 2.
- Resilient routing to the primary system (CEID 1) is invoked as described in [Call Originating on Primary System](#). The resiliency information also described in this section accompanies the call forwarding model sent to CEID 2.

In the call diversion model:

- CEID 1 invokes call diversion to the call originating MiVoice Business system (CEID 3) with resilient routing information and requests that CEID 3 divert the call to the destination party on its secondary system (CEID 2). This call diversion request contains resilient routing information.

If the call originating element (CEID 3) is a 3300 ICP Release 4.0 or Release 3.x for which the cluster has an element ID number, it already possesses the resiliency information required to divert the call to CEID 2. However, if the cluster does not have an element ID for CEID 3, and if CEID 3 is a 3300 ICP Release 3.x, the diverted call is routed to CEID 1 before being routed to CEID 2, using the basic call routing mechanism described in [Call Originating on Primary System](#).

Forwarding and Destination Device on Secondary System

Refer to [Non-home Call Forwarding](#). In this scenario, the forwarding and destination device are in service on the secondary system of the resilient (destination) device.

1. Extension 1000 on CEID 3 dials 1001 on CEID 2 and is forwarded to 1002, also on CEID 2.
2. In accordance with on- MiVoice Business system call forwarding, CEID 2 attempts to reach 1002 locally.

If 1002 has gone out of service on CEID 2 (and may have homed to CEID 1), the attempt to reach 1002 through on- MiVoice Business system call forwarding fails. In this case, the call routing and call diversion behavior is the same as that described above, in [Forwarding and Destination Device on Primary System](#), with the secondary system acting as the boundary node in the call.

4.3.5 Off System Call Forwarding

In off-system call forwarding or rerouting the three devices involved in the call forwarding operation (calling, forwarding, and destination) are hosted by different MiVoice Business systems. When the destination device is resilient, the destination device digits are translated into an RDN. [Off-System Call Forward and Reroute](#) illustrates an off-system call forwarding scenario in which the destination device is resilient.

In this scenario, the primary and secondary systems of the resilient destination device must be 3300 ICPs (Release 4.0 or higher) or MiVoice Business systems. The elements hosting the other devices can be running any 3300 ICP or MiVoice Business release.




Figure 34: Off-System Call Forward and Reroute

Off System Call Forwarding and Rerouting

Refer to ["Off-System Call Forward and Reroute"](#).

1. Extension 1000 on CEID 4 dials 1002 on CEID 3 and is forwarded by off-system call forwarding to extension 1001, which is translated into an RDN.
 - The calling system is CEID 4.
 - The forwarding system is CEID 3.
 - The destination system is CEID 1, but this can change to CEID 2, depending on where the device is in service.
2. A new boundary node is established (replacing any prior boundary node). For example, if 1002 is resilient, and CEID 3 or CEID 4 is Release 4.0, one of these becomes the new boundary node.

3. The fact that 1001 is now an RDN may invoke resiliency-specific behavior, depending on the software release running on CEID 3. Other behavior depends on the configuration of the forwarding model.

Call forwarding model:

- a. The forwarding call is made from CEID 3 without informing the call originating system (CEID 4).
- b. If the forwarding system (CEID 3) is a 3300 ICP Release 4.0 (or higher), CEID 3 treats the forwarded call as a new call and invokes a new boundary node for the call, replacing CEID 4 as the boundary node (if it was previously established as such when the call originated there) and clearing system knowledge of the previous boundary node for the original call. Note that if the call is forwarded for Always or Busy, CEID 3 sends a DPNSS Clear Request Message back to CEID 4, and if the call is forwarded on No Answer, CEID 3 continues to ring 1002.
- c. If the forwarding system (CEID 3) is a 3300 ICP Release 3.x, it treats the forwarded call as a new call, but as a Release 3.x ICP, CEID 3 is not resilient and would not establish a new boundary node for the new call. If no boundary node is established prior to this point (that is, if CEID 4 is not the boundary node for the original call), the forwarding behavior is the same as the routing described in ["Resilient Call Routing Operation"](#). However, if a boundary node has previously been established at CEID 4, CEID 3 sends the forwarding and boundary node information to CEID 1, preventing the creation of a new boundary node for the call. In some cases, this may prevent the call from being routed to the device on its secondary system, if it is failed over. Also, if the boundary node has been established prior to reaching CEID 3 where the forwarded call is treated as a new call, since CEID 3 does not clear the system of an existing boundary node, this could cause resilient routing to be performed on the original boundary node (for example, on CEID 4).

Call divert model:

- a. The forwarding system (CEID 3) sends a diversion request to the call-originating system (CEID 4) for diversion on Immediate, Busy, or No Reply.
- b. When the calling system (CEID 4) receives a Divert-Immediate or Divert-Busy message, it immediately clears down the call along the original path to CEID 3, thus clearing the call's resiliency information at all transit nodes and at the forwarding system (CEID 3).
- c. When the calling system receives a Divert-No-Reply message, it clears down the call only after diverting it. If the diversion fails, CEID 4 maintains the original call path to extension 1002 on CEID 3 until it is answered or cleared by the caller.
- d. The diverting call is made from the calling system (CEID 4). If CEID 4 is running 3300 ICP Release 4.0 (or higher), it clears the previous boundary node for the original call destination (1002 on CEID 3), if one has been established, invoke a new boundary node for the diverted call, and begin basic resilient call routing as described in ["Resilient Call Routing Operation"](#). If CEID 4 is running 3300 ICP

Release 4.0 and CEID 3 fails while CEID 4 is processing the new diverted call as the result of a Divert-No-Reply message from CEID 3, CEID 4 does not initiate resilient routing to the original destination (1002 on CEID 3).

This chapter contains the following sections:

- [Introduction](#)
- [About Resilient Consoles](#)
- [Call Routing](#)

5.1 Introduction

The 5540 IP Console and MiVoice Business Console are supported as resilient devices.

This chapter provides information about:

- Resilient IP Consoles
- behavior
- resilient user interface
- operator and service status
- emergency call handling
- Call routing
- queue handling
- out of service handling
- known routing conditions

5.1.1 For More Information

For installation, programming, and end-user information on the IP Consoles, refer to the following documents, at Document Center:

- *5540 IP Console Installation Guide* (look under “Installation Guides”)
- *MiVoice Business Console Installation Guide*
- *5540 IP Attendant Console User Guide* (look under “User Guides”)
- *MiVoice Business Console Installation Guide*
- *MiVoice Business System Administration Tool Help* (programming information)

5.2 About Resilient Consoles

5.2.1 Console Behavior

Like other resilient IP devices, if a resilient IP Console is hosting an active call stream when its primary MiVoice Business system or the link between the console and the MiVoice Business system fails, the console experiences call resiliency, that is, the call survives. The console does not fail over to its secondary MiVoice Business system until the call is ended and the console is in idle state (see [Emergency Call Handling](#) on page 186, and [Non-idle Devices](#)).

Console Application or PC Failure

If a MiVoice Business Console application fails, or if the entire PC fails, the only recourse to ensure continued operation is to use another PC preloaded and configured with the console application.

5.2.2 User Interface

Console features and functionality are available only when the console is in service on a MiVoice Business system. Keys and features do not function during call survival or while the console is in the process of registering with an MiVoice Business system. Any attempt at feature activation during a resilient call or registration prompts an error message on the console screen. For information about resilient console screen messages, refer to the console user guide.

5.2.3 Operator Status

The console resumes the same operator status (Present / Absent) it was in before failing over or back, with the exception that if the console resets when it re-registers, the operator status defaults to Absent.

5.2.4 Service Status

Service status (Day, Night1, Night2) is controlled by call control on each MiVoice Business system. The console assumes the current service status on any MiVoice Business system it registers with.

5.2.5 Emergency Call Handling

The ability to monitor and respond to emergency calls to the IP Console is a vital part of the MiVoice Business solution.

A resilient IP Console can respond to emergency calls while in service on both of its primary and secondary systems. If the console is in service on its secondary system and its primary system recovers, it rehomes to its primary as soon as it becomes idle; however, if the console is a designated emergency-call monitor and has unacknowledged emergency calls, it will remain on the secondary system until all emergency calls have been acknowledged, before failing back to its recovered primary system.

The IP Console is considered not-idle if it meets any of the following conditions:

- The console is in talk or dialing state.
- The console is an emergency-call monitor and there are unacknowledged emergency calls in the system.

i Note:

If there are no unacknowledged emergency calls, IP console Fail-back is allowed to proceed even if there are calls ringing the console. These calls are routed to the console on the primary system once the console has rehomed.

Also see [Non-idle Devices](#).

The following conditions do not prevent the console from rehoming:

- The console is using Guest Services.
- The console is viewing Trunk Status.

5.3 Call Routing

Call routing for the console is the same as that described for IP devices in [“About this Chapter”](#); however, extra consideration must be given to the following topics when planning and programming call routing for resilient consoles:

- Queue handling
- Out of service handling
- Known routing conditions for consoles

5.3.1 Queue Handling

Active IP Console calls and incoming new calls behave in the same way as active IP phone calls (see [About this Chapter](#) on page 167). The behavior of calls queued to the resilient console is described in the following sections.

Queued calls are cleared down if either the console's primary MiVoice Business system or the link to it fails; these calls are only preserved when the console fails back to its primary from its secondary system. In a Fail-back situation, queued calls follow the console back to its primary system, but these calls do not retain priority queueing.

5.3.1.1 Example

Assume that Resilient IP Console 102 is currently homed at the primary ICP CEID 1, with 10 calls queued in the call waiting queue and two calls in the hold queue. Each of the calls came in through ICP CEID 3 on the A-B link.




Figure 35: IP Console Queue Handling

In [IP Console Queue Handling](#), the A-B link is an IP trunk. If the A-B link is a TDM trunk, the call would be cleared down.

Link A-B fails

Hold Queue

Calls clear down at CEID 3, CEID 1, and on 102.

Call Waiting Queue

Calls clear down at CEID 3, CEID 1, and on 102.

ICP CEID 1 fails

If the console's primary ICP (CEID 1) fails, console 102 fails over to its secondary ICP (CEID 2) when the console is idle. Queued calls are cleared down.

Hold Queue

Calls clear down at CEID 3.

Call Waiting Queue

Calls clear down at CEID 3.

Link CEID 1 – 102 fails

If the link between the console and its primary ICP (link CEID 1 – 102) fails, the console fails over to its secondary ICP, CEID 2 when the console is idle. Queued calls are cleared down.

Hold Queue

The Hold timer continues running. Upon expiry, resilient routing is attempted at CEID 1 or CEID 3, wherever the hold timer is configured to run. If the routing is unsuccessful, exception handling begins.

Call Waiting Queue

Resilient routing is attempted on a call-by-call basis at CEID 1. If the routing is unsuccessful, exception handling begins.

Note:

If the link comes back up just as resilient routing is failing and exception handling begins, exception handling continues.

Console 102 fails

Hold Queue

The Hold timer continues running. Upon expiry, resilient routing is attempted at CEID 1 or CEID 3, wherever the hold timer is configured to run. If the routing is unsuccessful, exception handling begins.

Call Waiting Queue

Resilient routing is attempted on a call-by-call basis at CEID 1. If the routing is unsuccessful, exceptional handling on the primary ICP begins, since the console is out of service (OOS) on both ICPs.

Note:

If the console rehomes to CEID 1 just as resilient routing is failing and exception handling begins, exception handling continues.

Console 102 fails back from CEID 2 to CEID 1

When CEID 1 recovers, console 102 fails back to CEID 1 (its primary ICP) only when the console is idle and regardless of the number of calls it has in queue.

Note:

This Fail-back is not guaranteed since it is possible that while CEID 2 is aware of the CEID 1 recovery, the console 102 may not be.

Hold Queue

The Hold timer continues running. Upon expiry, resilient routing is attempted at CEID 2 or CEID 3, wherever the hold timer is configured to run. If the routing is unsuccessful, exception handling begins.

Call Waiting Queue

Calls remain queued at CEID 2 with timers running. New calls are allowed to queue, but no further calls are presented to the attendant once the rehoming process has begun. Every five seconds, CEID 2 queries the status of console 102 on CEID 1. When CEID 1 is ready to accept calls, the queued calls on CEID 2 are routed to CEID 1 and queued there. If console 102 is not ready to accept calls on CEID 1, its status is checked on CEID 2, just in case the console has re-registered on CEID 2. If it has, queue processing continues on CEID 2; if it has not, the five-second timer resets.

Note:

When a call timer pertaining to a queued call expires, traditional processing for that timer applies for that particular call.

5.3.2 Out of Service Handling

It can take up to 3 minutes for all devices affected by a MiVoice Business system or network outage to fail-over to their secondary system. During this time, a console is out of service (OOS) and cannot receive new calls. Depending on whether you have a centralized attendant console or multiple consoles on different MiVoice Business systems, you can program out of service handling to ensure that calls coming in to the console during this time are answered at an alternate extension instead of being sent to voice mail.

See [Out of Service Handling for Resilient Devices](#).

Centralized Attendant Scenario

If you have one IP console that functions as a centralized attendant console, it is important to ensure the handling of new calls while the console is out of service. To do this, you can program daytime OOS handling on the console's secondary system to route calls to a hunt group (alternate extension on the secondary system) until the console is fully registered on the secondary. The hunt group can be any appropriate alternate extension such as a security desk that can handle calls and that calls can queue to while the console registers on the secondary system. A recorded announcement device (RAD) can be attached to this hunt group to ensure that the Fail-over situation is transparent to callers.

Calls that are queued to the hunt group extension cannot be programmed to route to the console when it comes back into service on the secondary system. These calls can only be answered by the hunt group.

Multiple Attendant Scenario

In a multiple attendant scenario you can program OOS handling for a console to route calls to another console rather than another extension that is not normally an attendant station (as in [Centralized Attendant Scenario](#)).

In [Out of Service Handling \(Multiple Consoles\)](#), resilient console 1000 fails over to a secondary ICP (ICP B) that is also the primary ICP for another console (2000). You

can program out of service handling to route calls for 1000 to 2000 until 1000 is fully registered on the secondary ICP. Calls queued to 2000 by OOS handling must be handled by 2000.

Note:

If you have multiple consoles, it is recommended that they be configured to have the same primary and secondary ICPs. This out-of-service handling scenario for multiple consoles is applicable in cases where multiple consoles do not have the same primary and secondary ICPs.




Figure 36: Out of Service Handling (Multiple Consoles)

5.3.3 Known Routing Conditions

DNI and IP Console Interaction

Setup

A DNI console and a resilient IP console are programmed on 3300 ICP A, which is the primary ICP of the IP console. (Refer to [DNI and IP Console Routing](#)) 3300 ICP B is the secondary ICP of the IP console. ICPs A and B are connected by TDM trunking. The consoles have an appearance of the same listed directory number (LDN), 7000.

Conditions

If the ethernet connection to the IP console's primary ICP is lost, the resilient IP console fails over to its secondary ICP (B). As a result, there is an appearance of LDN 7000 on the DNI console on A and also on the IP console on B. In this scenario,

- If a call is placed to LDN 7000 before the cluster is updated to show the IP console on its secondary ICP, the call is routed to the primary ICP. Since the DNI console has an appearance of 7000 and is on the primary ICP, the call is either completed or given out-of-service handling there. The call is not routed to the IP console on the secondary.
- If a call is placed to LDN 7000 when the cluster is updated to show the IP console on its secondary ICP, the call can be routed to the IP console on ICP B or to the DNI console on ICP A.




Figure 37: DNI and IP Console Routing

This chapter contains the following sections:

- [About this Chapter](#)
- [Network Configuration Types](#)
- [Recommended Resilient Topologies](#)
- [IP Device Distribution](#)
- [Planning a Resilient Network](#)

6.1 About this Chapter

This chapter provides an overview of the steps involved in planning a resilient network. Carefully consider the requirements, guidelines, and constraints provided in [Network Configuration Types](#) on page 194 before proceeding to install and configure your network.

6.2 Network Configuration Types

There are three types of network configuration that offer different levels of reliability. In order of increasing reliability, these configurations are

- Non-resilient
- Enhanced reliability
- Resilient

6.2.1 Non resilient Configurations

There are three basic non-resilient configurations:

- Standalone
- Distributed
- Hybrid

These configurations form the basis for the recommended resilient topologies (see [Recommended Resilient Topologies](#) on page 200).

Standalone Environment

[Non-resilient Standalone 3300 ICP Environment](#) illustrates a typical non-resilient, standalone configuration with IP phones connected to a single ICP that has PSTN connectivity. All phones are registered with the one ICP that provides call control and routes calls. The ICP in this figure may be running 3300 ICP Release 3.x or 4.0 software.

Solution description

- Inexpensive but vulnerable to ICP or network failure
- Hosts multiple applications on one ICP
- Can serve extensions in multiple locations (local and remote phones)

For illustrations of basic enhanced reliability and resilient topologies, see [Enhanced Reliability Standalone Network](#), and [Resilient Standalone 3300 ICP Environment](#).




Figure 38: Non-resilient Standalone 3300 ICP Environment

Distributed Network

In a distributed network, devices are divided between several ICPs. Non-resilient Distributed Network Hybrid illustrates three ICPs (any release) connected through a LAN/MAN. Only ICP A has PSTN access. As a result, A is a single point of failure for the cluster: if a failure occurs at A, all devices in this cluster lose PSTN access. Since devices are not provisioned with a secondary ICP, they go out of service for the duration of a failure affecting their ICP or their link to it.

This system can be made more reliable, while not being resilient, with the addition of a second or third PSTN connection at Node B and/or C. For example, if Node B also has PSTN connectivity and the ICP at Node A fails, devices at Node C that normally access

the PSTN through Node A can access the PSTN through Node B, if this alternate routing is programmed.

For even greater reliability, this system can be made resilient as in [Resilient Standalone 3300 ICP Environment](#).





Figure 39: Non-resilient Distributed Network Hybrid

Hybrid Network

A hybrid network combines the standalone configuration with a distributed network topology. [Non-resilient hybrid network](#) illustrates a hybrid network for a campus-type environment in which the network is distributed throughout separate departments in the same building and to which another department at a separate branch site is connected. In this example, each call control ICP has PSTN connectivity, ensuring that phones can find alternate routes to the PSTN in the event of network failures.

For information about a resilient hybrid network, see ["Resilient Hybrid Network"](#).



IP0514

Figure 40: Non-resilient hybrid network

6.2.2 Enhanced Reliability Configuration

While not a resilient configuration, an enhanced reliability network offers a greater degree of network reliability than a non-resilient network but less reliability than a resilient one. Increasing the reliability of a non-resilient system such as the basic standalone site in [Non-resilient Standalone 3300 ICP Environment](#), requires the addition of a second ICP (any release) with either the new Embedded Digital Trunk Module.

[Enhanced Reliability Standalone Network](#) illustrates an enhanced reliability standalone site in which the registration of devices is split between Nodes A and B. The two ICPs are clustered for dialing and feature transparency and so that alternate routing to PSTN trunks can be programmed between the ICPs. In this way, the enhanced reliability configuration can ensure PSTN access to devices on both ICPs in most cases.

Even though this type system is still vulnerable to ICP and network failures, it is more reliable than a non-resilient system since only half the devices are affected by any given outage. This solution is scalable, and the more ICPs you have, the fewer devices need to be registered on any one ICP. This distribution further minimizes the percentage of devices affected by an outage.




Figure 41: Enhanced Reliability Standalone Network

6.2.3 Resilient Configurations

Resilient networks may contain more MiVoice Business systems than a non-resilient network. These controllers can perform specialized, dedicated functions (see [Dedicating MiVoice Business Systems to Specific Functions](#)) or participate in call-control load sharing.

In the event of an ICP or network failure, devices in a resilient cluster recognize the failure and work around it: devices requiring backup call control Fail-over or home to their secondary ICP, controllers in the system identify the Fail-over and the new ICP of the failed over devices and then use resilient call routing to route around the point of failure and continue to provide service to these resilient devices at their new location.

The following basic resilient configurations offer a high degree of network reliability and are highly scalable:

- Minimum resilient configuration
- Resilient cluster with dedicated ICPs

Minimum Resilient Configuration

A minimum of two 3300 ICPs is required to implement basic, single-cluster resiliency and resilient call routing. To ensure that no single point of failure exists in the cluster, both ICPs must provide devices with PSTN access.

[Minimum Resilient Configuration](#) illustrates the minimum resilient cluster configuration. This cluster contains non-resilient phones and a resilient IP phone (extension 102) that has a primary and secondary ICP. Both ICPs can provide PSTN access.




Figure 42: Minimum Resilient Configuration

Resilient Cluster with Dedicated ICPs

The more MiVoice Business systems you have in a resilient system, the more options are available to you. [Resilient Standalone 3300 ICP Environment](#), builds on the minimum resilient configuration, showing the addition of two more ICPs to the cluster. It also shows the type of trunking required between ICPs within the cluster, between the resilient cluster and external ICPs/clusters, between ICPs and devices, and between ICPs and the PSTN. Note that a resilient cluster interacts normally with external ICPs and other resilient or non-resilient clusters. As a result, your network can contain any number of resilient and non-resilient clusters.

In this example, all ICPs are 3300 ICP Release 4.0 or later ICPs A and B share the call-control load for the devices in the cluster, while C and D are dedicated to other functions (see [Dedicating MiVoice Business Systems to Specific Functions](#)) but can also be programmed as secondary ICPs. Devices can be programmed to fail over to any Release 4.0 or later ICP in the cluster. Resilient configurations include:

- ICPs A and B function as primary and secondary ICPs. That is, devices on A can fail over to B and devices on B can fail over to A.
- Alternatively, the cluster can be configured so that devices on A and B can fail over to either C or D. Note that while C and D are dedicated to specific functions and do not function as primary ICPs in this illustration (although this is also possible), they can be programmed as secondary ICPs.




Figure 43: Resilient Cluster with Dedicated ICPs

6.3 Recommended Resilient Topologies

We recommend the following three highly-scalable, resilient topologies:

- Resilient standalone site (small to medium-sized businesses)
- Resilient distributed network (large enterprises)
- Resilient hybrid network (combination of standalone and distributed)

6.3.1 Resilient Standalone Environment

In a resilient standalone site, a number of MiVoice Business systems are clustered together, and each system can function independently if it has its own PSTN access. ICPs are connected through a local area network (LAN), via IP trunks or TDM trunks. This type of topology is suitable for a small to medium-sized business, organization, or institution that is a standalone site (no branch offices).

In [Minimum Resilient Configuration](#), the two ICPs share the call control load to ensure that a minimal number of devices are affected by any one failure. For example, if ICP A goes out of service, there is no impact to the devices on ICP B. If this resilient configuration is scaled upwards to contain more ICPs, the resilient devices can be distributed among more elements so that even fewer devices are affected by an individual ICP or network failure.

Resilient Standalone Environment builds on **Non-resilient Standalone 3300 ICP Environment** to show how to achieve the minimum resilient configuration required for the same standalone site. In accordance with the minimum resiliency requirements, the resilient cluster in **Resilient Standalone Environment** shows the addition of a second ICP with PSTN access.

PSTN trunk access is available from both ICPs. If alternate routing is programmed, devices on either ICP can access the PSTN through the other ICP if their own ICP is healthy but the link from their current ICP to the PSTN has failed. When planning a resilient installation, consideration must be given to providing PSTN access for all phones on their secondary ICP.

If all trunks on one of the ICPs are busy, the ICP can route to the other over IP networking to seize a trunk, if it is programmed to do so.




Figure 44: Resilient Standalone 3300 ICP Environment

6.3.2 Resilient Distributed Network (Local)

A locally distributed network connected through a local area network (LAN) or metropolitan area network (MAN) is ideal for larger enterprises, organizations, or institutions that consist of several local branch offices or departments. These branches can be dispersed locally throughout a city as in the case of a restaurant chain or throughout a campus environment, as in the case of an educational institution.

Note:

If connecting system elements across a MAN, they must be located within the same area code to ensure that local emergency services such as 911 calls remain available to devices that have failed over to secondary ICPs at other sites.

In a distributed system, you might have a number of ICPs in a central location, some functioning as dedicated TDM gateways and others as central voice mail servers for group controllers or ICPs that are dispersed throughout branch offices, buildings, or departments. For example, in a campus environment each building or department might possess a group ICP and local devices but obtain PSTN access and voice mail service from dedicated ICPs in a separate communications building.

Alternatively, each branch site could possess its own PSTN access, voice mail service, and group controllers and devices, so that it functions independently of the cluster except during ICP or network failures where devices could fail over to an ICP at another site.

All ICPs can be located on a single site to provide resiliency to an office or campus location, or the ICPs can be dispersed throughout a metropolitan area. In either case, storing the resilient hardware in separate locations minimizes the possibility of a total system outage caused by water damage or fire that can affect a more centralized system. For this reason, dispersed PSTN connectivity is less vulnerable than centralized PSTN connectivity.

[Resilient Distributed Network](#) builds on [Non-resilient Distributed Network](#), adding PSTN access at node B. Node B can provide PSTN access to the devices it hosts, and if automatic route selection (ARS) is programmed to do so, B can also provide PSTN access to devices at C. This example has the minimum requirement of two PSTN access points, but a third can also be installed at C for greater resiliency (optional).

Nodes A, B, and C form one resilient cluster in which devices from ICP A fail over to ICP C, devices from C fail over to B, and devices from B fail over to A. Note that each site can also become independently resilient with the addition of a second ICP, so that devices fail over to an ICP at the same site.

This solution is cost-effective because it makes use of existing ICP capacity to ensure that no single point of failure exists in the network and that no extra hardware is required to make the network resilient. All that is required is a second PSTN access point, and sufficient IP User licenses.

Part of the flexibility of a resilient network is that not all resilient devices on any given ICP must fail over to the same secondary ICP. They can be programmed to Fail-over to different ICPs in the resilient cluster. This further minimizes the effect of a failure in a given area of the network. For example, devices from Node A may be divided to fail

over to Nodes B and C. In the event that another failure affected Node B before Node A comes back into service, only those devices from A that are now on B are affected. The devices from Node A that are currently in service on Node C are not affected. The only licensing constraint is that each resilient device must have one IP User license.





Figure 45: Resilient Distributed Network

6.3.3 Resilient Hybrid Network

In a resilient hybrid network, you might have a main office with a distributed system (connected through a LAN throughout several departments) and one or more branch or standalone sites connected to the main network through a MAN. These branch sites can be part of the resilient cluster at the main office and have corporate PSTN access while also retaining PSTN access through a local group controller. In this way, branch site(s) can continue to operate independently and retain PSTN service for things like emergency services in the event of a failure of the LAN/MAN link to the main office.

Resilient Hybrid Network illustrates a resilient hybrid network in which the branch office has corporate and local PSTN access. The dotted lines illustrate optional resilient clustering in this topology. For example, sites A and B can both be resilient clusters; devices at site A can fail over to a secondary ICP at site A (same for site B). Alternatively, devices at site A can fail over to an ICP at site B and vice versa—in this way sites A and B are part of the same resilient cluster. Branch site C can be clustered with A and B to fail over to an ICP at the main office, or it can be independently clustered and resilient with the addition of a second ICP at the branch office.



IP0584

Figure 46: Resilient Hybrid Network

6.4 IP Device Distribution

As shown in [Distribution of IP Devices in an IP-trunk Resilient Cluster](#), IP phone and device distribution is flexible within a resilient cluster as long as

- All IP phones (whether resilient or not) are associated with a single primary ICP.
- Resilient IP Phones are also provisioned on a single secondary ICP for backup control.




Figure 47: Distribution of IP Devices in an IP-trunk Resilient Cluster

6.5 Planning a Resilient Network

Resiliency requires careful planning. The following table provides an overview of the steps that you should follow to plan your resilient network

Table 8: Planning a Resilient Network

Planning Step		Where to Find Information
1.	Familiarize networking guidelines.	Refer to the <i>3300 ICP Engineering Guidelines</i> for a summary of design guidelines, guidelines for maintaining voice quality of service, guidelines for maintaining the availability of connections, and the recommended DHCP Option settings and DHCP Lease Time settings. For more information about Resiliency network guidelines, see Familiarizing Yourself with the Engineering Guidelines on page 207, and

Planning Step		Where to Find Information
		Appendix Appendix: "Engineering Guidelines for Resiliency," .
2.	Identify a suitable resilient network topology.	See Identifying the Network Topology on page 208.
3.	Plan PSTN access.	See Resilient Call Routing ARS on page 167.
4.	Plan resilient licensing (IP User licenses).	See Planning License Requirements on page 214.
5.	Plan hot desking and resilient clustered hot desking requirements.	See Planning Resilient Clustered Hot Desking on page 216.
6.	Plan ACD Agent resiliency configuration.	See and ACD Resiliency Solution on page 106 and Resilient ACD Configurations on page 135 .
7.	Plan VLANs and DHCP options.	Refer to the <i>Network Guidelines for Voice over IP Installations</i> and the <i>3300 ICP Technician's Handbook</i> .
8.	Plan IP trunking, internal trunking, cluster programming and ARS.	Refer to <i>Voice Networking-> Configure Network</i> in the <i>System Administration Tool Help</i> for instructions on how to plan a cluster, and program the ARS routes, CEID forms, IP Networking and XNET forms, and configure the telephone directory forms. See Upgrading from Nodal Non Clustered Hot Desking for guidelines specific to call routing in resilient networks.
9.	Plan resilient hunt groups	See Hunt Group Resiliency on page 91. See Configure Resilient Hunt Groups on page 237.

Planning Step		Where to Find Information
10.	Plan resilient ring groups	See Ring Group Resiliency on page 96. See Configure Resilient Ring Groups on page 242.
11.	Plan T1/E1 trunk resiliency.	Refer to T1/E1 Trunk Resiliency on page 97 for a list of the conditions that apply to this feature.
12.	Plan voice mail and other system options.	See Feature Resiliency on page 48 for feature information and considerations; and see Planning Voice Mail and Other Options on page 223, for conditions that affect voice mail (type, setting up on primary and secondary ICPs). Also refer to the <i>MVoice Business System Administration Tool Help</i> for information on voice mailbox types, supported features, default settings and so forth.
13.	Plan wireless resiliency.	See Planning the DECT IP Wireless Solution (EMEA) Networks on page 225
14.	Consider known upgrading or migration issues.	See For example, if you need resiliency for the Riffs you can deploy two Riffs to provide the same radio frequency coverage. For further information on how to engineer and ensure resilient operation of the above components, please refer to the IP DECT Wireless Solution (EMEA) Site Survey Guide and IP DECT Wireless Solution (EMEA) Technical Manual..

6.5.1 Familiarizing Yourself with the Engineering Guidelines

Resiliency is an IP-based solution that relies on the capabilities of the IP network in which it is deployed. To provision a resilient system, you must identify and implement the proper LAN/WAN requirements for a given installation. For current LAN/WAN guidelines, see Appendix C: “Resiliency Engineering Guidelines” and *3300 ICP Engineering Guidelines* at Document Center.

Proper Layer 2 and Layer 3 infrastructure is critical to building a resilient network. For Layer 2, see [Appendix: "Engineering Guidelines for Resiliency"](#).

6.5.2 Identifying the Network Topology

By choosing appropriate ICP capacities (AX, LX, MX, CX(i), Mx, Mx Server, MiVoice Business for Industry Standard Servers, MiVoice Business Virtual, or MiVoice Business Multi-instance), you can ensure that a minimal number of devices are affected by a failure, and in this way, determine the potential size and effects of any given outage.

When planning your Resilient network topology, consider the points outlined in [Resilient Topology Considerations](#).

Table 9: Resilient Topology Considerations

Considerations		
1.	Cluster element requirements (for minimum resiliency configuration)	- Do you have a minimum of two 3300 ICPs (3300 ICP Release 4.0 or later) or MCD Release 4.x or later systems? This minimum is required to support Device resiliency because one primary and one secondary ICP is required for each device that is configured for resiliency. We strongly recommend that you upgrade all 3300 ICPs in your network to a minimum of 3300 ICP Release 4.0 or MCD Release 4.x. - You need at least two access points to the PSTN. One PSTN side should be configured as a hunt group so that if one access link (or the attached ICP) fails, the other access link can still receive calls through the PSTN side. Similarly, cluster ARS should be configured with the ability to access

Considerations		
		both PSTN access point for outgoing calls. For an example of resilient configuration, see Minimum Resiliency Configuration.
2.	Failure scenarios	• What is the potential impact on your network of a single failure of any one of the network elements (MiVoice Business systems)? • Do you have adequate redundancy of Layer 2 switches and Layer 3 routers? Any single failure of a Layer 2 switch or Layer 3 router will remove from service all MiVoice Business systems and IP phones that are dependent on these ICPs for IP connectivity.
3.	3300 ICP capacity	• Larger-capacity 3300 ICPs (700 and 1400 users) support more users than small-capacity 3300 ICPs (100 users); therefore, their failure removes a larger population from service. From a Resiliency perspective, we recommend the distribution of small-capacity devices throughout the network to minimize the number of users and devices

Considerations		
		affected by any single ICP failure.
4.	Control device distribution	- Does your network have centralized or distributed control devices? In general, centralization is less desirable from a Resiliency perspective because centralized systems are more vulnerable to single failures that affect a large proportion of a network. For example, data network failures can block all WAN access. Centralized systems are also more vulnerable to multiple failures such as power failure, fires, and floods. In contrast, widely distributed systems usually fail only in specific network areas. Note that you can reduce the higher risks associated with a centralized networking approach by incorporating UPS power, redundant L2/L3 data networking, and so on. - You can use any combination of small- and large-capacity 3300 ICPs as well as ASUs, depending on your required system capacity, desired granularity of failure protection, geographic

Considerations		
		distribution, cost, and so on.

Since a resilient system is highly scalable, many topologies are possible. The following examples identify recommended topologies for certain business models and provide an overview of network configurations, advantages, considerations, and constraints for each. For more information about recommended topologies, see ["Recommended Resilient Topologies"](#).

Scalable Resilient Topology for Locally Dispersed Networks

[Scalable, Locally Dispersed Resilient Network](#), illustrates a scalable resilient topology that is ideal for large corporations, organizations, or institutions that

- Require large deployments of phones
- Want to concentrate their network at a main office or headquarters and retain the option of including branch offices (local or remote) in the main resilient cluster or making branch offices separate resilient clusters within the network.

Advantages

- Resiliency is a function of scale, allowing you to scale only the resources you need.
- Allows ICP function concentration or dedication (see [Dedicated MiVoice Business Systems](#) and [Dedicated ICP Connectivity](#)).

Important Considerations

Configurations in which devices fail over to elements in different time zones are not supported. If your main and branch offices are in different time zones and/or different critical services areas, they should also belong to separate resilient clusters in order to ensure correct time synchronization and 911 service during a failure.




Figure 48: Scalable, Locally Dispersed Resilient Network

Scalable Resilient Topology for Campus Environment

"[Scalable Resilient Topology for Campus Environment](#)," illustrates a scalable, resilient topology that is ideal for local campus-style deployments in which ICPs function as controllers for different functional groups or departments.


Advantages

- Ideal for small to large deployments
- Resiliency is a function of scale, allowing you to scale only the resources you need.
- Fully distributed system is extremely resilient, each ICP individually survivable.

Important Considerations

- Which users need to be resilient (for example, critical services), and which do not? These decisions affect hardware requirements.

Figure 49: Scalable Resilient Topology for Campus Environment



Dedicating MiVoice Business Systems to Specific Functions

The flexibility and scalability of the Resiliency solution is enhanced by the ability to integrate or dedicate controller functionality. [Dedicated MiVoice Business Systems](#) and [Dedicated ICP Connectivity](#) illustrate the various functions to which you can dedicate your controllers.




Figure 50: Dedicated MiVoice Business Systems




Figure 51: Dedicated ICP Connectivity

6.5.3 Planning License Requirements

User licenses are provisioned on each MiVoice Business system during initial provisioning, using the License and Options form in the System Administration Tool. Additional licenses can be purchased and added to a controller after it has been programmed. Note that the required number of user licenses must be available on the appropriate ICPs before you can provision resilient users.

Each MiVoice Business system has a maximum usable license capacity. Refer to the Engineering Guidelines for details.

Required User and Device Licenses

Users, both resilient and non-resilient, require one IP User license on their home ICP..

i Note:

External Hot Desk users require an IP User license and an External Hot User license.

Resilient networks consist of a cluster of MiVoice Business systems that act as hosts for resilient devices on the network. To determine your license requirements:

1. Determine how many 3330 ICPs will be in the resilient cluster. The appropriate number of ICPs is determined by the nature of failures your system can tolerate and also by the distribution of key resources such as PSTN trunks. For example, if you want to limit exposure to certain types of outages such as those caused by fire or water damage, you may want to distribute resources among several ICPs in separate locations (recommended).
2. Determine how many user licences you will need. Estimate how many IP phone users your system will have, allowing some room for future growth.
3. Based on the number of resilient and non-resilient users you will have, determine the number of device licenses you will need.
4. Determine the maximum number of resilient seats available on a given ICP by subtracting the number of IP user licenses from the max capacity of the ICP. For example, a 700-user ICP with 400 IP-user licenses can host a maximum of 300 resilient devices from other ICPs in the cluster.
5. For T1/E1 trunk resiliency, ensure that you obtain licenses for the T1/E1 trunks for both the primary and secondary controllers. You require licenses for the T1/E1 trunks on the secondary controller even though they will only be used during a failure of the primary controller.
6. Determine which ICPs will support each resilient device. There are no restrictions regarding where resilient devices can be supported. For each resilient device, the administrator can designate any 3300 ICP or MiVoice Business system in the cluster to serve as its secondary system (call routing must be appropriately configured). For example, in a cluster of three ICPs, if an ICP with 400 resilient users fails, you can program 200 devices to Fail-over to each of the two remaining ICPs, and in this way split the Fail-over load.

Note:

1. It is recommended that devices belonging to designated workgroups be supported by the same primary and secondary ICPs. This helps to preserve the workgroup capabilities of these devices during a failure and also helps to minimize traffic on the secondary ICP.
2. It is recommended that you minimize the distance between a resilient device and its secondary ICP. ICPs and their devices do not have to be located in the same room or even in the same building, but they must be located in the same time zone and emergency services area to ensure that users retain access to local emergency services and appropriate day/night service. A device's PSTN access and time settings reflects what is programmed on the ICP it is registered on.

6.5.4 Planning Resilient Clustered Hot Desking

Clustering Requirements

In addition to existing clustering requirements for resiliency, the following conditions apply to resilient clustered Hot Desking:

- Only 3300 ICPs (Release 5.0 or later) can act as primary and secondary ICPs for resilient clustered hot desk users and devices.
- If any hot desk enabled phones in the cluster are dual mode IP phones, then all clustered ICPs hosting hot desk phones must be Release 5.0 or later.
- 3300 ICP Release 4.0 and 4.1 controllers can act as primary and secondary ICPs for nodal (non-clustered) Hot Desking devices only.

Resilient Clustered Hot Desking Elements

To implement resilient clustered Hot Desking, you must have the following elements:

- Resilient cluster
- Minimum of two 3300 ICP (3300 ICP Release 5.0 or later) or MiVoice Business systems in the cluster
- Remote Directory Number Synchronization must be supported by the cluster.
- Devices that support resiliency and Hot Desking (For a list of devices that support resiliency, see [Resilient Devices](#) on page 25, and for a list of devices that support Hot Desking, refer to the *MiVoice Business System Administration Tool Help*.)

Licensing

[Licensing Requirements for Resilient Clustered Hot Desking](#) details the licensing requirements for resilient clustered Hot Desking:

Table 10: Licensing Requirements for Resilient Clustered Hot Desking

Requires License	Number & Type of Licenses	License Allocation
Resilient hot desk user	1 IP User license	1 user license is required on the user's primary ICP.

Device Distribution

Within a resilient cluster, hot desk device distribution is flexible. Resilient clustered Hot Desking has the same distribution requirements as Resiliency. See [IP Device Distribution](#) on page 204.

Device and Call Behavior

Resilient call routing and call behavior, as described in this document, apply to resilient hot desk users and devices.

The following scenarios further characterize call routing and device behavior for resilient clustered Hot Desking:

Table 11: Device and Call Behavior

Scenario	Behavior
Out of service (OOS) handling	Out of service handling is invoked on the primary ICP of the resilient hot desk user, except in cases where the user is in service on the secondary ICP and the primary is unreachable.
Resilient hot desk user logs in to a resilient hot desk device.	Resilient hot desk users who log in to resilient hot desk devices experience full Hot Desking resiliency.

Scenario	Behavior
	• If the hot desk user's primary ICP fails or becomes unreachable • the resilient hot desk device fails over to the hot desk user's secondary ICP • the user experiences persistent login. When the hot desk user's primary ICP returns to service • the resilient hot desk device fails back to the hot desk user's primary ICP • the user experiences persistent login. • If both the primary and secondary ICP of the hot desk user are unreachable • the resilient hot desk device fails over to its primary ICP (or secondary ICP if the primary is unreachable) • the user is logged out and cannot log back in to a hot desk device until either their primary or secondary ICP returns to service.
Resilient hot desk user logs in to a non-resilient hot desk device.	On a non-resilient hot desk device, resilient hot desk users experience call resiliency but not persistent login. • If a resilient hot desk user logs in to a non-resilient device and the ICP hosting the user fails or becomes unreachable • the call in progress survives • the hot desk device fails over to its home ICP • the hot desk user is logged out and cannot log back in to any non-resilient hot desk device until the hot desk user's primary ICP returns to

Scenario	Behavior
	service. The user can, however, log in to a resilient hot desk device.
Non-resilient hot desk user logs in to a resilient device.	Non-resilient hot desk users experience call resiliency but do not experience persistent login. • If a non-resilient hot desk user logs in to a resilient hot desk device and the hot desk user's home ICP fails or becomes unreachable • a call in progress survives • the hot desk device fails over to its primary ICP (or secondary ICP if the primary is unreachable) • the user is logged out and cannot log back in to a hot desk device until the hot desk user's home ICP returns to service.
Resilient hot desk user logs in over already logged in user (without logging out the user first).	• If either the primary or secondary ICP of the new mobile DN are healthy and reachable, the new login succeeds, and the previous user is logged out. • If neither the primary nor secondary ICP of the new mobile DN are reachable, the new login fails, and the previous user is not logged out.

Implementation Considerations

- Resilient hot desk users and hot desk devices can be provisioned on different ICPs (primary and secondary) in the cluster.
- No additional user licenses are required for resilient Hot Desking; however, to experience Hot Desking resiliency
 - the hot desk user must have resiliency enabled and be provisioned on a secondary ICP.
 - the resilient hot desk user must log in to a hot desk device that has been provisioned to be resilient. That is, the device must be provisioned with resiliency and with Hot Desking enabled.

Upgrading from Nodal Non Clustered Hot Desking

All provisioning associated with nodal Hot Desking on a 3300 ICP Release 4.0 or 4.1 controller is preserved when you upgrade to 3300 ICP Release 5.0 or later. The hot desk PIN (new in Release 5.0) is set to the default (blank) for every user affected by the upgrade.

6.5.5 Planning VLANs and DHCP Options

For information on VLANs and DHCP options, see [Appendix Appendix: "Engineering Guidelines for Resiliency"](#).

6.5.6 Planning Call Routing

CAUTION:

Do not configure any cluster elements to absorb CEID routing digits. The absorption of CEID routing digits can cause calls to route improperly. It may also cause message waiting indicators on resilient phones to continue flashing for Callback messages that have already been returned.

Call routing within a cluster is similar to dialing a system speed-call, representing an ARS digit string to a destination on a remote network element. Consider the following clustering requirements and recommendations:

1. All participating ICPs must have the same node ID.
2. All participating ICPs must be listed in the cluster element table on each ICP. The table holds a unique Cluster Element ID (CEID) for each ICP, a unique string of routing digits for each ICP, and identifies which ICP is the local or "THIS" ICP. The information in the cluster element table must be the same on each ICP, except for the ICP designated as "local".
3. Cluster routing digit strings must be the same length on all cluster elements. Note that it is the responsibility of the system administrator to manually program these digit strings and ensure their lengths match.
4. Every ICP must have exact ARS leading digits defined for each string of routing digits in the cluster element table, except for the routing digits of the local ICP. The ARS digits are required to select an ARS destination that routes to the corresponding ICP.
5. The routing digits assigned to the local ICP are in the local ICP's digit tree, pointing to a local cluster element entity. When these routing digits are dialed on the corresponding local ICP, the digits are ignored, and digit collection/processing continues as though the digits had not been dialed.

6. The DNs that are to be dialable from any ICP in a cluster (referred to as Portable Directory Numbers, PDNs), must be identified on every ICP in the cluster. A different subset of these PDNs on each ICP represents local devices hosted by the ICP. On each ICP, the remaining PDNs are stored in a Remote Directory Number (RDN) table. For each number in the RDN table, the ICP must be identified. Local PDNs are not identified in the RDN table.

i Note:

For a given ICP, a local PDN is analogous to a local DN, and a remote PDN is analogous to a remote DN or RDN.

7. On a given ICP, when a local DN is dialed, the digits translate to a local resource table, and call control (CC) attempts to seize the referenced device. When an RDN number is dialed, the digits translate to a RDN table entry.
8. For an RDN destination, CC retrieves the corresponding ICP routing digits, fills a digit register with the routing digits followed by the RDN digits, and repeats digit translation. The second invocation of digit translation will result in an ARS destination that routes the call to the ICP hosting the device referenced by the RDN.
9. It is recommended, and required for DPNSS Route Optimization, that the ARS routes used for cluster routing deliver the routing digits and the RDN digits.
10. The routing digits steer the call to the appropriate host ICP, and the RDN digits represent the local device digits on the host ICP.
11. Routing calls between clusters is typically based on node ID. That is, each non-cluster ICP has a unique node ID and ARS routing based on leading digits equal to the remote ICP node IDs. Each entire cluster in such a network acts as another ICP with a unique node ID.

6.5.7 Planning ARS Routes

When a call is made to a resilient device, the originating ICP attempts all ARS routes to the primary ICP. If those attempts fail, the originating ICP attempts all ARS routes to the secondary ICP. For this reason, you must configure ARS routes properly to ensure correct and efficient operation of ARS routing at the cluster level.

To program ARS routes for resilient call routing

1. Design a fully meshed IP Trunking environment.
2. In the configuration tables of each ICP, program ARS routes to both the primary ICP and secondary ICP CEIDs of all resilient devices.

3. Provide alternate (non-IP) routes to any high-priority primary and secondary ICPs in the cluster to reduce the likelihood of total failure due to IP network failures.

Planning ARS Routes for Boundary Nodes

Use the following Guidelines when planning and programming ARS routes for a resilient cluster that includes a boundary node:

- For calls that originate from outside the resilient cluster, program ARS routes to include at least one ICP running 3300 ICP Release 4.0 or later in the destination cluster. This ICP should be the first one in the list.
- For calls that originate from any ICP in the cluster with a release prior to 4.0, program ARS routes to contain at least one ICP running 3300 ICP Release 4.0 or later in the routes to all other ICPs in the cluster. These more recent ICPs should be high in the list.

Note:

Because the SX-2000 LIGHT is capable of running LIGHTWARE 32 Release 1.1, which is compatible with 3300 ICP Release 4.0, these nodes may also participate in resilient call routing as a boundary node, where the boundary is not also a primary or secondary ICP.

About Resilient Multicall and Key Line Appearances

The following network setup causes a known routing condition that impacts resilient multicall and key line appearances:

Network Setup

TDM and IP trunking connect ICP A (primary) and ICP B (secondary). An IP phone with the extension number 5000. 5000 has a resilient Multicall or key line appearance of extension 4000. A DNI phone in service on ICP A, extension 6000, also has a Multicall or key line appearance of extension 4000.

Routing Condition

If 5000 fails over to ICP B, there is a Multicall or key line appearance of 4000 on an IP phone in service on its secondary ICP (ICP B) and on a DNI phone on ICP A. In this situation, if a call originating from ICP B is placed to 4000, the system uses the

TDM trunking to ring 6000, the DNI phone on ICP A, instead of ringing the extension appearance on 5000, which is in service on ICP B, the call-originating ICP.

6.5.8 Planning T1 E1 Trunk Resiliency

Refer to [T1/E1 Trunk Resiliency](#) on page 97 for a list of the conditions that apply to this feature.

6.5.9 Planning Voice Mail and Other Options

See below for information on planning and setting up voice mail types. For information about other system features, see [Feature Resiliency](#) on page 48, and also refer to the *MiVoice Business System Administration Tool Help*.

Note:

It is recommended that you use call routing to route calls to voice mail. Call forwarding cannot be used to route calls to voice mail when a phone is on its secondary system.

Embedded Voice Mail

The embedded voice mail in each MiVoice Business system is set up as a hunt group with up to 750 mailboxes and up to 30 ports. The number of ports you have determines the number of users who can access the embedded voice mail on an ICP at the same time.

- Each voice mail port must be configured as part of a voice mail hunt group. The pilot number of this hunt group is programmed as a (non-resilient) RDN in the Remote Directory Numbers form of all 3300 ICP Release 4.0 or later system in the resilient cluster, other than the ICP that is hosting the voice mail. The CEID of the ICP hosting the voice mail is programmed in that ICPs Remote Directory Numbers form. There is no CEID for the secondary ICP.
- If programming non-centralized embedded voice mail, two embedded voice mailboxes must be assigned for the user of a resilient device: one on the primary ICP and one on the secondary ICP. Each mailbox must be individually configured (password, greeting, name) by the user. The user must set up pastiches and greetings on both the primary and secondary ICPs.
- Call rerouting must be programmed to the local voice mail on a resilient device's primary and secondary ICP so that the user does not need to do anything to make the mailbox active in the case of Fail-over.

For users to be notified of messages waiting on either of their controllers, you must amalgamate the callback message queues on the primary and secondary controllers.

External Voice Mail

Resilient users obtain voice mail service from their centralized voice mail server, regardless of whether their device is registered on their primary or secondary ICP. In either case, the user is notified of waiting messages by the Message Waiting Indicator. Notification is attempted at both of the device's ICPs, and is received by the device through the ICP it is in service on.

You must amalgamate the callback message queues on the primary and secondary ICPs, so that once devices fail back to their primary ICPs, users can access messages recorded on the secondary ICP.

Programming Speed dial Keys to Access Voice Mail

If a user is using speed-dial keys to access voice mail, two speed-dial keys must be programmed: one for the primary ICP and one for the secondary ICP.

Routing Calls to Voice Mail

You must program voice mail pilot numbers in the Remote Directory Numbers form so that the callback message can be routed to the correct voice mail ICP when a device user accesses voice mail messages.

In the following example, 1234 is the directory number of a resilient device that has ICP 5 and ICP 9 as ICPs. The pilot number for embedded voice mail on ICP 5 is 90005 and the pilot number for embedded voice mail on ICP 9 is 90009.

90005 must be defined as an RDN on all cluster elements other than ICP 5. The cluster element ID (CEID) of the primary ICP is that of ICP 5; there is no secondary CEID. Similarly, 90009 must be defined as an RDN on all cluster elements other than ICP 9. The CEID of the primary ICP is that of ICP 9, and there is no secondary CEID.

ICP Number	Directory Number (DN)	Rerouting Pilot Number (No Answer, Busy, DND...)
5	1234	90005
9	1234	90009

With centralized voice mail, all devices in the cluster use the same voice mail pilot number. In large systems, there may be more than one centralized voice mail server, and

consequently, there may be more than one pilot number. Each centralized voice mail pilot number must be defined as an RDN or provided an ARS route on all cluster elements except the one that is hosting the voice mail.

Voice mail pilot numbers are not resilient, so only the primary CEID is provided. The pilot number(s) can be programmed as an RDN or as an ARS destination. For any given device with resilient centralized voice mail, the way you route calls to voice mail at both the primary and secondary ICPs is identical. Rerouting to voice mail can be programmed in the first or second alternate system rerouting.

6.5.10 Planning the DECT IP Wireless Solution (EMEA) Networks

In 3300 Release 6.0 or later, the IP DECT Wireless Solution (EMEA) can be registered as a resilient device.

Note:

IP DECT wireless telephones do not obtain their firmware from the TFTP server. The firmware or application software for the IP DECT phones is downloaded from a PC through a USB interface. For more information, please refer to the *IP DECT Wireless Solution (EMEA) Technical Manual*.

The IP DECT wireless components and services are not by default operating as resilient or redundant devices. The MiVoice Business System Administrator determines the required level of resiliency that is required and takes the IP DECT components and services into consideration.

For example, if you need resiliency for the Riffs you can deploy two Riffs to provide the same radio frequency coverage. For further information on how to engineer and ensure resilient operation of the above components, please refer to the *IP DECT Wireless Solution (EMEA) Site Survey Guide* and *IP DECT Wireless Solution (EMEA) Technical Manual*.

6.5.11 Considering Upgrading and Migration Issues

Upgrading an ICP to Support Resiliency

If you upgrade a non-resilient ICP to support resiliency within a cluster that is already resilient, you must ensure that the RDN table of the newly upgraded ICP is updated with the secondary CEIDs for resilient users. The ICP must have this information in order to route resilient calls.

In clusters that have been migrated to support Remote Directory Number (RDN) Synchronization, the RDN tables are updated automatically after you add the non-resilient ICP to the migrated cluster. Note that to support RDN Synchronization, all elements in the cluster must have 3300 ICP Release 4.0 or later software.

Implementing Resiliency

7

This chapter contains the following sections:

- [Implementation Overview](#)
- [Assigning User and Device Licenses](#)
- [Installing or Upgrading the MiVoice Business Systems](#)
- [Programming the Cluster and ARS](#)
- [Setting up VLANs and DHCP Options](#)
- [Implementing IP Trunking](#)
- [Configuring T1 E1 Trunk Resiliency](#)
- [Migrating the Cluster to Support RDN Synchronization](#)
- [Provisioning Resilient Devices](#)
- [Configure Resilient Hunt Groups](#)
- [Configure Resilient Multi-device User Groups](#)
- [Configure Resilient Personal Ring Groups](#)
- [Configure Resilient Ring Groups](#)
- [Programming Voice Mail](#)
- [Configuring the IP Consoles](#)
- [Programming Listed Directory Numbers](#)
- [Programming Class of Service Options](#)
- [Programming Resilient Keys](#)
- [Verifying Resiliency](#)
- [Maintaining a Resilient System](#)

7.1 Implementation Overview

After you have planned your resilient network, complete the steps in the following table to install and configure it.

Table 12: Implementing a Resilient Network

Implementation Step		Where to Find Information
1.	Install or upgrade MiVoice Business systems	See Installing or Upgrading the MiVoice Business Systems on page 230.

Implementation Step		Where to Find Information
2.	Assign user licenses (resilient and non-resilient users).	See Assigning User and Device Licenses on page 230.
3.	Define the cluster elements and program ARS to support resiliency.	See Programming the Cluster and ARS on page 230.
4.	Set up VLANs and DHCP options.	See Appendix Appendix: "Engineering Guidelines for Resiliency" and Setting up VLANs and DHCP Options on page 231.
5.	Implement IP trunking.	See Implementing IP Trunking on page 232.
6.	Configure T1/E1 trunk resiliency.	See Configuring T1 E1 Trunk Resiliency on page 232
7.	Program voice mail and other options.	See Configuring T1 E1 Trunk Resiliency on page 232.
8.	Migrate the cluster to support RDN Synchronization.	See Migrating the Cluster to Support RDN Synchronization on page 235.
9.	Provision resilient users and devices.	See Provisioning Resilient Devices on page 236.
10.	Configure resilient hunt groups, resilient personal ring groups, resilient multi-device user groups	See Configure Resilient Hunt Groups on page 237, Configure Resilient Ring Groups on page 242, Configure Resilient

Implementation Step		Where to Find Information
	and resilient ring groups (optional).	Multi-device User Groups on page 241 and Configure Resilient Personal Ring Groups on page 241.
11.	Provision resilient ACD agents (optional).	See Configuring ACD Resiliency on page 118
12.	Configure IP consoles.	See Configuring the IP Consoles on page 249.
13.	Program Listed Directory Numbers.	See Programming Listed Directory Numbers on page 250
14.	Program Class of Service Options.	See Programming Class of Service Options on page 251
15.	Program resilient keys.	See Programming Resilient Keys on page 251
16.	Verify the resilient configuration.	See Verifying Resiliency .
17.	Troubleshoot the resilient system.	See Troubleshooting a Resilient System in the 3300 ICP Troubleshooting Guide.
18.	Maintain the resilient system.	See Maintaining a Resilient System in the 3300 ICP Troubleshooting Guide.

7.2 Assigning User and Device Licenses

Assign user licenses to both resilient and non resilient users with the System Administration Tool. At each MiVoice Business system that you plan to include in the resilient network

1. Launch the System Administration Tool (refer to the *3300 ICP Technician's Handbook* for instructions).
2. In the License and Options form
 - In the **IP User Licenses** field, specify the number of purchased IP user licenses.

Note:

The License and Options form displays the number of purchased IP user and device licenses only. You can view the numbers of both purchased and used licenses, in the System Capacity Display.

7.3 Installing or Upgrading the MiVoice Business Systems

If you are implementing resiliency in a new site, install the MiVoice Business systems. Refer to the *3300 ICP Technician's Handbook* and *Hardware Technical Reference* guide for instructions.

If you are implementing resiliency in an existing site, upgrade the existing ICPs to 3300 ICP Release 4.0 software or higher. Refer to the *3300 ICP Technician's Handbook* for instructions on how to perform a software upgrade.

7.4 Programming the Cluster and ARS

A resilient network must first be configured as a standard cluster. You must perform cluster programming at each element Refer to *Voice Networking -> Configure Network* in the *System Administration Tool Help* for instructions on how to program and set up a standard cluster.

For system that are running 3300 Release 6.0 or later software, use the System Data Synchronization (SDS) feature to create the cluster definition and make all the required system data consistent across the all the cluster elements.

7.4.1 Guidelines for Programming a Resilient Cluster

The following additional programming steps must be performed on each cluster element through the elements' System Administration Tool:

1. Program identical values on all ICPs in the cluster for the following items:
 - Set Registration Access code in the System Options form
 - Set Replacement Access code in the System Options form
 - The way that COS or COR indices map to specific sets of services and restrictions.
2. Configure the Cluster according to the existing configuration requirements for IP Trunking, using the following rules:
 - The Cluster Element ID in the Cluster Elements form must have the same value (range is 1–256) as the PBX Number in the ICP/PBX Networking form.
 - Each cluster element must have all ICPs in the cluster configured in its ICP/PBX Networking form.
 - It is highly recommended that COS/COR settings and options be identical on all cluster nodes.

For additional information pertaining to cluster configuration refer to *Voice Networking -> Configure Network* in the *System Administration Tool Help*.

7.5 Setting up VLANs and DHCP Options

DHCP Options

The DHCP server in the MiVoice Business system provides options 66 (TFTP Server IP address for E2T, same as option 129 or its equivalent tag in options 43 and 125) and 67 (Name of file on external FTP server—Boot File Name) to support multiple internal DHCP servers on the same subnet. Set up these options to ensure that the E2T obtains an IP address from the correct ICP. These options are available on 250- and 700-user controllers that have an E2T card. The options do not apply to the 100-user controller because it does not have an E2T card. For more information about DHCP options, see [DHCP Servers](#) on page 314.

Administrators must still program a static IP (reservation) address, but now they can also apply options 66 and 67 to the static IP range (scope) to ensure that the E2T of any

system within the same subnet obtains the correct IP address and Boot from its local controller.

For instructions on how to program E2T information into the DHCP server in a MiVoice Business system, see [Program E2T Information for DHCP](#).

7.6 Implementing IP Trunking

Complete the required programming for IP trunks and/or XNET trunks at each element in the cluster. Refer to the Programming Procedures book of the *MiVoice Business System Administration Tool Help*.

When programming the ICP/PBX Networking form at each element:

- Assign the local element a PBX Number. This PBX Number (range is 1–256) must match the CEID index that you assigned to the element in the Cluster Element form.
- Assign a PBX Number to each remote element. The PBX Number that you enter for a remote element must match the CEID index that you assigned to it in its Cluster Element form.

7.7 Configuring T1 E1 Trunk Resiliency

To configure T1/E1 trunk resiliency

1. Review the list of conditions that apply to this feature (see [T1/E1 Trunk Resiliency](#) on page 97).
2. Ensure that 3300 Release 7.0 software or later is installed in both the primary and secondary controller.
3. Install the T1/E1 Combo MMC (PN 50005160) in the primary controller. Refer to the Technician's Handbook for instructions on how to install MMC modules.
4. Install the T1/E1 Combo MMC or T1/E1 Dual MMC in the secondary controller.
5. Connect the trunk from the PSTN to the Input port of the T1/E1 Combo MMC in the primary controller (see [Configuring T1/E1 Trunk Resiliency](#)). Use Category 5 cable with an RJ-45 connector to make the connection.
6. Connect the Failover link port in the T1/E1 Combo MMC on the primary controller to the Input port in the T1/E1 MMC on the secondary controller. Use a Category 5 (straight-through) cable with RJ-45 connectors. This cable must not exceed 10 meters (30 feet) in length.

**Warning:**

The Failover port on the T1/E1 Combo MMC does not provide secondary protection. Only connect the Failover port to a T1/E1 Combo MMC or Dual T1/E1 Combo MMC in a secondary controller.




Figure 52: Configuring T1/E1 Trunk Resiliency

Figure 53: Configuring T1/E1 Trunk Resiliency

7. Log into the System Administration Tool of the primary controller and program the T1/E1 Combo MCC module. Refer to the System Administration Tool Help for instructions.
8. Log into the System Administration Tool of the secondary controller and program the T1/E1 Dual Framer or T1/E1 Combo MCC module.
9. Program the T1/E1 digital trunks on both the primary and secondary systems. Program the trunks with the same Link Descriptor settings. Refer to the System Administration Tool Help for instructions on how to program trunks.

10. In the Digital Links form of the primary controller:

- Select the T1/E1 link. The interface type of the link must be UNIVERSAL T1 or UNIVERSAL E1.
- Click **Change**.
- Check the "Resilient Link" box.
- From the Resilient Link ID drop-down menu, select a link identifier (1 to 4) for the primary link. This link ID must match with the link ID that you assign to the secondary link. The link ID must be unique on the primary controller. To simplify the configuration, make the link ID that same as the module ID on the primary controller.
- Set the "Primary System Name" to the system name of the primary (Local) controller.
- Set the "Secondary System Name" to the name of the secondary controller that contains the standby T1/E1 MMC.

 Note:

You cannot set the "Primary System Name" and "Secondary System Name" to the same value. In addition, you must program at least one of these fields with the name of the Local system.

11. In the Digital Links form of the secondary controller:

- Select the T1/E1 link and click **Change**.
- Check the "Resilient Link" box.
- From the Resilient Link ID drop-down menu, select a link identifier (1 to 4) for the secondary link. This link ID must match with the link ID that you assigned to the primary link.
- Set the Primary System Name to the system name of the primary controller.
- Set the "Secondary System Name" to the name of the secondary (Local) controller.

Note:

If you do not complete T1/E1 Resiliency programming in the Digital Link Assignment on the secondary controller, the feature is only partially supported. Although the T1/E1 trunk fails over to the secondary controller if the primary controller goes out of service, the following deficiencies are present:

- Alarms are displayed on the secondary link when the secondary is not active.
- As soon as the primary controller recovers, it seizes control of the link back from the secondary. It will not wait until all circuits are idle.
- The Programmed Failback maintenance command is not supported

12. Program ARS. Ensure that you program the route lists on both controllers so that all users can access the resilient T1/E1 trunk when the trunk is controlled by either the primary or secondary controller. The route list that you program on the controllers must have routes that are programmed as follows:

- On the primary controller, program a route list with a first choice route that seizes a resilient T1/E1 trunk. Program the second choice route in the route list to seize an IP trunk to the secondary controller.
- On the secondary controller, program a route list with a first choice route that will seize a resilient T1/E1 trunk. Program the second choice route in the route list to seize an IP trunk to the primary controller.

13. Verify that you have configured the T1/E1 trunk resiliency correctly by forcing the trunk to fail over to the secondary controller. Use the EDT Force Failover (EDT F FO) maintenance command.

7.8 Migrating the Cluster to Support RDN Synchronization

To migrate a cluster to support Remote Directory Number Synchronization, all elements in the cluster must be upgraded to MCD Release 4.0 or higher software. If any of the elements in your cluster have pre-MCD Release 4.0 software, you must use OPS Manager.

Refer to the Voice Networking -> Manage Network -> Remote Directory Number Synchronization book in the System Administration Tool Help for instructions on how to migrate the cluster to support Remote Directory Number Synchronization.

7.9 Provisioning Resilient Devices

To program resilient devices from the System Administration Tool (RDN Synchronization must be enabled for cluster):

1. Log into a System Administration Tool session on the element that you want to be the home (primary) element of the device. When you create a resilient device, the local element that you are logged into is the device's home element.
2. Access the User and Services Configuration form.
3. Complete the fields (Personal Information, Device Type and so forth) for the new user.

Note:

- a. You can assign different Mailbox Numbers and Extension Numbers to the same user through the System Administration Tool. SDS cannot replicate data from this view unless these fields match. Ensure that the Mailbox Number and Extension Number are identical when you provision an embedded voice mailbox.
- b. When configuring an 5560 IPT for resiliency, you must assign the 5560 IPT master and slave DNs to the same primary and secondary controllers.

4. Under "Service Profile", select the name of the secondary element from the drop-down list of cluster members.
5. Click **Save**. RDN Synchronization automatically updates the other cluster elements including the secondary element with the remote directory number of the new device; System Data Synchronization automatically synchronizes the agent's user and device data and begins sharing it between the primary and secondary controllers.

7.9.1 Provisioning SIP Devices

You provision a resilient SIP device manually on both ICPs. The configuration process to make a SIP device resilient is exactly the same as the process to make a MiNET device resilient.

The "OPSless" resiliency feature allows the automatic propagation of the device programming from the primary ICP to the secondary ICP. Again, the data that must be propagated for SIP devices is the same data as that must be propagated for MiNET devices.

The primary and secondary controllers of a device as identified by its DN must be defined in the Remote Directory Numbers form.

The SIP Device Capabilities form allows variations in the handling of different SIP endpoints.

7.10 Configure Resilient Hunt Groups

To configure voice hunt group, voice mail hunt groups, or recorder hunt groups with resiliency:

1. Program device resiliency for the users that you want to include as members in the resilient hunt groups. See [Provisioning Resilient Devices](#) on page 236.
2. Configure System Data Synchronization (SDS) to support resilient Hunt Groups. You must enable the sharing of data between the primary and secondary controllers. Ensure that
 - SDS is enabled on the primary and secondary controllers and that they are configured to share data across the cluster.
 - The Class of Service (COS) form is shared at the network or cluster level.
 - The System Speed Calls form is shared at the cluster level so that any hunt group members that are identified as speed call numbers will exist in the databases of both the primary and secondary controller. However, you can also manually program the System Speed Call numbers on both the primary and secondary controller.

Refer to System Data Synchronization in the Features Reference book of the System Tool Online help for instructions on how to configure SDS. When you configure SDS to share data between the elements in a cluster, the system shares resilient user and device data between primary and secondary controllers of a resilient pair by default.

3. Enable resiliency for the hunt groups (See [Enabling Resiliency for a Hunt Group](#) on page 238 for instructions).
4. If you created a new resilient hunt group, add the hunt group pilot number in the Telephone Directory form of the primary controller. SDS automatically updates the telephone directory on the secondary controller with the hunt group pilot number.
5. Add the resilient hunt group members to the resilient hunt group on the primary controller. See [Adding a Member to a Resilient Hunt Group](#) on page 239. SDS automatically updates the secondary controller with the new members data. For this feature to be useful, all or most of the hunt group members that you add to the resilient hunt group should be configured with resiliency.
6. Program the Recorded Announce Device (RAD) greetings for the hunt group on **both** the primary and secondary controllers. Refer to the System Tool online help for instructions on how to program RAD greetings. SDS only shares the data that

is programmed in the RAD fields of the Hunt Groups form between the primary and secondary controllers.

Note:

You cannot program RADs against network hunt groups. Uniform Call Distribution (UCD) network hunt groups are not supported.

7.10.1 Enabling Resiliency for a Hunt Group

You enable resiliency for a new hunt group or for an existing hunt group from the Hunt Groups form of the primary controller:

1. Log into the primary controller.
2. In the Hunt Groups form
 - select an existing voice hunt group and click **Change**, or
 - click **Add** to create a new hunt group. Enter the Hunt Group Pilot Number, set the Hunt Group Mode (Circular is recommended), set the COS values, program the RAD, and set the Hunt Group Type to "Voice", "VoiceMail" or "Recorder".
3. To enable resiliency, select the name of the secondary controller from the drop-down menu.

Note:

If the secondary controller has pre- 3300 Release 7.0 software, or if SDS is not enabled on the secondary, or if SDS is enabled but data sharing has not been enabled between the primary and secondary, SDS will not distribute the data to the secondary controller.

4. Click **Save**. SDS automatically adds the hunt group to the secondary controller and then updates the secondary controller with the resilient group member data.

7.10.2 Adding a Member to a Resilient Hunt Group

You can add members to a resilient hunt group from either the primary or secondary controller. Members can be local members or remote members (defined by remote directory numbers or system speed call numbers).

1. Log into the System Administration Tool of either the primary or secondary controller.
2. Provision the user that you want to add to the resilient hunt group with the resiliency. See [Provisioning Resilient Devices](#) on page 236. You must program a local user with the same secondary controller as the resilient hunt group.
3. In the Hunt Groups form, select the resilient hunt group.
4. Click **Add Member**.
5. Enter the directory number of the resilient user.
6. Click **Save**. SDS distributes the "add member" update to the other resilient peer. When the resilient peer receives the update, it applies it to its database. If the update fails, SDS records an error in the SDS Distribution Errors form on the controller that initiated the update.

7.10.3 Deleting a Member from a Resilient Hunt Group

You can delete members from a resilient hunt group from either the primary or secondary controller.

1. Log into the System Administration Tool of either the primary or secondary controller.
2. In the Hunt Groups form, select the resilient hunt group.
3. Select the member that you want to delete.
4. Click **Delete Member**.
5. Click **Delete**. SDS distributes the "delete member" update to the other resilient peer. When the resilient peer receives the update, it deletes the entry from its database. If the update fails, SDS records an error in the SDS Distribution Errors form on the controller that initiated the update.

7.10.4 Disabling Hunt Group Resiliency

You disable hunt group resiliency from the primary controller by clearing the check from the **Enable Resiliency** check box in the Hunt Groups form. After you click **Save**, SDS distributes the "delete hunt group" update to the secondary controller which then deletes the hunt group, pilot number, and hunt group member list from the database of the secondary controller.

7.10.5 Deleting Resilient Hunt Groups

You can delete a resilient hunt group from either the primary or secondary controller. After you delete a resilient hunt group from the primary or secondary controller, SDS automatically deletes the hunt group from the resilient peer.

If you delete a hunt group (resilient or non-resilient) the telephone directory entry of the pilot number remains in the Telephone Directory form.

7.10.6 Changing the Secondary of a Resilient Hunt Group

You change the secondary controller of a resilient hunt group from the Hunt Groups form of the primary controller.

1. Log into the System Administration Tool of the primary controller.
2. In the User and Services Configuration form, change the existing secondary controller for each resilient member to the new secondary controller.
3. Perform a full synchronization to distribute the resilient directory numbers to all the elements in the cluster
4. In the Hunt Groups form, select the resilient group.
5. Click **Change**.
6. Select the name of the desired secondary element from the selection menu.
7. Click **Save**. SDS updates the hunt group data on the new secondary controller. SDS adds the hunt group pilot number to the Remote Directory Numbers form of the new secondary.

7.10.7 Migrating from Pre- 3300 Release 7.0 Hunt Group Configurations

On systems with pre- 3300 Release 7.0 software, you may have identical hunt groups programmed on both the primary and secondary controllers to provide a level of hunt group resiliency. After upgrading the systems to Release 7.0, these hunt groups will remain local to the controllers. To achieve full hunt group resiliency, you will still need to enable resiliency for the groups (see [Enabling Resiliency for a Hunt Group](#) on page 238).

In addition, in pre- 3300 Release 7.0 software, the COS of the first member is applied to the hunt group. If you upgrade to Release 7.0, the new COS fields in the Hunt Groups form will be blank and the COS of the first member will still be applied to the hunt group. However, if you want to add a remote member to the hunt group, you must first enter values in the COS fields in the Hunt Groups form. Typically, you would enter the existing COS values that are programmed against the first member of the group.

7.11 Configure Resilient Multi-device User Groups

The procedure for configuring resilient ring groups, multi-device user groups and hunt groups is virtually identical. To configure a resilient multi-device user group, start by following the instructions in [Configure Resilient Hunt Groups](#), omitting step 6. Continue with the instructions for enabling resiliency and adding members to the resilient group, using the Multi-device User Groups form. Ignore parts that pertain specifically to hunt groups, such as selecting the hunt group type and programming Record-A-Call resiliency.

Note:

If the secondary controller has pre-MCD 5.0 software, or if SDS is not enabled on the secondary, or if SDS is enabled but data sharing has not been enabled between the primary and secondary, SDS will not distribute the data to the secondary controller.

To delete members, disable group resiliency, delete resilient groups or change the secondary of a resilient group follow the hunt group instructions exactly as stated.

7.12 Configure Resilient Personal Ring Groups

The procedure for configuring resilient personal ring groups and hunt groups is virtually identical. To configure a resilient personal ring group, start by following the instructions in [Configure Resilient Hunt Group](#), omitting step 6. Continue with the instructions for enabling resiliency and adding members to the resilient group, using the Personal Ring Groups form. Ignore parts that pertain specifically to hunt groups, such as selecting the hunt group type and programming Record-A-Call resiliency.

Note:

If the secondary controller has pre-MCD Release 4.0 software, or if SDS is not enabled on the secondary, or if SDS is enabled but data sharing has not been enabled between the primary and secondary, SDS will not distribute the data to the secondary controller. If necessary, perform the upgrade and a start sharing operation before you perform this procedure.

To delete members, disable group resiliency, delete resilient groups or change the secondary of a resilient group follow the hunt group instructions exactly as stated.

7.13 Configure Resilient Ring Groups

The procedure for configuring resilient ring groups and hunt groups is virtually identical. To configure a resilient ring group, start by following the instructions in [Configure Resilient Hunt Groups](#), omitting step 6. Continue with the instructions for enabling resiliency and adding members to the resilient group, using the Ring Groups form. Ignore parts that pertain specifically to hunt groups, such as selecting the hunt group type and programming Record-A-Call resiliency.

Note:

If the secondary controller has pre- 3300 Release 8.0 software, or if SDS is not enabled on the secondary, or if SDS is enabled but data sharing has not been enabled between the primary and secondary, SDS will not distribute the data to the secondary controller.

To delete members, disable group resiliency, delete resilient groups or change the secondary of a resilient group follow the hunt group instructions exactly as stated.

7.14 Programming Voice Mail

The procedures for configuring resilient users with voice mail support depend on the type of voice mail system:

- **MiVoice Business Embedded Voice Mail:** You can provide resilient users with access to voice mail while they are on their secondary controller by programming a separate voice mailbox on the secondary controller.
- **Centralized voice mail:** You can program call rerouting on both the primary and secondary ICPs to direct calls to the hunt group pilot number of the centralized voice mail system.
- **NuPoint Unified Messenger (Release 10.0 and later software):** You can configure voice mail port resiliency on the NuPoint Unified Messenger server.

Note:

Resiliency is not supported for MiVoice Business embedded voice mail ports.

7.14.1 Embedded Voice Mail

MiVoice Business does not support resiliency for Embedded Voice Mail. However, resiliency can be achieved (to an extent) by manually duplicating the **Auto Attendant** and **Voice Mailboxes** programming for each resilient user on both primary and secondary controllers.

Note:

- For Embedded Voice Mail resiliency on a Container-based MiVoice Business (MiCloud Flex on Google Cloud Platform (GCP)), see *MiCloud Flex on GCP Deployment Guide*.
- For Container-based MiVoice Business, Mitel recommends MiCollab Visual Voice Mail (VVM) for managing voice mail messages rather than desk phone VVM or Telephone User Interface (TUI) as MiCollab Visual Voice Mail (VVM) offers full resiliency support. For more information, see [Desk Phone VVM/TUI and MiCollab VVM](#) on page 246.

7.14.1.1 Auto Attendant

Before you begin, ensure that you have at least one **Embedded Voice Mail** license configured in the **License and Option Selection** form on the secondary controller.

To implement Auto Attendant resiliency:

1. Program Embedded Voice Mail on the primary controller. For more information, see **System Applications > Messaging > Voice Mail (Embedded) > Programming > Programming Embedded Voice Mail** in the *System Administration Tool Help*.
 - In the **VM Ports** form, set **Local-only DN** to **yes** for all ports. This allows you to reuse the same Directory Numbers on the secondary controller.
 - Configure the Auto Attendant feature. For more information, see **System Applications > Messaging > Voice Mail (Embedded) > Features > Auto Attendant** in the *System Administration Tool Help*.
2. Repeat step 1 on the secondary controller.
 - Configure the Auto Attendant feature as programmed in the primary controller.
3. Program call rerouting on the primary controller to direct calls to the hunt group pilot number of the embedded voice mail. See [Routing Calls to Voice Mail](#).

7.14.1.2 Voice Mailboxes

Before you begin, ensure that you have at least one **Embedded Voice Mail** license configured in the **License and Option Selection** form on the secondary controller for every user requiring a resilient Voice Mailbox.

To implement Voice Mailbox resiliency:

1. Program Voice Mailboxes on the primary controller. For more information, see **System Applications > Messaging > Voice Mail (Embedded) > Programming > Mailboxes > Programming Voice mailboxes** in the *System Administration Tool Help*.
2. In the **SDS Form Sharing** form, select **VM Mailboxes**, and set **Share From With** as **Resilient Pair**.

Note:

Before sharing the **VM Mailboxes** form, ensure that the **VM Options** form on primary and secondary controllers are identical. In particular, the **Passcode Length** and **Mailbox Length** fields must match.

3. (Optional) To avoid re-recording greetings, names, and call flows on the secondary controller, copy the following from the primary controller to the secondary controller.

Item	Location on the primary controller
Auto Attendant greetings	/vmail/c/vm/prompts/<language>/ where, language is your primary and secondary languages.  Note: Pday.vox is the opening greeting and Pnite.vox is the closing greeting.
Voice Mailbox greetings	- User greetings - /vmail/d/vm/int/int<DN>.vox - User name - /vmail/d/vm/name/nam<DN>.vox where, DN is the Directory Number of the user.
Call Flows	- Call Flows - /vmail/d/vm/callflow - Call Flows uploaded recordings - /vmail/d/vm/int/cf

4. Initialize Voice Mailboxes on the primary and secondary controllers by changing your passcode, recording your name, and greeting. You can initialize Voice Mailboxes on the secondary controller by one of the following ways:

- Failover devices to the secondary controller.
- Accessing Embedded Voice Mail features on the secondary controller.

To access EMEM features on the secondary controller, such as recording greetings or listening to voice mails through the TUI of sets connected to the primary controller, do the following:

- On the secondary controller, create a new **Local-only DN** or non-resilient Voice Mail Hunt Group.
- To the new Voice Mail Hunt Group, add the same voice mail ports that are currently in the actual Voice Mail Hunt Group.
- Program a Speed dial key on the user's set with CEID digits of the secondary controller followed by the new Voice Mail Hunt Group created in step i).

7.14.1.3 Desk Phone VVM/TUI and MiCollab VVM

The following table describes the resiliency offered by desk phone VVM or TUI and MiCollab VVM.

Note:

MiCollab VVM is supported only with Containers-based MiVoice Business.

Scenario	Desk phone VVM or TUI	MiCollab VVM
Both primary and secondary controllers are active.	You can read and delete the voice mail messages on both the primary and the secondary controllers, provided you have programmed the secondary controller to access Embedded Voice Mail features (see Embedded Voice Mail).	You can read and delete all voice mail messages on both primary and secondary controllers.
Only primary controller is active.	You can read and delete only the voice mail messages on the primary controller. You can read and delete voice mail messages on the primary controller.	You can read voice mail messages that are already read (or cached in MiCollab) on the secondary controller, but you cannot delete. You cannot read or delete unread voice mail messages on the secondary controller.
Only secondary controller is active.	You can read and delete voice mail messages on the secondary controller.	You can read voice mail messages that are already read (or cached in MiCollab) on the primary controller, but you cannot delete.

Scenario	Desk phone VVM or TUI	MiCollab VVM
	You can read and delete voice mail messages on the secondary controller.	You cannot read or delete unread voice mail messages on the primary controller.
Both primary and secondary controllers and inactive.	You cannot read or delete voice mail messages.	You can read only the messages that are previously read or cached, but cannot delete.

7.14.2 Non Resilient Centralized Voice Mail

If your cluster is configured with MiVoice Business embedded centralized voice mail or with an external centralized voice mail server (resilient users obtain voice mail from same ICP or server, regardless of whether or not they are on their primary or secondary ICP), perform the following steps to program voice mail for resilient users:

1. Configure each voice mail port as a member of a voice mail hunt group.
2. From OPS Manager, program this hunt group pilot number as a non-resilient RDN and propagate the change to the Remote Directory Number Assignment forms of each 3300 ICP Release 4.0 or later system in the cluster, other than the system that is hosting the embedded voice mail. Do not assign secondary CEIDs to this hunt group at any of the ICPs.
3. Program the required voice mail boxes and settings. For MiVoice Business embedded voice mail refer to the *MiVoice Business System Administration Tool Help* for instructions.
4. Program call rerouting on both the primary and secondary ICPs to direct calls to the hunt group pilot number of the embedded centralized voice mail or external centralized voice mail system. See [“Routing Calls to Voice Mail”](#) for an example.

Note:

If embedded voice mail is configured as centralized voice mail, the voice mailboxes and the sets should be hosted on different controllers.

5. Amalgamate the callback message queues on the primary and secondary ICPs.

7.14.3 Resilient NuPoint Unified Messenger Voice Mail Ports

NuPoint (Release 10 or later) supports voice messaging port resiliency when it is configured with IP integration ports:

- Resiliency programming configured on each MiVoice Business system establishes the path to the secondary controller. No additional programming is required on the NuPoint server.
- You must program the voice mail ports as resilient. In addition, the resilient voice mail hunt group, the resilient hunt group that the ports belong to, and the resilient sets that invoke voice mail must be programmed with the same secondary controller.
- 3300 ICP Release 7.0 software or later provides resilient voice mail hunt groups. Any call to the NuPoint hunt group after a failover, routes to the same hunt group on the secondary controller.
- If the primary controller fails, calls that are connected to voice mail are maintained. However, only audio streaming and call termination events are supported. No other telephony features (transfer, initiate Record a Call) are available.
- If the primary host fails and recovers while a messaging port is engaged, the channel fails over to the secondary controller first, then fails back to the primary.

If your cluster is configured with NuPoint (Release 10 or later), perform the following steps to configure voice mail resiliency:

1. Ensure that both the primary and secondary controllers are installed with 3300 ICP Release 7.0 software or later and that the SDS feature is configured to share data.
2. If you are configuring more than 60 ports, ensure that both the primary and secondary controllers have **Extended Hunt Group** support enabled in their License and Option Selection forms.
3. Program the required voice mail boxes and settings. For MiVoice Business embedded voice mail, refer to the *MiVoice Business System Administration Tool Help* for instructions.
4. Program switch integration ports and standard IP integration parameters (System Options, Feature Access Codes, Class of Service) on the primary controller.
5. Enable resiliency for the switch integration ports either from OPS Manager, or through the System Administration Tool (if RDN Synchronization is enabled).
6. Configure a resilient voice mail hunt group and add resilient switch integration ports as members (see [Configure Resilient Hunt Groups](#) on page 237 for instructions):
 - You can create a new resilient voice mail hunt group or enable resiliency for an existing voice mail hunt group.
 - If you create a new resilient voice mail hunt group, you must manually add the pilot number to the Telephone Directory form to allow OPS Manager to add the Remote Directory Number (RDN) entry on all the cluster elements. SDS distributes the

telephone directory number to the resilient peer; OPS Manager propagates the telephone directory number to the other cluster elements.

- After you save the configuration, SDS distributes the hunt group information and its member data to the secondary controller.
7. Complete the standard programming tasks to integrate the NuPoint server with the primary MiVoice Business system:
 - Configure IP address for the MiVoice Business system
 - Define Line Groups
 - Map MiVoice Business extension numbers to lines in the Line Group
 - Apply the appropriate Line group Only Applications
 - Define the Message Waiting Indication (MWI) parameters
 - Save and activate the configuration changes.
 8. Program call rerouting on both the primary and secondary ICPs to direct calls to the hunt group pilot number of the centralized voice mail system. See Routing Calls to Voice Mail for an example.

7.15 Configuring the IP Consoles

Creating a New Resilient IP Console

1. Log into the element that will be the home (primary) element for the resilient device. The element that you are logged into when you create the new device is the device's home (primary) element.
2. Access the IP Consoles form.
3. Click **Add** to create a new device and complete the required fields in the form.
4. Under "Phone Service Settings", select the name of the secondary element from the drop-down list of cluster members.
5. Click **Save**. RDN Synchronization automatically updates the other cluster elements including the secondary element with the remote directory number of the new device; System Data Synchronization automatically synchronizes the agent's user and device data and begins sharing it between the primary and secondary controllers.

Enabling Resiliency for an Existing IP Console

1. Log into the local element for the device. The local element for the device lists the device's directory number in the User and Services Configuration form or IP Consoles form. After you enable the resilient device, the local element becomes the device's home (primary) element.
2. Access the IP Consoles form.

3. Select the directory number of an existing, non-resilient IP Console and click **Change**.
4. In the "Secondary Element" field, select the name of the secondary element from the drop-down list of cluster members.
5. Click **Save**. System Data Synchronization begins sharing the device data at the Resilient Pair scope between the device's primary and secondary element.

Configure a Resilient Non prime line Directory Number

Resilient non-prime-line directory numbers can only be configured on 3300 ICP Release 4.0 or later systems. To configure a resilient, non-prime line

On the Primary ICP

1. Program the non-prime line DN in either the Multiline Set Keys form, or the Console Softkeys form.
2. In the Remote Directory Numbers form, enter the non-prime DN with its primary and secondary CEIDs.

On the Secondary ICP

3. Repeat steps 1 and 2 to program identical information on the secondary ICP.

Telephone Directory Form

In Customer Data Entry (CDE) on the System Administration Tool on the MiVoice Business system, the Telephone Directory form indicates a non-prime DN as network type, **Int** (Internal). In the corresponding view, however, this changes to uniquely associate a new network type, **Key App** (Key Appearance), to a non-prime DN.

Pending Changes Commands

On both the primary and secondary ICPs, the Pending Changes command displays resilient non-prime-line data as having an entry type **Key App** instead of **Local**.

7.16 Programming Listed Directory Numbers

Listed Directory Numbers (LDNs) specify how an operator answers the phone. For example, the LDN, or key, on which an incoming call rings at Mitel, determines whether the call is answered "Mitel" or "switchboard".

The term "non-prime line" or "non-prime DN" denotes either a key line appearance on a multiline set, or for the IP Console, a programmed call answering softkey other than the "recall" or the prime directory number softkey.

Non-Prime DNs are allowed in the Remote Directory Numbers form if they have been programmed in either the Console Softkeys form or in the Multiline Set Keys form.

7.17 Programming Class of Service Options

It is strongly recommended that you maintain consistent COS options and settings throughout the cluster.

- Using the System Administration Tool in the primary and secondary ICPs, program identical COS options on the primary and secondary.

7.18 Programming Resilient Keys

Programming Clustered Resilient Busy Lamp Field Keys

The following sections provide examples on how to program clustered resilient BLF keys.

Cluster with Two Controllers

[Cluster with Two Controllers](#) shows an example of a cluster that consists of two 3300 ICP controllers.




Figure 54: Cluster with Two Controllers

To program resilient clustered busy lamps for a cluster with two controllers:

1. Log into the System Administration Tool on Controller 2.
2. In the Multiline Set Keys form, select extension 3000 in the upper frame of the form.
3. In the lower frame of the form, select a button that is "Not Assigned" and then click **Change Member**.
4. Set the Line Type to "DSS/Busy Lamp",

5. Set the Button Directory Number to 4000. This is the extension number of the monitored set.
6. Select a Ring Parameter and click **Save**. The Remote Busy Lamps form of Controller 2 is automatically updated to send lamp updates to Controller 1 when Controller 2 is hosting extension 4000.
7. In the Remote Busy Lamp form of Controller 2, select extension 4000, Directory Number 3000 will be programmed as a remote lamp and displayed as a local set hosting a busy lamp.
8. Repeat the above programming steps on Controller 1.
9. In the Remote Busy Lamp form of Controller 1, select extension 4000, Directory Number 3000 will be programmed as a remote lamp and displayed as a local set hosting a busy lamp.

Cluster with Three Controllers

[Cluster with three Controllers](#) shows an example of a cluster that consists of three 3300 ICP controllers.




Figure 55: Cluster with three Controllers

To program resilient clustered busy lamps for a cluster with three controllers:

1. Log into the System Administration Tool on Controller 2.
2. In the Multiline Set Keys form, select extension 3000 in the upper frame of the form.
3. In the lower frame of the form, select a button that is "Not Assigned" and then click **Change Member**.
4. Set the Line Type to "DSS/Busy Lamp",
5. Set the Button Directory Number to 4000. This is the extension number of the monitored set.
6. Select a Ring Parameter and click **Save**.
7. Repeat the above programming steps on Controller 3.

8. The Remote Busy Lamps form of Controller 3 is automatically programmed to send lamp updates from Controller 3 to Controller 2 when Controller 3 is hosting extension 4000.
9. In the Remote Busy Lamps form of Controller 3, select extension 4000. Directory Number 3000 will be programmed as a remote lamp. This form shows the host busy lamps of the local host for a particular monitored device; so, extension 3000 will appear in this section also.
10. Log into the System Administration Tool on Controller 1.
11. In the Remote Busy Lamps form, select extension 4000.
12. In the Remote Busy Lamps form, select an available "Remote Lamp" entry.
13. Add RDN 3000 as a "Remote Host Set Directory Number" and then click Save.
Controller 1 is configured to send remote busy lamp updates to Controller 2 and automatically configured to send remote busy lamp updates to Controller 3 -- the backup controller for extension 3000.

Cluster with Four Controllers

[Cluster with Four Controllers](#) shows an example of a cluster that consists of four 3300 ICP controllers.




Figure 56: Cluster with Four Controller

To program resilient clustered busy lamps for a cluster with four 3300 ICP controllers:

1. Log into the System Administration Tool on Controller 3.
2. In the Multiline Set Keys form, select extension 3000 in the upper frame of the form.
3. In the lower frame of the form, select a button that is "Not Assigned" and then click **Change Member**.
4. Set the Line Type to "DSS/Busy Lamp",
5. Set the Button Directory Number to 4000.

6. Select a Ring Parameter and click Save.
7. Log into the System Administration Tool on Controller 4 and repeat the above procedure.
8. Log into the System Administration Tool on Controller 1. In the Remote Busy Lamps form, select extension 4000 and then select an available "Remote Lamp" entry. Add RDN 3000 as a "Remote Host Set Directory Number" and then click **Save**.
9. Log into the System Administration Tool on Controller 2. In the Remote Busy Lamps form, select extension 4000 and then select a available "Remote Lamp" entry. Add RDN 3000 as a "Remote Host Set Directory Number" and then click Save.
10. Controller 1 and Controller 2 are now programmed to send remote busy lamp updates to extension 3000 on Controllers 3 and 4, whenever they are hosting extension 4000.

7.19 Verifying Resiliency

After you have configured the system, you should test the resilient IP phones:

1. Disconnect the resilient IP Phone's primary controller from the network while the phone is on hook.
2. Ensure that you can place calls from the resilient IP Phone (verifies Fail-over).
3. Ensure that you can make calls to the resilient IP Phone from another set on the secondary controller (verifies call routing).
4. Ensure that you can make calls to the resilient IP Phone from each of the other controllers in the cluster (verifies call routing in the cluster).
5. Plug the primary controller back into the network. After the primary controller boots, ensure that you can make calls (verifies that IP phone has failed back to primary).
6. While on a call at the resilient IP Phone disconnect the IP Phone's primary controller from the network. Listen for two beeps while you are in the call. If resiliency has been set up correctly, your call will be maintained. However, you will not have access to any other features (verifies call survival).
7. Hang up and then go off-hook. After you go off-hook your set should have rehomed to its secondary controller and you should have access to all the system phone features.
8. Reconnect the primary controller to the network; the IP phone fail backs to its primary.

Perform the above procedure on the IP console to verify its resiliency programming. If the any of the above test fail, refer to "Troubleshooting a Resilient System" in the 3300 ICP Troubleshooting Guide.

 Note:

The 5320, 5330, 5340, 5360, 6920, 6930, and 6940 IP Phones support applications on the set display. The first time these phones fail over to the secondary controller the applications take longer to load than they do on subsequent failovers. Therefore, if your resilient system includes these set models, it is recommended that you perform a managed failover and failback after the initial installation using the HST Courtesy Failover maintenance command. Then, if a failover does occur, the applications will be loaded faster.

7.20 Maintaining a Resilient System

Replacing IP Devices

 CAUTION:

Do not attempt IP device replacement while a device's primary ICP is not in service or while the device is homed to its secondary ICP. Doing so can introduce temporary configuration inconsistencies into the system. These inconsistencies are corrected during the next SDS data synchronization. Replace IP devices only while they are in service on their primary ICP.

Replace an IP Phone

1. Disconnect the old phone.
2. Erase a stored PIN that may be programmed in the replacement phone:
 - Disconnect power from the phone.
 - Hold down the star (*) key while you power-on the phone.
 - Disconnect the phone.
3. Replace the device and program it with the proper PIN:
 - Connect the new device.
 - To register the device, enter the Set Registration Access feature code followed by the extension number of the phone, and then press either **Superkey** or **OK**.

7.20.1 Upgrading MiVoice Business System

For information on how to perform an MiVoice Business software upgrade, refer to the *3300 ICP Technician's Handbook*.

Upgrading a Resilient MiVoice Business System

During an upgrade of a resilient MiVoice Business system, its resilient devices fail over to their secondary system. Once the upgrade is complete and the system has rebooted, the health check determines that the system is healthy. At this point, resilient devices will fail back to their upgraded primary system unless you force them to remain registered on their secondary system.

In the Controller Registry form in the System Administration Tool, the **Resiliency Allow Return to Primary ICP** default setting is **Yes**. If you do not change this setting, the resilient devices and services will automatically fail back to their primary ICP once the upgrade is complete.

To prevent the devices and services from failing back to their upgraded primary ICP, in the Controller Registry form, set the **Resiliency Allow Return to Primary ICP** to **No**.

Note:

Remember to reset this option to Yes to allow devices to fail back to their primary ICP once your operations are complete.

Upgrading a 3300 ICP to Support Resiliency

For information on how to upgrade a non-resilient ICP to support resiliency in an already resilient cluster, see [Considering Upgrading and Migration Issues](#) on page 225.

 CAUTION:

At the end of the software upgrade procedure, you are instructed to issue the Load IP Device command to ensure that the IP devices are upgraded with the latest device load provided in the 3300 ICP software upgrade. Do not issue this command until all resilient devices are back in service on their newly upgraded primary ICP; otherwise, some devices may not be upgraded. To determine the state and location of IP devices, you can use IP diagnostics (see the *3300 ICP Troubleshooting Guide*) or the IP Phone Analyzer (see the IP Phone Analyzer Help).

This chapter contains the following sections:

- [About this Appendix](#)
- [General Engineering Guidelines](#)
- [Guidelines for Resilient Network Design](#)
- [Resiliency, Functional Description](#)
- [Link Failure Detection and Management](#)
- [Connecting and Configuring IP Network Devices](#)
- [3300 ICP](#)
- [Software Version Control](#)

8.1 About this Appendix

This appendix provides engineering guidelines for resilient networks. These guidelines are intended to supplement the *3300 ICP Engineering Guidelines*. For current 3300 ICP LAN/WAN guidelines and for general 3300 ICP engineering guidelines, refer to the *3300 ICP Engineering Guidelines* on Mitel Online.

This appendix provides general resilient engineering guidelines only. For complete engineering information, refer to the detailed version of the Resiliency Guide.

8.2 General Engineering Guidelines

Considerations for a Resilient Network

- Each IP device must have a secondary ICP to fail over to, in case of an ICP or network failure.
- Ensure that startup services such as DHCP and TFTP are available from primary and secondary servers. You must consider the location of the primary and secondary DHCP/TFTP servers in the network because different topologies can influence network behavior.
- Consider reasons why the system could be taken out of service (maintenance, upgrades, and so on).
- Provide multiple access points and nodes to the same PTT or PSTN provider.
- Provide multiple access points and nodes to different PTT or PSTN providers.

Note:

ICPs (RTCs) do not accept multiple gateway addresses. They rely on ICMP redirects from routers to provide information on valid gateway addresses.

8.2.1 Performance

Set Registration

The following table provides approximate initial set registration times for the sets in 3300 Release 8.0 UR3.

Set Type	# of Sets	Approximate Time for Sets to Register
5224	1400 (MXe) 3000 (MXe Server, MiVoice Business ISS, MiVoice Business Multi-instance, MiVoice Business Virtual)	• 145 seconds until all sets are able to place calls • 180 seconds until all sets are able to place calls
5320, 5330, 5340, 5360, 6920, 6930, 6940	1400 (MXe) 3000 (MXe Server, MiVoice Business ISS, MiVoice Business Multi-instance, MiVoice Business Virtual)	• 115 seconds until phone is ready to place call • 12 minutes until applications and key programming are available on all sets (see Note under Failover/Failback times) • 190 seconds until phone is ready to place calls • 21 minutes until applications and key programming are available on all sets (see

Set Type	# of Sets	Approximate Time for Sets to Register
		Note under Failover/ Failback times)
5560 IPT	32 (MXe) 3000 (MXe Server, MiVoice Business ISS, MiVoice Business Multi-instance, MiVoice Business Virtual)	• 25 seconds until 5560 IPT is ready to place calls • 5 minutes until applications and key programming are available • 14 seconds until 5560 IPT is ready to place calls • 3 minutes until applications and key programming are available on all sets (see Note under Failover/ Failback times)

Failover/Failback

The failover and failback times in 3300 Release 8.0 and later are significantly reduced compared to previous software releases. The following table provides approximate failover times for the sets in Release 8.0.

Set Type	Number of Sets	Approximate Time
5224	1400 (MXe)	• 145 seconds until all sets are able to place calls
	3000 (MXe Server, MiVoice Business ISS, MiVoice Business Multi-instance, MiVoice Business Virtual)	• 180 seconds until all sets are able to place calls
5320, 5330, 5340, 5360	1400 (MXe)	• 120 seconds until phone is ready to place calls

Set Type	Number of Sets	Approximate Time
		12 minutes until applications and key programming are available on all sets (see Note under Failover/Failback times)
	3000 (MXe Server, MiVoice Business ISS , MiVoice Business Multi-instance, MiVoice Business Virtual)	190 seconds until phone is ready to place calls 21 minutes until applications and key programming are available on all sets (see Note under Failover/Failback times)
5560 IPT	32 (MXe)	10 seconds until 5560 IPT is ready to place calls 10 seconds until applications and key programming are available.
	3000 (MXe Server, MiVoice Business ISS, MiVoice Business Multi-instance, MiVoice Business Virtual)	14 seconds until 5560 IPT is ready to place calls 30 seconds until applications and key programming are available.

Note:

The 5320, 5330, 5340, 5360 and the 5560 IPT Phones support user applications that are downloaded from the host controller. The first time these phones fail over to the secondary controller the applications take longer to load than they do on subsequent failovers. Therefore, if your resilient system includes these set models, it is recommended that you perform a managed failover and failback after the initial installation using the HST Courtesy Failover maintenance command. Then, if a failover does occur, the applications will be loaded faster.

Managed Handoff

3300 Release 8.0 and later supports managed handoff for resilient IP phones. Managed handoff provides the System Administrator with the ability to move resilient IP phones from the primary controller to the secondary controller in a controlled fashion. This allows for controller software upgrades or reboots to take place without disrupting IP phone availability, for details refer to the System Administration Online Help.

8.3 Guidelines for Resilient Network Design

Ideal Resilient Network Design

This section offers an example of an ideal resilient network and provides the network designer with some general resilient network design guidelines and highlights areas where overall network resiliency can be enhanced.

The diagram below is an example of an ideal resilient network. This is not the only possible configuration, but it includes a number of key elements. The level of resiliency and the network configuration can change, depending upon local requirements such as number of connected devices and their location within the overall network.

Figure 57: Ideal Resilient Network Topology

AddChangeChange PageChange AllDelete

PreviousPage 1 of 1NextGo to: value:

DHCP Static IP			
Name	IP Address	Subnet	Client ID
7100MAP_NPI_1	10.35.12.19	P59WAN Data (10.35.12.0)	11020068
Scott's PC	10.35.29.19	P59 Data (10.35.29.0)	000874d446e5
E2t	10.35.30.12	P59 Voice (10.35.30.0)	08000f119f47

DHCP Static IP

Name:7100MAP_NPI_1

Subnet:P59WAN Data (10.35.12.0)

IP Address:10.35.12.19

Protocol:BOOTP or DHCP

Hardware Address

Type:MAC Address

Other - Type:

Address:00:02:c6:00:07:fa

Other - Address Length:

Client ID:11020068

IP0755

As shown in the above diagram networks can be divided into three distinct groups, the Core Network, the Distribution Network and the Access Network.

In general when designing a resilient network the end goal should be to design resiliency into the core network and the distribution network, the access network is usually non-resilient. The reasoning behind this approach has to do with balancing costs and complexity with the benefits (network reliability) that the customer requires or expects.

In a non-resilient network

- A failure at an end point will deny one user network access.
- A device failure in the access network will deny network access to all the users connected to the failed network device, this will generally be a small number of users, possibly a work group/department.
- A device failure in the distribution network will deny network access to more than one work group/department.
- A device failure at the network core will have major impact in that a large number of users will lose network access.

In a resilient network such as is depicted in the above diagram users will only lose network access as a result of failures occurring at the end point or in the access network and the number of users affected by such a failure will be relatively small.

The design approaches that have provided this level of resiliency are

- A Structured Network Layout has been applied to both the core network and the distribution network. As a result, devices residing in the core and distribution networks are interconnected with primary and back up connections.

- Backup connections are provided between devices in the distribution network and the access switches in the access network. Backup connections are not provided to the actual end-devices. The number of end-devices that are connected to an access switch will determine the number of users that will go out of service should one of the access switches fail.
- There are multiple routers between different core networks.

8.3.1 Distribution of Servers Across the Network

The distribution of servers throughout the network must be taken into consideration when designing the network; otherwise, the potential for the following problems arises:

- If a server is located remotely from its clients, excessive WAN bandwidth might be required to allow communication between clients and servers. Locating servers locally (or closer) to their clients can alleviate this issue.
- If a server is located remotely from its clients, excessive network latency might be introduced between the client and its server. Locating servers locally (or closer) to their clients can alleviate this issue.

The network designer must consider the following elements of network design to ensure adequate distribution of servers across the network:

- To support basic IP telephony, the key elements for the network designer to consider are distribution of the ICPs and the DHCP and TFTP servers:
 - Adequate LAN/WAN bandwidth must exist between the primary and secondary ICP to support the traffic that will be generated as a result of a Fail-over. Locating backup ICPs, DHCP and TFTP servers on the same subnet as the primary devices will alleviate demands on WAN bandwidth if a device fails.
 - The primary and secondary ICP for a work group should be located in the same time zone.
 - To ensure correct operation of 911 emergency services, the primary and secondary ICP for a work group should be located in the same emergency services region.
- If a Layer 3 device exists between IP devices, the network designer must ensure that the Layer 3 device supports:
 - The Internet Control Message Protocol (ICMP), specifically ICMP Echo and Echo Reply.
 - Forwarding of Dynamic Host Configuration Protocol messages (DHCP-Relay).

Note:

In general, Layer 3 devices meet the above requirements, however, if a particular Layer 3 device does not meet these requirements, the network designer must ensure that phones and their associated servers exist within the same subnet.

- When deploying application servers the network designer must consider if adequate LAN/WAN bandwidth exists between the work group and the application server.

8.3.2 Distribution of Devices Across the Network

IP Devices

Within a resilient cluster, all resilient IP devices are distributed according to the following rules:

- All IP phones must be provisioned with one primary ICP.
- You must provision each resilient IP device with one secondary ICP for backup call control.
- You can map devices to primary or secondary ICPs in any way you choose, within a resilient cluster; however, your options are limited by licensing requirements.
- You can have pre-Release 4.0 ICPs in a resilient cluster, but these ICPs do not offer resilient functionality; however, they can function as boundary nodes and transit nodes.
- An IP device must support resiliency to be programmed as a resilient device (see [Device Resiliency](#) on page 24).

IP Console

The MiVoice Business Console cannot be configured to be resilient across mu-law / a-law boundaries.

8.3.3 IP Trunking

A well designed IP network will have been designed to be fully meshed, meaning that there is always more than one IP path between IP network devices. A fully meshed IP network is inherently resilient, higher layer software is able to take advantage of this inherent resiliency.

When ICPs are introduced into a fully meshed IP network it follows that IP trunks within this network will be resilient at the physical level, it should be noted that additional trunk resiliency can be provided by TDM trunks.

If ICPs are deployed in an IP network that is not fully meshed and a TDM trunk failure is encountered, call resiliency between nodes will be lost due to the fact that there is now only one IP path between nodes.

To provide the highest level of trunk resiliency the network designer should ensure that:

1. The IP network is fully meshed.
2. That alternate routes are programmed and available via TDM trunks.

8.3.4 TDM Trunking

In non-IP networks it is standard practice to use alternate TDM trunk routes to provide trunk resiliency. This practice remains valid in IP/TDM based networks and the level of network resiliency will increase if these TDM paths are employed as backup paths for IP trunks.

In general the network designer should:

- Make use of TDM trunks as backup paths for IP trunks.

8.3.5 Network Power Provisioning

If a resilient network has a single power source (not recommended), and if that source fails, all devices stop working.

Ideally the critical elements of the network should retain power in the event of failure of main power feed. This can be through an alternative power source such as a secondary main supply, local UPS or local generator.

Primary network devices (ICPs, L2 switches, TFTP servers and DHCP servers) should be powered from a different branch circuit than the branch circuit that is used for powering the secondary network devices. This will require special attention if the primary and secondary network devices are co-located in the same wiring closet.

Critical phones should also be provided with backup power. If the phones are powered locally, then this will be needed at each phone. It might be more prudent therefore to provide power feed from the access switch, or wiring closet, where a common UPS, or backup, can maintain power. This may require the addition of in-line power feeders after the access switches.

IP consoles should be considered as critical phones, and at least one should always be provided with power backup. This is mainly to provide a callback point for 911 and emergency services.

Note:

Certain devices and applications cannot be powered from the data lines. Unless local power is provided, these devices will fail. When allocating such devices, ensure that they will not be required for use in a mission critical situation, or where 911 backup is required.

The *MIVoice Business Engineering Guidelines* provide information on powering phones locally and remotely. The Engineering Guidelines discuss proprietary Power over Ethernet and IEEE 802.3af compliant Power over Ethernet. For planning purposes individual phone power requirements are also provided in the Engineering Guidelines. See the section in the Engineering Guidelines called Power.

8.3.6 Providing Adequate Network Bandwidth

A number of aspects need to be considered to ensure adequate bandwidth between different locations within a network. There is the signaling bandwidth and the voice streaming bandwidth. In a resilient system it is also important to consider the additional bandwidth that might be needed when a particular group controller fails and the phones rehome to new controllers. As well as spreading the overall load between a number of controllers, spreading the location of these controllers around the network will also reduce the sudden increase in bandwidth needed between any two points.

Another consideration, especially with WAN links, such as Frame Relay, is how much bandwidth to allocate on a committed (more expensive) rate and how much to allow to overflow to a best effort. For example, a headquarters location may have a number of local controllers supporting a number of remote sites. It might be decided at this location that the CIR will accommodate the maximum sized system in the event that *one* of the remote site fails. The remote sites however may only allocate enough CIR bandwidth for normal operation, and following a failure will accept some calls may fail. Also to consider is the CIR bandwidth and the actual pipe bandwidth. The pipe needs to be larger than the CIR if advantages in best effort are to be realized.

Phone Signaling Bandwidth

In general, an extra 10% margin will allow for adequate bandwidth.

IP Trunk Signaling Bandwidth

The minimum bandwidth for IP trunking is a little different. If there are no calls in progress then no bandwidth is needed. However as calls are created and removed bandwidth is needed for these. Also, then messages are more frequent and more complex. Typically this is of the order of 5 times larger, than for normal phone traffic.

With an IP-Trunk busy 100% of the time, the number of calls processed will increase. Thus the 100bits/s will increase to about 600bits/s, based on previous calculations of a phones being busy 16% of the time. With the additional IP-Trunk overhead, this will increase to 3kbits/s per IP trunk.

Typically, then, the signaling overhead as a percentage is:

Table 13: IP Trunk Signaling Overhead Bandwidth Requirements

CODEC	IP Trunk Signaling Overhead %	How Derived
G.711	3.0%	3k / 100k
G.729a	7.5%	3k / 40k

In general the 5% to 10% margin normally considered allows for these conditions.

Voice Streaming Bandwidth

Generally, an extra 10% overhead for signalling will provide enough voice streaming bandwidth for a resilient network

TFTP Downloads

In order to get optimal performance from the TFTP server it is recommended that at least 7kbits/s per phone is allocated within the network. Therefore, in a network with 700 phones, and one TFTP server, a minimum speed of 5Mbits/s to the server is recommended.

DHCP Access

The bandwidth needed to handle DHCP to a phone is fairly insignificant. In a LAN network the number can be regarded as 'noise'.

Typically, for a server handling 700 phones, a minimum bandwidth of 7kbits/s is recommended.

Examples

The following are some examples of how to calculate signaling bandwidth for different connections.

Typical Business Signaling Bandwidth Per Phone

This assumes a standard phone, running at 6 calls per hour with a typical hold time of 100 seconds. The calculations are based on a full hour. Typically 3 calls will be outgoing, and 3 incoming. In the worst case scenario, these would all be trunk calls.

Table 14: Typical Business Signaling Bandwidth Requirements Per Phone

Call / Function	Type of Signaling	Required from Phone (Bytes)	Required from Controller (Bytes)
Incoming Call	Off-Hook action	1410 x 3	2420 x 3
	Keep Alive during call	310 x 10 (3 calls at 100 seconds, keep alive at 30 seconds = 10 messages)	310 x 10
Outgoing Call	Off-Hook action	1410 x 3	2420 x 3
	Dial Digits (15)	230 x 15 x 3	60 x 15 x 3
	Keep Alive during call	310 x 10	310 x 10
Display Updates	Time and Date	80 x 60 (60 minutes in 1 hour)	230 x 60
Transfers	Transfer Key	230 x 3	60 x 3
	Dial Digits (6)	230 x 6 x 3	60 x 6 x 3

Call / Function	Type of Signaling	Required from Phone (Bytes)	Required from Controller (Bytes)
Total	34640 Bytes	38480 Bytes	

As seen in [Typical Business Signaling Bandwidth Requirements Per Phone](#), there are 38480 Bytes per hour per phone, which is 308kbits/hour, which in turn is equivalent to 85bits/s. In this case, 100bits/s per phone is recommended for signaling.

Minimum Bandwidth Following Switch Between ICPs

When a number of phones rehome from one controller to another it is necessary to understand the signaling requirements for those devices. In reality, this is also the minimum bandwidth needed.

Suppose that controller A has 100 IP phones and all of these will rehome to controller B in the event that controller A fails. The signaling bandwidth can simply be calculated from the example in section 3.3.10.6.1 above, assuming that the traffic is the same (6 calls per hour per phone).

Thus with 6CCS (6CPH with 100 seconds hold time) the 100 phones would require a minimum bandwidth of 100×100 bits/s. That would be 10kbits/s.

If the traffic were to increase, for example if these were ACD phones, then the amount of signaling would increase by a proportional amount, as a rough guide. ACD phones are typically rated at 27CCS (0.75e), whereas the business phone is rated at 6CCS. Thus for 100 ACD phones, the minimum bandwidth needed would be $10k \times 27/6 = 45k$ bits/s.

When compared against the required voice streaming requirements, these numbers are small in comparison.

For the office scenario (6CCS) there will be trunk traffic but little internal traffic to controller B. Typically trunk traffic is 2/3 of all traffic. Thus 100 phones will generate 400CCS of traffic ($100 \times 6 \times 2/3$). This would require 19 channels with trunk blocking at P.01 (1 in 100).

For ACD virtually all traffic is destined for trunks. With 100 ACD users, this will require 91 channels at 27CCS.

Table 15: Signaling Versus Voice Streaming Bandwidth Requirements

CODEC	Voice Streaming Bandwidth	Signaling Bandwidth	Signaling (%)
G.711 Office	1.9Mbps/s (19 x 100k)	10kb/s	0.5%
G.729a Office	760kb/s (19 x 40k)	10kb/s	1.3%
G.711 ACD	9.1Mbps/s (91 x 100k)	45kb/s	0.5%
G.711 ACD	3.64Mbps/s (91 x 40k)	45kb/s	1.3%

8.4 Resiliency, Functional Description

Resiliency is a means of providing continued service to an end device should there be a failure within the system. Typically only one failure is considered, although failure due to a power loss has the potential to remove a number of devices in one go.

A resilient phone will be able to continue operation by virtue of the fact that it can rehome, or re-register, with a secondary ICP, rather than the designated primary ICP. Through use of IP this is a fairly 'simple' operation, simply changing the IP address. However, there are a number of other devices in the network that provide services to the phones to ensure that they continue to work, or register with the appropriate controller and also initialize to the correct state. These involve a number of other IP services within the network as well as sound network architectures.

There are a number of network tasks that need to occur in order for resiliency to work. These require both the phones and the controllers to be involved with a number of network devices in order to start up correctly and also to continue to work in the event of a failure. The main tasks are

- Acquiring and maintaining an IP address through DHCP
- Obtaining an application load to run the end devices through TFTP
- Registering with an appropriate controller for call control

The behavior of the phones and controllers during a resiliency handover, are described in the sections below.

To maintain operation, you must consider the following issues:

- How an end device gets an IP address
- Once an IP address is acquired, how this is maintained through lease time-outs
- How a new phone obtains its service, application, or code
- How to locate a working controller

Network Considerations

In planning the network, the paths needed for the various end devices to reach the appropriate servers needs to be considered. For instance how does the phone reach a DHCP server? Also to be considered is the effect of a single failure within this network and the effect of this occurring. For example, would it be acceptable to lose 50% of the devices for a period of time? Should that split be 50% in office A and 50% in office B or 0% in office A and 100% in office B?

The location of various servers should also be considered. In a larger network it may be more advisable to split the various functions into a number of independent blocks. In this way a single failure will result in a small part of the network failing. The ICP can provide a number of services, i.e. DHCP, TFTP and Call Control in one box. However it should be considered that if the power were to fail to this unit, that although it is only one box, there are in effect three services lost. For a small branch environment this may not be an issue, but it may have a larger effect on a larger business.

The bandwidth requirements through various parts of the network should be considered in light of a number of devices failing over to a backup server. It may be prudent to share a number of devices from one controller with several other controllers for call control. However TFTP and DHCP may be independent from these paths, so each must be considered.

For optimum working the loading on the ICP controller should also be considered. Ideally the system should be provisioned to handle the number of devices that would expect to be present following a handover. Or, it may be acceptable to accept the number of phones with reduced functionality, such as increased blocking to the trunks.

In a large installation it is also recommended that the IP Phones of one controller be allocated to a number of secondary controllers. In this way the additional load added to the overall system can be distributed. In this way there is minimal incremental change needed on the secondary units. It is possible, if desired, to still allocate one controller entirely as backup, but then this unit would provide little functionality under normal operation.




Figure 58: Network Server Distribution

Network Server Distribution is an example of an ideal network with a couple of controllers and appropriate network servers. TFTP and DHCP are available internally on an ICP but can also be supplied by independent servers.

The picture also shows minimal number of IP sets. Controller 'L' (Local/Left) is shown with only one IP subnet. With increased number of sets, the controller 'L' may be dealing with a number of sub-nets. As well, there may only be a single DHCP server and TFTP server covering all of these sub-nets. There are not shown for simplicity.

Configuration guidelines are highlighted in the sections below.

Obtaining IP Addresses

The phones can obtain IP address by two main methods:

- Manual programming of each set with a static IP address
- DHCP

In the manual programming of an IP address there are few issues to consider with respect to obtaining and maintaining an IP address. Since this is pre-allocated and 'permanent' then these addresses are already defined and unlikely to be lost. It does add administration burden though when phones are moved to different areas of the business, or phones upgraded or removed.

Use of DHCP is the normal method of controlling and distributing IP addresses. The main IP end points in the system are the RTC and E2T within the ICP, and all of the phone devices.

The RTC address is predefined and programmed as a static-address through the maintenance port. This address should not appear within a DHCP Scope.

The E2T address also needs to be predefined, but this can be obtained through DHCP with other information. The E2T should have its IP address predefined and marked as permanent, or very long duration. The address should be filtered using the E2T MAC address

The phones need to be able to gain access to the general pool of IP addresses. Also the phones may have different lease times depending upon the business requirements. A permanent address could be used, but then if phones move between subnets, the entry will need to be manually removed from the DHCP server so that another device can use this address. It also assumes that the administrator knows which randomly selected IP address given to the phone is to be removed. Where phones move often, a shorter lease time is recommended.

The phones operate in the following manner:

- Broadcast to locate DHCP server for initial IP address
- Unicast to DHCP server that provided IP address at first timer T1 expiry
- Broadcast to renew IP address at expiry of timer T2, and if unsuccessful, the phone will reset from initial startup

If a DHCP server fails, then a backup is needed. This must continue to hand out IP addresses within the appropriate subnet, as well as being able to renew the lease on an address. With multiple DHCP servers there is also the potential conflict of handing out the same IP addresses from two different servers. The servers do a check for this through ICMP Echo (Ping) and the phones will check, once they have been given an IP address, through ARP. Thus conflicts should be resolved.

To provide DHCP backup, two local DHCP servers could be used. Another alternative, shown in the picture above, is to use servers assigned to the primary 'Local' ICP and a backup in the secondary 'Remote' ICP. Each DHCP server can back up the other. Each DHCP server will have a Scope and rules pertaining to the other subnet. This also requires that DHCP forwarding be enabled on the router. The router will append information to the DHCP broadcast, and it is this that the DHCP servers will use to identify where the request originated and hence how to reply.

Note that although the information in both DHCP servers pertaining to the same subnet is generally identical, it may be that the IP address for the TFTP servers is different. This is dependent upon how the IP addresses are partitioned, and whether the TFTP server is a central resource or within an ICP. See also under TFTP in section 5.1.4 IP Phone Interaction with Routers, DHCP and TFTP Servers.

The phones will take the first DHCP reply and if satisfied with the information will acknowledge this back to the server giving out the address. If the second server tries to give out this same address again, the ICMP echo will result in a reply so this server will know to busy out that address.

The address range within each DHCP server can be partitioned in the following ways:

- The address ranges in each DHCP server can be overlapped
- The address ranges can be split without overlap
- IP addresses can be over provisioned (reduce number of devices per subnet to ensure more IP addresses than devices)

- The address ranges can be split without overlap

Over-provisioning IP Addresses Between DHCP Servers

This is more the scenario that might be encountered in a data environment where DHCP Server backup is required.

Typically 80% of the needed IP addresses are on the primary server and 20% on the secondary server, those on the primary server being used first by virtue of the fact those messages get to the requesting devices first. In a data environment there is also the implied condition that not all of these addresses are required at startup, and so there is also a level of redundancy.

In a more populated environment, with phones, the number of phones and the split of addresses needs to be considered.

Where a phone needs to be resilient then there is a requirement to provide 2 possible IP addresses, one on the primary DHCP Server and a backup on the secondary DHCP server. The size of the subnet also needs to be considered for example, a Class C address would typically allow up to 250 addresses.

Thus, where there are 100 phones and all need to be backed up, there is a requirement for 100 IP addresses on the primary and 100 different IP addresses on the secondary, thus a total of 200 address for 100 phones. This represents a 50:50 split rather than the data 80:20 split.

In this situation then, the scopes of the two DHCP Servers will refer back to the same subnet, but the addresses will be non-overlapping.

For the case where all phones need to be backed up, it can be seen that 50% needs to be available on the one DHCP server. Don't forget that the ICP may be separate from the DHCP server in which case a reserved address should also be made for the RTC and E2T interfaces within the ICP.

A table is shown below as an example:

Table 16: Over-provisioning of IP Addresses Between DHCP Servers

Interface	Primary DHCP Server Addresses	Secondary DHCP Server Addresses
RTC – Fixed address, not available via DHCP	10.10.1.1 – Exception	10.10.1.1 - Exception

Interface	Primary DHCP Server Addresses	Secondary DHCP Server Addresses
E2T – Permanent address	10.10.1.2 – Permanent Lease	10.10.1.2 – Permanent Lease
Phones Basic (100)	10.10.1.3 to 10.10.1.127 (81% used)	
Phones Backup (100)		10.10.1.128 to 10.10.1.254 (80% available)
Gateways, TFTP, etc.	10.10.1.251 to 10.10.1.254 – Exception	10.10.1.251 to 10.10.1.254 – Exception

Although this latter scenario is much more sound in terms of being able to provide adequate IP addresses it also has the disadvantage of reducing the number of IP Phones that can be connected to a particular subnet. This may not be an apparent issue but it increases the number of subnets, and potentially VLANs, that are needed to handle a larger system. Also, if the sub-nets connect to a router through physical ports, rather than a common virtual port with multiple VLANs, then additional physical ports will be required.

TFTP

The TFTP server is used to provide the application code to the IP Phones. There is an internal TFTP server within the MiVoice Business system. An external TFTP server can also be used. For a larger system it is suggested that having a central TFTP server would be of benefit in reducing administration, especially when upgrades or changes are needed. The more dedicated TFTP server also has the advantage of potentially providing increased number of accesses and so the overall system will boot quicker than using the on board TFTP server.

Whether the internal, or an external, TFTP server is used, it will be necessary to ensure that the same application code and version are used for both the primary and secondary servers. This will ensure that there are no compatibility issues when the phones transition from their primary to secondary ICPs. Should there be a difference in the software versions and the ICP requires new code to be downloaded, it will override the DHCP option and force the phone to reboot after loading new code.

The entry in the appropriate DHCP Server should identify the associated TFTP server. So where the information in the scopes between the primary and secondary DHCP servers are virtually identical, this is where they differ. Thus, the primary DHCP should

reference the primary TFTP server, and the secondary DHCP should reference the secondary TFTP server. In this way if the DHCP server and TFTP server are linked and they fail, new devices will be directed to the secondary DHCP and hence secondary TFTP server.

There is one snag to consider though. If the TFTP server is independent of DHCP and fails, any new phones will not be able to obtain their application code and hence will not be brought into service. Phones that are already operational and maintain an/their IP address will not be required to re-boot and hence will not require access to the TFTP server.

FTP for E2T

The FTP server for the E2T is dedicated within the RTC that is physically co-located with the E2T. The E2T will therefore use the RTC address to get its correct application, or running, code.

The only potential difference is with DHCP programming for the E2T card. When the internal DHCP within the ICP is used, dedicated information is provided and so does not need to be programmed. If an external DHCP server is needed then options 66 and 67 need to be defined. These are covered elsewhere in this document.

Locating a Working ICP

Once the phones have obtained IP addresses and their application code, they will also have information regarding the IP address of their primary ICP. Upon registering with that ICP, the phones become operational.

Both the primary and secondary DHCP servers will identify the same primary RTC address for the same scope.

Should the ICP fail to respond, the IP phones are initially equipped with three additional IP addresses that they can use to find an ICP.

Should a secondary ICP not be found from this list, the phone will reset, and further attempts will be made, following DHCP response.

Should a secondary ICP be located, it will have information to identify the primary ICP for this device and also whether it has homed to the appropriate secondary ICP. If it has, the device registers there. If the phone has arrived at the incorrect secondary ICP, it is informed of the address of the correct secondary ICP. The phone then homes to the correct secondary ICP.

Following successful registration with either the primary or secondary ICP, the list of ICP (RTC) addresses is updated. The new list is limited to two IP addresses for the primary and secondary ICPs. The third and fourth IP addresses that the phone had initially received from the DHCP server are purged.

In a large system installation, it is expected that a number of subnets will be in use at one ICP, for example, 700 users and 100 users per subnet would suggest seven subnets will be required. Should a system fail, these phones would need to locate a controller elsewhere in the system. The primary ICP will have already defined the secondary ICP address, assuming that the phones were able to register. In the event that the phones are still in start up they will locate an ICP from the addresses given within the DHCP message. To distribute the loading on the chosen secondary ICP, as well as improve the registration speed, it is recommended to distribute the order of the addresses within DHCP. For example, phones in subnet A might try ICP A followed by ICP B, then ICP C, then ICP D. Phones in subnet B might try ICP A (primary), then D, then A, then B. Each subnet contains the same information, just in a different sequence.

A phone that is designated as non-resilient will hunt for an initial controller to register with. If successful, it will be given the address (single since it is designated as non-resilient) of the primary ICP. If this ICP is unavailable, the phone will reset and re-attempt to register with its ICP.

8.5 Link Failure Detection and Management


Link Failure Detection Mechanism

The link failure detection mechanism is composed of 3 key elements:

- A symmetric ICP-to-phone and phone-to-ICP heartbeat mechanism enables either side to detect when the other has failed, during idle periods. Each direction is decoupled, can be independently controlled and configured, and minimizes bandwidth usage by neither sending nor requiring the receipt of heartbeat messages if other call-control traffic is present.
- The mechanism causes the phone to assume Early Warning behavior if the ICP becomes unresponsive. Early Warning behavior precedes Fail behavior and presents an alternate user interface on affected devices (see [Resilient IP Device Behavior and User Interfaces](#) on page 28). If the ICP becomes responsive again, the devices resume normal operation.
- The mechanism provides a means to specify "special keys" that are guaranteed to end streams in progress and send phones into Fail-over handling while the phones are in either Warning or Fail state.

[Interaction of Parameters in Link Failure Management](#) depicts the timing of the link failure management mechanism.

Figure 59: Interaction of Parameters in Link Failure Management



8.5.1 Link Failure Management Mechanism

- Link maintenance parameters (ConfigureLinkManagementParameters message) are sent to the IP phone once it registers on an ICP.
- Once the parameters are received, if configured to do so, the 2-way heartbeat process begins. If no other messages have been sent, either end will send a null KeepAlivePoll message to the other side. If no messages are received from the other end within the specified interval, either the device or the ICP will close the control link. At this point, two things happen:
 - The phone assumes Fail handling behavior.
 - The phone is removed from service on the ICP.
- Also, once the link maintenance parameters message is received, if configured to do so, the phone makes use of key acknowledge (KeyAck) messages returned by the ICP after every key-press message sent. Failure to receive KeyAck messages within the specified time intervals sends the phone first into Early Warning user interface behavior and then into Fail handling status.
- If configured in the maintenance parameters message, special keys terminate any existing streams as well as the control link, when pressed at any time during either the Early Warning or Fail handling intervals. Special keys may also have a separate forced link-closure timer associated with them if they are not acknowledged by the ICP.
- If the ConfigureLinkManagementParameters message is not sent to the phone (all administrative controls have been set to “off”), the phones underlying TCP mechanisms will remain in effect.

i Note:

The default overall TCP link failure time is set at 30 seconds on both phone and ICP sides, with the phone polling the ICP.

Controller Registry Configuration Timers

! CAUTION:

Do not change the default settings for these parameters without consulting a Mitel representative. For all of the following parameters, the sending side response and polling intervals must account for both the actual timeout value set on the receiving side as well as the maximum acceptable link propagation delay between sender and receiver (including TCP retries). For example, the sending-side heartbeat interval must be shorter than the corresponding receiving-side heartbeat timeout MINUS the maximum expected propagation delay.

The following parameters are configurable in the Controller Registry form, in the MiVoice Business System Administration Tool. For more information, see [Controller Registry Configuration](#).

HeartBeat Rx Link Failure Time-out (1 /10 seconds)

If the device has not received any messages from the system for the time specified by this timer, the device considers the link to have failed, and closes it.

IP Trunk Heartbeat Rx Link Failure Time-out (1 /10 seconds)

If the IP Trunk has not received any messages from the device for the time specified by this timer, the trunk considers the link to have failed, and closes it.

IP Trunk Inactivity Tx Timer (1 /10 seconds)

This timer sets the polling period (interval) at which a KeepAlivePoll message is sent by the IP Trunk when it is idle (no other messages transmitted).

KeyAck EarlyWarning Timer (1 /10 seconds)

If the device does not receive a Key Ack message from the ICP for a SendScanKey message, for the duration of this timer, the device assumes early warning behavior (set specific).

KeyAck Link Failure Time-out (1 /10 seconds)

If the device does not receive a Key Ack for a SendScanKey message before this timeout elapses, the device considers the link to have failed, and closes it.

Phone Inactivity Tx Timer (1 /10 seconds)

This timer sets the polling period (interval) at which a KeepAlivePoll message is sent by the device when it is idle (no other messages transmitted).

Resiliency - Allow Return to Primary ICP (Yes/No)

Controls whether failed-over IP devices re-home to this MiVoice Business system.

Resiliency - Secondary ICP Health Check Interval (seconds)

Sets the waiting period between HealthChecks attempts with a given primary system. Health Checks occur only if a secondary system has IP devices registered with it as a result of a fail-over from a primary system.

Resiliency - Successful Health Checks Prior to Hand Off (1 - 20)

Sets the number of successful Health Checks required from a secondary to a primary system before IP devices are passed back to the primary system after a recovery.

Special Key Time-out (1 /10 seconds)

If one of the special keys is pressed (see [Special Keys List](#)) and no Key Ack is received from the ICP by the time this timer elapses, the device tears down the MiNET link and streams (if any) to the ICP and rehomes.

Special Keys List

This list defines the special keys for use by a resilient phone. The default keys in the phone list are

- Hookswitch
- Cancel

i Note:

Only the hookswitch is configured as a default setting. This setting can be modified by an administrator.

SIP Endpoint Resiliency – Device Heartbeat Interval

Enter the interval at which SIP devices that support the Mitel P-Alternate-Server header sends Options heartbeats to their current controller.

i Note:

Parameter values are assigned in seconds.

SIP Endpoint Resiliency – Peer Heartbeat Interval

Enter the interval at which NOTIFY heartbeats are sent between primary and secondary ICPs. The peer ICP heartbeat interval determines the amount of time during which a SIP device goes to voice mail between failure of the primary ICP and failover to the secondary ICP. The primary ICP sends NOTIFY heartbeats to the secondary ICP at the rate requested by the secondary ICP, and the secondary ICP sends NOTIFY heartbeats at the rate requested by the primary ICP. The rates may be different.

Note:

Parameter values are assigned in seconds.

8.5.2 Link Failure Detection by IP Phones

If a phone fails to receive messaging from the ICP within the time specified by the timers, the phone assumes the control link has failed, closes the link, and assumes fail handling behavior.

Timer Control

If the phone is idle (no user activity), the HeartBeat Rx early warning timer controls the phone's link failure detection. If this timer is set, the ICP is expected to send either normal control messaging or (if the ICP is idle) an intentional poll message within the specified period to let the phone know it is still functioning properly.

During user activity (key presses), the KeyAck link failure timeout controls the phone's link failure detection. A KeyAck message is required for any key press to ensure key-press signals are being received by the ICP, within the specified time interval.

8.5.3 Detecting an ICP or Link Failure Through the Phone Display

For information about the IP Console user interface, see Appendix A.

Failure Detection

When a resilient phone detects an ICP or network failure (or potential failure), its display is frozen.

Registering on Secondary ICP

Boot sequence messages are displayed on IP devices when they are registering on their alternate ICP.

In Service on Secondary ICP

Once a resilient device has registered with its secondary ICP and is in service there, a flashing square appears in the top right-hand corner of the display.

Failure During an Active Call

If a user is engaged in an active call when an ICP or network failure occurs or is about to occur

- The IP device user interface is frozen, with the exception that programmed softkeys are no longer displayed.
- The user hears two quick beeps (warning tone) indicating a failure or potential failure. The user continues to hear the warning tone at twenty-second intervals for the duration of the call.
- If the failure is averted, resilient devices resume normal operation.
- If a failure occurs, the voice stream is maintained until the call is ended. At this time, the device rehomes to its alternate ICP.

Failure in Idle State

If a resilient device is idle when an ICP or network failure occurs or is about to occur

- The IP device user interface is frozen, with the exception that programmed softkeys are no longer displayed.
- If the failure is averted, resilient devices resume normal operation.
- If a failure occurs, the device rehomes to its alternate ICP.

Phone Behavior for Special Keys

If the phone is in either early warning or failure state

- If a special key is pressed, all active streams (whether priority streams or not) are terminated immediately, and the TCP link (if still open) is immediately closed,
- The phone rehomes, regardless of the settings of the special keys timers.

If the phone is not in either early warning or failure state

- If a special key is pressed, and if the SpecialKeys timer value is “off” or not specified, then the regular KeyAck mechanism applies.
- If a special key is pressed, and if the SpecialKeys timer value is specified, then phone behavior is controlled by the SpecialKeys timer (for example, if time expires with no

Acknowledge message received from the ICP, the phone closes both the control link and call streams and fails over to its secondary ICP).

8.5.4 MiVoice Business Health Check

The ICP-to-ICP health check is a polling mechanism between two ICPs. It is used to determine if an ICP is ready (healthy) to register previously handed-off or failed-over IP devices.

When an IP device registers on its secondary ICP, the Health Check process is initiated to that device's primary ICP, that is, the secondary begins sending Health Check messages to the primary. In a Health Check, the secondary is the sending or polling ICP, and the primary is the receiving or polled ICP. The Health Check message interval is a configurable parameter.

A Health Check message consists of three things:

- The IP address of the polling ICP
- The number of devices that the sender would like to return to the receiver (polled ICP)
- The number of Health Checks already sent

When the primary ICP comes back into service and is ready to re-register resilient devices, it begins responding to incoming Health Check messages.

A Health Check response consists of three things:

- The Health Check attempt number to which the ICP is responding
- The IP address of the polled ICP
- The number of devices the ICP can re-register

Once the secondary ICP receives a certain number of consecutive Health Check responses, it initiates the Fail-back process, and sends devices back to the primary ICP when they are idle. The number of responses required is a configurable parameter.

If the secondary ICP cannot establish health check communications with the primary ICP (receives no response) or does not receive the required number of consecutive health check responses, it does not attempt to send devices back to their primary.

Overview

The health check process is initiated when an IP phone registers with its secondary ICP; the secondary ICP opens a socket and starts sending health check messages to the IP phone's primary ICP to determine if the primary ICP is healthy.

Configurable Parameters

For information about resiliency parameters in the Controller Registry form, refer to the System Administration Tool Help and ["Controller Registry Configuration Timers"](#).

ICP Health Check and Handover Maintenance Logs

Three maintenance logs are generated to allow a system administrator or technician to track the behavior of resilient devices. For details on accessing these logs refer to *3300 IP Communication Platform Technician's Handbook*.

Idle Devices

Devices that are not in an active call state are considered idle and are forced by the secondary ICP to rehome to the primary ICP once the primary achieves healthy status. Devices do not Fail-back to or re-register on their primary ICP in order of any priority.

Non-idle Devices

Devices that are in an active call state are not forced to rehome (Fail-back) to their primary ICP until they become idle. An IP device is considered not-idle if it meets any of the following conditions:

- The device is in talk or dialing state (the handset or headset is off-hook, or the speaker is on).
- The device is audibly ringing, excluding threshold alert (any line, device or attached PKM).
- The device is an emergency-call monitor and has unacknowledged emergency calls in the system.
- The device has a call on Hold.
- The device is in a Superkey session.
- The device is being group paged.
- The device has a DSS/BLF key that is lit or flashing.

Fail-back Delay

If a device on its secondary ICP is in any of the above states it will not fail back to its recovered primary ICP until the device becomes idle.

Regarding emergency calls, if the last emergency call in the system is acknowledged by a device that is in Fail-back delay, the pending Fail-back is initiated as soon as the emergency call log is cleared; however, if the last emergency call in the system is acknowledged by a third party, Fail-back is not initiated. If this happens, the device will not Fail-back until it enters a non-idle state and then becomes idle again.

Note:

1. For pre-Release 4.0 ICPs, if the phone is not instructed to perform Health Checks by the MiNET Configure Link Management Parameters message, it will perform control handling only when the TCP socket is determined to have failed.
2. When a phone receives a message to enable polling, any existing polling-period setting is overwritten, and timing restarts. For example if a polling interval of 5 seconds is received in the middle of a polling interval of 60 seconds, the 60-second timer is cancelled and a new one begins. This behavior prevents timeouts on the first request to the phone.
3. Timer values are configured in tenths of seconds, from 0.1 to 6553.5 sec, with 0 being interpreted as “off”.

8.5.5 MiVoice Business Behavior During Link Failure Detection

MiVoice Business behavior for link failure detection is the same as for the IP phones. If an ICP fails to receive messaging from a phone within the time specified by the timers, the ICP assumes the control link has failed, closes the link, and assumes Fail handling behavior, that is, the phone is removed from service on that ICP.

Timer Control

While idle (no other set messaging activity), ICP link failure detection is controlled by the phone inactivity Tx timer. If this timer is set, the phone is expected to send either normal messaging or (if idle) an intentional poll message within the specified period, to let the ICP know it is still functioning properly.

8.5.6 IP Phone Behavior During Resilient Mode Operation

For more information, see [Resilient IP Device Behavior and User Interfaces](#) on page 28.

IP Phone Behavior During Primary ICP Failure

If a failure occurs on a phone's primary ICP during a call, the user hears two quick beeps indicating that

- The phone's primary controller has failed, or the network link between the phone and the controller has failed.
- The current call will be maintained, but for the remaining duration of the call, the user cannot use the keypad or phone features.
- Once the user ends the call, the IP phone begins the process of registering on its secondary controller

If a failure occurs on a phone's primary ICP while the phone is on-hook, the phone immediately begins the process of homing to and registering on its secondary ICP.

Failure Detection and Fail-over Times

Fail-over times can vary, depending on the time it takes a phone to detect a failure and also on the loading of the secondary ICP that the phone is failing over to.

A phone takes 30 to 60 seconds to detect that its ICP has failed or is unreachable. This time is based on the default value for the Heartbeat Rx Link Failure Timeout Parameter.

Once it has detected an ICP failure, a phone takes an average of 45 seconds to fail over to a secondary ICP that with medium loading and a maximum of 90 seconds to fail over to a secondary ICP with heavy loading.

Resilient Phone Functionality

Phone functionality on the secondary ICP consists of

- Basic call handling, incoming and outgoing calls.
- Emergency Call, Phonebook and Messaging Superkey sessions
- Feature access keys and codes

Functions that are not supported include

- Applications that were running on a failed ICP.

The effect of not supporting certain applications has an impact on devices such as IP appliances. Following a Fail-over, these devices rehome to their secondary ICP. Basic telephony functions will be maintained, and limited graphics and directory support are provided. There is a potential delay, however, in providing graphics support following a switch between controllers.

Typically the graphics support takes about 2 minutes to render per device (rendered sequentially). This means that for 100 phones, it takes approximately 200 minutes for all WebSets to have graphics after they fail back to their primary ICP (this is because the primary has just come back into service, and graphics are being re-rendered). If

graphics are needed within a certain period of time, this constraint affects the number of Websets that can be supported on particular controller.

In contrast, when WebSets fail over to a secondary ICP that has been running for a number of hours, graphics are already rendered on that controller, and phones that fail over will immediately get graphics.

- Phone features access codes

For a list of the Feature Access Codes that cannot be accessed or invoked when an IP device is registered with it's secondary ICP see [Resilient Feature Support](#) on page 49.

- Feature access keys (multiline sets)

For a list of the Feature Access Keys on multiline sets that do and do not support resiliency, see [Resilient Feature Support](#) on page 49.

- DPNSS Features

See [Resilient Feature Support](#) on page 49, for a list of the DPNSS Features that support resiliency.

Conference Calls

Analog Trunk Lockup

Trunk lockup could occur if a party using a trunk to access a conference call is the first party to hang up and leave the conference and the CO does not disconnect the trunk.

To safeguard against this trunk lockup possibility the installer should request that the CO (The service provider) provides a trunk clear down indication. For example the CO should provide a line break or a battery polarity reversal as an indication that the trunk should be cleared down.

Note:

At the present time the ICP does not provide a Dial Tone detection mechanism for conference calls so that trunks can be automatically cleared down.

Privacy Holes Created When an ICP Fails

A privacy hole is a situation where a user's audio transmission continues to stream to "somewhere" and the user is not yet aware that he/she has been disconnected from the other parties. Two failure situations are described below.

Scenario: ICP Hosting the Conference Call Fails

During a conference call, if the ICP hosting the conference fails then the conference bridge also fails. Phones that are connected to the ICP hosting the conference call disconnect and rehome to their secondary ICPs; these phones are not subject to a privacy hole.

If a remote phone is connected into the conference via a TDM trunk it will enter resilient talk state and request End Of Stream (EOS) detection to the E2T. This failure detection is based on the configurable heartbeat setting for the trunk and the EOS detection which is configurable as well, thus the time before the trunk connection is taken down will be the trunk heartbeat failure time + EOS detection time.

The EOS detection default value is 5 seconds so in this scenario the remote TDM calls will be dropped 5 seconds after the IP trunk failure is detected.


In the case where a remote IP phone is streaming directly to the ICP hosting the conference, this phone will stay connected until the user goes on-hook. In this situation there is a possibility of a privacy hole.

Scenario: The Non-Hosting ICP Fails

In the case where the non-hosting ICP fails, the ICP hosting the conference will drop the phones connected to the non-hosting ICP from the conference. When the phones connected to the non-hosting ICP detect EOS they will rehome to their secondary ICPs.

The following figure depicts the ICP fault scenarios that are described above:

Figure 60: Resilient Conference Phone Behavior



IP0772

Music on Hold

When a user has been placed on MOH and the controlling ICP fails, the call is cleared down when the local ICP determines that the IP trunk has failed.

Non-resilient Phone Operation & Configuration

Special consideration must be given to network and phone configuration in a mixed cluster (containing both resilient and non-resilient devices).

At initial startup, phones have no knowledge of whether or not they are resilient. As a result, they operate in the same manner to network services such as DHCP and TFTP. A phone is resilient-aware, once it registers with an ICP.

Two main aspects to consider are whether the DHCP server is external to the ICP or whether it is integrated within the ICP. With integrated DHCP, if the ICP fails, the DHCP is also likely to fail. With external DHCP, there might be a DHCP failure but not an ICP failure, in which case call control resiliency does not apply.

The fact that the phones make no distinction in network operation requires you to consider, as with DHCP services, that a phone may register with a DHCP server and require an IP address, even though it might not register with a working or designated ICP. This means that non-resilient devices could acquire IP addresses that would be needed for resilient devices, negating any call control settings and locking out resilient devices. To avoid this scenario, you can consider segregating various operations and over-provisioning IP addresses to ensure adequate available addresses for phones.

Some methods of segregation include

- Segregate resilient and non-resilient devices into different subnets and VLANs, and restrict DHCP Forwarding between subnets.
- Provide MAC filtering within DHCP, so primary gives IP address, secondary does not
- Static IP address for non-resilient phones
- Over-provision IP addresses so non-resilient treated as resilient until registered with call control

Segregating Non-resilient Phones into Another Subnet

A larger system is likely to require a number of subnets in order to handle the quantity of phones. The non-resilient phones can be identified and placed exclusively into one of these subnets.

The non-resilient subnet can have its own DHCP server, or it can have forwarding across the router enabled. Forwarding to the primary DHCP server only should be enabled.

The primary DHCP will include a scope for the particular subnet with IP addresses.

Where a secondary DHCP server exists, this need *not* include any scope information for the non-resilient phones. If the primary DHCP server is separate from the primary ICP, then it may be that the primary DHCP server has experienced a failure. In this way, the phones would still be able to register with the primary ICP even with an unreachable primary DHCP server.

Where the primary DHCP server is located with the primary ICP then adding scope information into secondary DHCP server is irrelevant since there is no ICP to register with.

Where a single DHCP server is used then the full range of IP addresses can be used. Where a secondary DHCP server is used then over provisioning and/or splitting the IP addresses is recommended to improve handover once the IP leases start to expire.

MAC Filtering in DHCP Scope

MAC filtering in DHCP scope identifies those phones destined to be resilient and non-resilient. It requires that the MAC addresses of all the non-resilient phones be entered into both the primary and secondary DHCP servers.

The information in the two DHCP servers would essentially be the same. The IP address range could be split to allow only some of the non-resilient sets to continue to operate should the DHCP server fail and the ICP still be working. For example, there may be 100 IP addresses allocated on the primary DHCP, but only 30 on the secondary, or even 0 (zero). Resilient sets would have a 1:1 mapping between the two servers.

Where the primary DHCP server is located with the primary ICP then adding scope information into the secondary DHCP server is not necessary.

Static Programming of IP Addresses

The non-resilient phones can be assigned IP addresses statically, in this way the phones will always have IP addresses and there will be no issues regarding DHCP support for resilient and non-resilient sets. However, if a particular network has a large number of non-resilient phones this method could be labour intensive. For details on statically programming phones, see [Static IP Configuration for IP Phones](#) on page 291.

Static IP Configuration for IP Phones

Level 1 Diagnostics: Static IP Setup Mode


Normally, IP information is retrieved via one or more DHCP sessions with a network server. However, for diagnostic purposes the user can enter IP information manually. Manually entered data is referred to as “static” data.

If static data has been entered, when the phone boots up it does not have to use DHCP; it uses the static data instead. Because the static data is stored in non-volatile FLASH memory, it is used every time the phone boots up, until such time as the data is deleted.

You can configure an IP phone to use a combination of static and dynamic data. You can manually enter any subset of parameters, and the phone will use DHCP to acquire values for the rest. However, if static data has been entered, it will ALWAYS be used instead of dynamic DHCP data – thus the user should triple-check to make sure that all static fields have proper values, and that fields for which DHCP data is desired are not active (i.e., their static values have been deleted).

To enter the static IP utility, press the Volume Up key and hold it for at least 3 seconds while powering-up the IP phone. The following is displayed on the phone for approximately 5 seconds:

Figure 61: Static IP Utility Display Prompt



- If you press the star key (*), the session is ended, and you must re-power the phone.
- If you press 0(zero), you can view previously enabled static parameters and use the volume up and down keys to scroll through the menu. In view mode, screens do not allow input.
- If you press the number sign (#), you can enter static IP parameters.

User Keys

As with most display sequences, the Volume Up/Down keys are used to traverse the menu list. The '*' key is used to back up in an entry (i.e., to correct an error), and the '#' is used to insert a decimal and move forward into another IP address field.

Figure 62: Menu User Keys




Delete / Enter Screen

If some or all of the static parameters are active, and the user wishes to return to using only DHCP, then '*' should be chosen in the Delete/Enter screen. In this case, the user will be taken immediately to the Save And Exit screen.


Note that when the user deletes settings, all static IP parameters are removed from Flash; thus the next time this utility is used, values must be re-entered from scratch. To delete a single field, the user can press the Superkey (OK key on the Webset) while the field's screen is being displayed

Figure 63: Delete / Enter Screen



USE PRESENT SETTING?
*** = DELETE # = ENTER**

Figure 64: Save and Exit Screen



SAVE AND EXIT?
*** = NO # = YES**

If the user wishes to enter IP data, the user should choose '#' at the Delete/Enter screen. In this case, the user will be taken to VLAN ID input.

VLAN ID Screen

Following the Delete/Enter screen, the user is prompted to "INPUT VLAN ID". This is the VLAN ID that will be inserted into packets sent by the phone. If VLANs are being used, the ID should be entered.

Figure 65: VLAN ID Screen



INPUT VLAN ID:

As with all screens, if no value is required then nothing should be entered. And whether or not a value is entered, the screen should be exited by pressing a volume key (the down-key to advance).

Priority Screen

Following the VLAN ID SCREEN, the user is prompted to "INPUT PRIORITY". This is the Priority that will be inserted into packets sent by the phone. If VLANs are being used, priority should be entered.

Figure 66: Priority Screen


A rectangular box representing a screen with the text "INPUT PRIORITY:" centered inside.

INPUT PRIORITY:

IP Address Screen

Following the PRIORITY SCREEN, the user is prompted to "INPUT IP ADDRESS". This is the IP address of the phone itself.

Figure 67: IP Address Screen


A rectangular box representing a screen with the text "INPUT IP ADDRESS:" and "XXX.XXX.XXX.XXX" centered inside.

INPUT IP ADDRESS:
XXX.XXX.XXX.XXX

Subnet Mask Screen

Following the IP ADDRESS SCREEN, the user is prompted to "INPUT SUBNET MASK".

Figure 68: Subnet Mask Screen


A rectangular box representing a screen with the text "INPUT SUBNET MASK:" and "XXX.XXX.XX.XXX" centered inside.

INPUT SUBNET MASK:
XXX.XXX.XX.XXX

Default Gateway Screen

Following the SUBNET MASK SCREEN, the user is prompted to "INPUT DEFAULT GATEWAY" (This is not a spelling mistake: some IP phones have a screen width maximum of 20 characters.)

Figure 69: Default Gateway Screen

A rectangular box representing a screen with the text "INPUT DEFAULT GATEWAY" and "XXX.XXX.XX.XXX" centered inside.

INPUT DEFAULT GATEWAY
XXX.XXX.XX.XXX

RTC Address Screen

Following the DEFAULT GATEWAY SCREEN, the user is prompted to “INPUT RTC ADDRESS”. This is the address of the RTC card on the PBX.


Figure 70: RTC Address Screen

TFTP Address Screen

Following the RTC ADDRESS SCREEN, the user is prompted to “INPUT TFTP ADDRESS”. This is the address of the TFTP server used to download the main and boot load images.

A rectangular box representing a screen. Inside, the text "INPUT TFTP ADDRESS" is centered in a bold, black, sans-serif font. Below it, the text "XXX.XX.XXX.XXX" is centered in a smaller, black, sans-serif font.

Figure 71: TFTP Address Screen

A rectangular box representing a screen. Inside, the text "INPUT DNS ADDRESS" is centered in a bold, black, sans-serif font. Below it, the text "XXX.XXX.XXX.XXX" is centered in a smaller, black, sans-serif font.

DNS Address Screen

Following the TFTP ADDRESS SCREEN, the user is prompted to “INPUT DNS ADDRESS”. This is server that will be used during Web browsing, to resolve host names into IP addresses.


Figure 72: DNS Address Screen

A rectangular box representing a screen. Inside, the text "INPUT DNS ADDRESS" is centered in a bold, black, sans-serif font. Below it, the text "XXX.XXX.XXX.XXX" is centered in a smaller, black, sans-serif font.

Manual Input of Technician IP Address

Following the DNS ADDRESS SCREEN, the user is prompted to enter “TECHNICIAN IP ADDR”. This is the address of a debugging utility, and if the utility is not being used, then this entry can be ignored.

Figure 73: Technician IP Address Screen




TECHNICIAN IP ADDR?
(optional) * = NO # = YES

Entering '#' takes the user to the INPUT TECHNICIAN ADDRESS SCREEN.

Input Technician IP Address Screen

Figure 74: Input Technician IP Address Screen



INPUT TECHNICIAN ADDRESS:
XXX.XXX.XXX.XXX


Once the address is entered, or if the user chose not to enter the address, we move to the NVRAM storage display.

Storing Static Parameters

Following the INPUT TECH ADDRESS menu item(s), the user is prompted to "STORE IN NVRAM?".

It is significant that we store in non-volatile RAM, or FLASH. Thus, when the set is powered-up, it will not use DHCP to retrieve IP information already entered as static data. Further, if the FLASH is upgraded via the TFTP server, the static data will be retained, and unless manually deleted, will be used at power-up. To delete static data, the user must use the Delete/Enter Screen described previously.

Figure 75: Store in NVRAM Screen



STORE IN NVRAM?
* = NO # = YES

Figure 76: Saving to NVRAM Screen





Figure 77: NVRAM Save Complete screen



The set will then reboot, and will use the static IP data.


Level 1 diagnostics: Teleworker Setup Mode

This is a shorter version of the IP STATIC SET-UP MODE.

This mode should be used by teleworkers. It allows the user to set both the RTC and TFTP IP addresses with a single entry. All other IP data will be retrieved from the first DHCP server to make an offer.

To enter the static IP utility, the user must press the '7' key and hold it for at least 3 seconds while powering-up the IP Phone. The following display will appear for approximately 5 seconds:

Figure 78: Configure Teleworker Prompt



If '*' is pressed, the session is ended, and the user must re-power the set.

And if '#' is pressed, the user can enter static parameters.


User Keys

The VolumeUp/Down keys are used to traverse the menu list. The '*' key is used to back up in an entry (i.e., to correct an error), and the '#' is used to insert a decimal and move forward into another IP address field.

Delete / Enter Screen


If the user wishes to return to using only DHCP, then '*' should be chosen in the Delete/Enter screen. In this case, the user will be taken immediately to the Save And Exit screen.

Figure 79: Delete / Enter Screen



A rectangular screen with a black border. The text is centered and reads: **USE PRESENT SETTING?**
*** = DELETE # = ENTER**

Figure 80: Save and Exit Screen




A rectangular screen with a black border. The text is centered and reads: **SAVE AND EXIT?**
*** = NO # = YES**

If the user wishes to enter data, the user should choose '#' at the Delete/Enter screen. In this case, the user will be taken to the Teleworker Gateway screen.

Teleworker Gateway IP Address Screen

The user is prompted to press # for enter or New. TELEWORKER GATEWAY appears on the display screen. This is the IP address of the TELEWORKER GATEWAY server acting as a Proxy for the PBX within the corporate firewall.

Figure 81: Teleworker Gateway IP Address Screen



A rectangular screen with a black border. The text is centered and reads: **INPUT TELEWORKER GATEWAY ADDRESS:**
XXX.XXX.XXX.XXX

Storing Teleworker Gateway IP Address

Following the Teleworker Gateway IP Address screen, the user is prompted to STORE CHANGES or STORE IN NVRAM.

Figure 82: Store in NVRAM Screen




Figure 83: Saving to NVRAM Screen





Figure 84: NVRAM Save Complete Screen



The set will then reboot, and will use the Teleworker Gateway data for both the RTC and TFTP addresses.

IP Phones and E911

IP phones that are being used for E911 support can be dynamically assigned IP addresses via the DHCP server.

In 3300 Release 5.2 or later, the 3300 ICP when configured correctly will perform automatic CESID updating for moved IP phones. The in th detects IP phones by analyzing data reported from the Spanning Tree Protocol/Rapid Spanning Tree Protocol or the Cisco Discovery Protocol.

To ensure the correct operation of E911 services and automatic CESID updating, you must follow certain network guidelines. For details refer to the Engineering Guidelines under "E911 Support and Location Change Indication".

Line Interface Module and E911

In 3300 Release 6.0 or later, the Line Interface Module (LIM) gives 5320, 5330, 5340, and 5360 IP Phones the ability to access an analog line via a POTS interface in addition to the regular ethernet connection.

The Line Interface Module (LIM) can be used with a 5320, 5330, 5340, or 5360 IP Phone to provide

- enhanced emergency services in the enterprise environment
- 911 support to the Teleworker.

Configuration of the LIM as a second channel device and the emergency call routing option is required on both the primary and secondary ICP to support LIM E911 interworking on a resilient IP phone.

For additional information on LIM usage, limitations, and the necessary ICP programming, refer to the MiVoice Business *System Administrator Tool Help*. For details on how to install and use the LIM, refer to the set installation guides and set user guides.

Unsupported Devices

Currently, the following devices do not fully support Enhanced 911 (E911) operation:

- Mitel Teleworker Solution
- Hot Desk users
- UC Express and UC Advanced users
- Mobile IP phone users or users who carry a phone from location to location.

Note:

E911 support for Teleworker can be provided via a Line Interface Module. See [Line Interface Module and E911](#) on page 299.

If an unsupported device is moved to a different physical location, the System Administrator must update the Customer Emergency Services Identifier (CESID) database with the new physical location.

Unsupported Configurations

CESIDs will not be automatically updated for IP phones that are moved if

- the IP phone is connected to an Ethernet hub.
- the IP phone is connected to a L2 switch that does not have CDP or STP/RSTP enabled.
- multiple IP phones report connectivity to the same L2 port (the system detects this condition upon device registration).

Local regulations that are presently under review for E911 suggest that a phone call should now be locatable within a 7000 square foot area (15 meter radius). One phone within such a defined area must maintain power, and both its location and exterior limits of the radius must be visibly marked. For recommendations on powering critical (E911) phones, see [Network Power Provisioning](#) on page 266 and the section titled "Power" in the 3300 ICP Engineering Guidelines.

Resilient Hot Desk Devices and E911 Service

A hot desk user can log in to any Hot Desking enabled device. All calls made from a hot desk device use the CESID of the registration DN of the device itself (not the Mobile DN of a user who may be logged in), regardless of whether or not a hot desk user is logged in.

If an emergency call is made from a mobile DN, the SMDR log indicates that the call originated from the mobile DN with the CESID of the registration DN. The call originates from the ICP hosting the mobile DN.

A mobile DN cannot be configured with a CESID. This is enforced by CDE when programming the CESID Assignment form. Mobile DNs are not listed in the form.

Resilient Devices and E911 Service

CAUTION:

The System Administrator should note that devices that are in Teleworker mode and are connected outside of the corporate firewall will not have E911 calls blocked. E911 calls placed from such devices may report an incorrect CESID, or may be outside the coverage area of the Public Service Access Point (PSAP).

CAUTION:

E911 service is not guaranteed to Teleworkers. Since a Teleworker can be geographically located away from the main office, if the Teleworker relocates, the system administrator must perform a manual database update to reflect the user's location. For this reason, it is strongly recommended that you not provide E911 services to a Teleworker phone.

Impact of Resilient Operation on Administrative Features

Note:

E911 support for Teleworker can be provided via a Line Interface Module. See [Line Interface Module and E911](#) on page 299.

For details on how resilient operation affects administrative features see [Administrative Feature Interactions with Resilient Routing](#) on page 83.

8.6 Connecting and Configuring IP Network Devices

This section provides the network designer and installer with configuration information and guidelines to ensure that Mitel products and third party networking devices work together in concert to provide a resilient network solution.

This section deals with the more detailed instructions and configurations that apply to network equipment that is needed to make the MiVoice Business system, phones, and applications resilient.

8.6.1 Connecting MiVoice Business Systems to L2 Switches

3300 Release 6.0 and Release 7.0 introduced the CXi and Mx controllers respectively. Both controllers support the Rapid Reconfiguration of Spanning Tree (RSTP) portion of IEEE Standard 802.1w. RSTP is backwards compatible with STP and serves the same purpose as STP. The main difference between the two protocols is that RSTP will reconfigure the network more rapidly than STP.

As of 3300 ICP Release 5.0, the following ICPs support the Spanning Tree Protocol (STP) portion of IEEE 802.1D:

- LX Controller
- 700-User 3300 Controller, if upgraded to 5.0 or later
- 250-User 3300 Controller, if upgraded to 5.0 or later

STP and RSTP enable physical layer resiliency between the ICPs and the Layer 2 Distribution Network. The STP/RSTP feature can be left disabled (default setting), or it can be enabled by the System Administrator.

The following sections describe how to connect an ICP to the L2 Distribution Network. For additional guidelines regarding networks that use STP/RSTP, [Spanning Tree Protocol STP and Rapid Spanning Tree Protocol RSTP](#) on page 335.

Physical Connection, Layer 2 Switch to ICP

- If Spanning Tree is not running on the L2 switch, there should be only one physical connection between the distribution network L2 switch and any version of the ICP. The connection should be made with a Category-5 or better UTP cable.
- If Spanning Tree is running on the L2 switch and you are using a controller with STP enabled, there can be multiple physical connections between the distribution network L2 switch(es) and the ICP. The connections should be made with a Category-5 or better UTP cable.

Note:

The System Administrator should enable STP on the ICP before completing the physical connections to the L2 switch(es).

- If STP/RSTP is running on the L2 Switch and you are using a controller with STP/RSTP disabled, use only one physical connection between the distribution network L2 switch and the ICP. This will ensure that loops are not created. Make the connection with Category-5 or better UTP cable.

Note:

If the System Administrator chooses to have multiple connections between the ICP and the L2 switch, the following requirements must be met:

- STP must be running on the L2 switch.
- The L2 port used to connect to the ICP must not be set to PortFast, this is to ensure that this port fully participates in the STP algorithm.
- The System Administrator must ensure that no one alters the above settings on the L2 switch at a future date.

Port Settings for Connecting L2 Switches to ICPs

The L2 switch port that is used to connect to the ICP should be configured with the following settings:

- If Spanning Tree is not running on the L2 Switch, the L2 Switch port should be configured to allow for speed and duplex Auto-Negotiation. The port should be capable of running 100Mbps/s FULL DUPLEX.
- If STP/RSTP is running on the L2 Switch but STP/RSTP is disabled on the ICP, then the ICP is not participating in STP/RSTP.

In this scenario, if there is only one physical connection between the ICP and the L2 switch, the ICP should be treated as an end point, and

- the L2 Switch port should be configured for PortFast operation (ensures that this port skips the first stages of the STP Algorithm and directly transitions to Forwarding mode).
- the L2 Switch port should be configured to allow for speed and duplex Auto-Negotiation. The port should be capable of running 100Mbps/s FULL DUPLEX.

In this scenario, if there are multiple physical connections between the ICP and the L2 switch the ICP should not be treated as an end point, and

- the L2 Switch port should not be configured for PortFast operation (ensures that this port participates in the STP/RSTP algorithm).
- the L2 Switch port should be configured to allow for speed and duplex Auto-Negotiation. The port should be capable of running 100Mbps/s FULL DUPLEX.
- If STP/RSTP is running on the L2 switch and STP is enabled on the ICP, then the ICP is participating in STP/RSTP. There are two or more physical connections between the ICP and the L2 switch(es), and as a result, the ICP must be treated as an STP/RSTP participant, and
 - the L2 Switch port should NOT be configured for PortFast operation (PortFast should be disabled to ensure that this port fully participates in the STP/RSTP algorithm).
 - the L2 Switch port should be configured to allow for speed and duplex Auto-Negotiation. The port should be capable of running 100Mbps/s FULL DUPLEX.

VLAN and Priority Configuration, Layer 2 Switch to ICP

These ICPs do not handle VLAN and priority directly at the local LAN connections. This is similar to many PC servers. Therefore the correct voice VLAN and priority need to be configured at the external L2 switch port. The data between the ICP and the external L2 switch port will be 'untagged'. Thus, VLAN conversion will take place at the external L2 switch port. Typically the ICP will be connected at the distribution layer within a larger corporate network. Ports should be configurable to accept untagged information, pass this on to a specified VLAN, as well as accepting tagged information. The link between the LAN and ICP must not be shared with other non-voice devices. Ideally, this should

be a dedicated link, as with most servers. Any variations should be carefully considered taking in to account the amount of data already on this connection.

All ICPs Except the CXi

The CXi controller handles VLAN and priority at the local LAN ports (Port 1 through Port 16). Only ONE VLAN exists on the product and this is VLAN1. Two priority queues are available at the local ports. Priority (COS) 0 to 3 in the low priority queue and Priority (COS) 4 to 7 in the high priority queue. Untagged data will be tagged with VLAN1 and low priority.

When the internal DHCP server is used on the CXi, the VLAN and Priority (COS = 6) will be sent to the phones and these options need not be configured.

When an external DHCP server is used with the CXi, the VLAN option MUST be set to VLAN1 and Priority option set to high, typically 6.

VLAN priority can be enabled at the uplink port (17) to maintain priority to voice, and is a recommended setting when phones and PCs share the connection path. Priority and VLANs should also be enabled between expansion units for the same reasons, such as voice priority.

Further details can be found in the *MiVoice Business Engineering Guidelines, Technician's Handbook* and the *System Administration Tool Help*.

8.6.2 Connection Scenario for Applications Servers

Physical Connection, Layer 2 to Application Servers

There should be only one physical connection between the Distribution Network L2 Switch and the application server.

The connection should be made with a Category-5 or better UTP cable.

Mitel Application Servers do not support redundant NIC configurations, however, it is possible that 3rd Party Application servers might utilize redundant NICs.

If a 3rd Party Application Server supports redundant NICs the following points should be considered:

- The Server might assume responsibility for preventing a network loop by disabling one port, monitoring link integrity on the live port and making the decision to switch over ports in the event of a failure.
- Some other Server implementations might provide two NICs with separate MAC addresses, so that there are in fact two network access paths into/out of the server.

In the cases described above, the Server would be connected to the L2 switching infrastructure with two physical connections, however the Network Administrator must verify that the Server is capable of preventing a network loop.

L2 Port Configuration, Application Servers

The L2 switch port that is used to connect to the application server should be configured with the following settings:

If STP/RSTP is running on the L2 switch:

- The port should be configured for PortFast operation, this will ensure that this port skips the first stages of the STP Algorithm and directly transitions to Forwarding mode.
- The port should be configured to allow for speed and duplex Auto-Negotiation.

If STP/RSTP is not running on the L2 switch:

- The port should be configured to allow for speed and duplex Auto-Negotiation.

Virtual LAN Connection, Application Servers

Typically Application Servers are PC based and do not support VLAN capable NIC cards.

In the case of a dedicated voice application server, then the LAN port should be configured to accept the 'untagged' data from the server and translate this to the appropriate voice-VLAN and with the appropriate priority.

8.6.3 Connecting IP Phones and Desktop PCs to Layer 2 Switches

This considers both the standard PC and desk top phone as well as specialist phones and PCs used with voice applications. Configurations may differ between these devices.

Physical Connection, Desktop PCs and IP Phones

For the standard IP Phone and PC, the PC should support a 10/100 NIC card and this should be set to Auto-Configure. This should be connected to the IP-Phone PC port via CAT5 wiring. Power feed will not be provided over this connection.

The other LAN port on the phone should be connected to the LAN port on the Layer 2 switch. Again, CAT5 wiring should be used. Power feed can be provided over this link. Where power feed is provided from the LAN Layer2 switch, this needs to be configured, or an optional power dongle can be connected in this path to request power from the Layer2 switch, e.g. Cisco Catalyst 3524-XL PWR. Where power is not provided from the Layer2 switch, this can be provided by other means such as an in line power hub (as a group, or individual units) or via an external power 'brick' connected directly to the phone.

Note:

The Layer 2 switch that is integrated into the CXi controller, supports Power Over Ethernet. For details on powering Phones over ethernet cables, refer to the 3300 ICP Engineering Guidelines.

Only one LAN connection is allowed from the phone back to the network since this is an end device connected to the access layer of the network.

L2 Port Configuration

The L2 port should be configured for auto speed and duplex configuration. Ideally the port should be capable of 100Mbps/s FULL Duplex.

The port should also support at least two priority queues, voice being sent to the higher priority and data to the lower priority.

If STP/RSTP is running on the L2 switch:

- The ports should be configured as STP/RSTP 'Portfast' or disabled. This LAN port is an access port and should not need to extend Spanning Tree Protocol since the connected device is an end device.

If STP/RSTP is not running on the L2 switch:

- There are no configurations for STP/RSTP protocol required.

Virtual LAN Connection, Desktop PCs and IP Phones

The port should be set up to handle both 'untagged' traffic destined for the data VLAN (default_VLAN or Native VLAN). It should also handle 'tagged' data destined for the Voice VLAN(s). The enabled VLAN should match the VLAN that the attached phone will receive through DHCP, i.e. if the phone will be told VLAN 2 from DHCP, then this same VLAN should be enabled at this port.

Certain voice applications running on PCs may appear on the data VLAN since these are incapable of running with 'tagged' protocol. In these situations the voice will also be marked as high priority in the Type of Service field. It may be a requirement for these ports to also enable Layer3 TOS to COS/Priority mapping, e.g. precedence of 5 may be mapped to COS of 6. Examples of such a voice application include the Unified Communicator phone and also the PDA phone.

8.6.4 MiVoice Business Console Registration Limitations

The Audio Coding Law (CODEC) that is to be used with the console is determined at the time of installation on the PC and is a fixed value. The MiVoice Business Console cannot register on a secondary ICP that uses a different coding law. As a result, the console cannot be configured to be resilient across U-law / A-Law boundaries; that is, consoles programmed with U-Law cannot register on a secondary ICP that uses primarily A-Law.

8.6.5 IP Phone Interaction with Routers, DHCP, and TFTP Servers

IP phones rely on the DHCP server to provide them with certain network information, this information is dynamically downloaded to the phones. The following information is provided to the IP phones by the DHCP server:

- DSCP value, Layer 3 Priority
- VLAN information
- Priority (COS/Layer 2 Priority) information
- Multiple RTC IP addresses, up to 4 RTC addresses are accepted by the phone
- One TFTP server IP address
- One Router (Gateway) IP address and the Router default mask

Note:

The Internal DHCP server is capable of providing multiple TFTP server and Router IP addresses. However, at the present time the IP phone will only accept the first TFTP server IP address and the first Router IP address that is provided, the IP phone will ignore additional addresses.

For information about programming the internal or external DHCP server, see [DHCP Servers](#) on page 314.

If a DHCP server is not available on the network the IP phones will have to be statically programmed with the network information that the DHCP server would have provided. The current phone load only allows a single device IP address, such as RTC or TFTP to be entered statically. For details, refer to [IP Phone Behavior During Resilient Mode Operation](#) on page 286.

Unsupported IP Phones

Early single port IP sets such as the 4015 IP and 4025 IP do not support clustering or resilient operation.

Multiple TFTP Servers

When there is more than one TFTP server on the network it is recommended that these servers should be loaded with a common software load for the phones.

IP Phones, DHCP IP Addresses, and Lease Timing

Lease time refers to the length of time that a DHCP server has dynamically assigned (leased) an IP address to a device. The device needs to communicate with the DHCP server before the lease time has expired to renew the lease. Guidelines regarding values for lease times are outlined below.

- Long lease times are useful to maintain device IP addresses in the event that an ICP fails, for example 3 days ensures enough time for the administrator to return the failed ICP to service.
- Short lease times are better where it is expected that there will be frequent Moves/ Adds/Changes, e.g. ½ hour
- Could provide 'static' IP addresses to phones, but this will require manual removal from the DHCP database if phones are moved between sub-nets. Furthermore this defeats the purpose of the DHCP mechanism, which is to automate the system administration function and eliminate human errors.
- The E2T in the ICP MUST be assigned a static IP address using the E2T MAC address as the filter.
 - The E2T is often programmed manually with a static IP address using a maintenance cable, thereby avoiding the need to program this into DHCP.)
- When deploying multiple DHCP servers in a network the DHCP servers must be programmed with common information i.e. Lease Time values.
- This is needed to allow the device lease timer to be renewed, even if the initial DHCP server is no longer present, i.e. one in an inactive ICP.
- Allows the IP address currently assigned to a phone to be maintained and used in the event that the initial DHCP is no longer present.
- The system administrator must ensure that the network routers allow ICMP Echo forwarding and DHCP forwarding.
- Phones may receive multiple DHCP replies after an initial broadcast. The phones will take the first DHCP response.

IP Consoles, DHCP IP Addresses, and Lease Timing

If an ICP is part of a resilient cluster any IP Console, which includes PC and the Console keypad must be configured as follows:

- For both devices the DHCP Lease Time should be set as permanent.
- For both devices the IP address should be statically assigned.

i Note:

If a console is marked as permanent and it moves out of the subnet, then the installer/maintenance will also need to be informed that they will need to access the DHCP server to mark the IP address as available.

IP Phones and DHCP IP Address Distribution

There are two models the system administrator can use for distribution of IP addresses from DHCP servers, overlapping IP address ranges or splitting the address range into separate groups within a subnet. Splitting the address range is the recommended method to use. Guidelines are as follows.

- Overlapping should be possible, but will require ICMP Echo and ARP to work across routers so that the DHCP servers and phones can determine when a duplicate IP address is already assigned.
- Splitting the range is the recommended method, this will restrict the number of duplicates since there are no common addresses between different DHCP servers. Splitting will also be advantageous since some devices will work in the event of a DHCP server failure.
- Splitting the range will likely result in different TFTP servers being specified for different sets. In order to maintain compatibility, the TFTP servers must have the same software revisions for the phone software downloads.
- A particular network could have multiple DHCP servers on the same subnet, however the same guidelines outlined above will apply.
- If a secondary DHCP server is located on a different subnet, then routers/layer3 devices must have DHCP forwarding and ICMP Echo enabled.
- In a split scenario with say 50% on one DHCP server and 50% on another, if the first DHCP server has run out of addresses it should now stay quiet when a further DHCP broadcast is sent out, i.e. all subsequent requests will be serviced by the secondary DHCP server. Additionally the first DHCP server should inform the system administrator that the IP address pool has been exhausted.
- In a scenario where the subnet is split between two DHCP servers and the phones per subnet are reduced to match the addresses on one DHCP server, then the addresses on the secondary will be held in reserve until the phones need to switch and request an address update.

IP Phone Software Load Updating and TFTP

IP phones can get loads from different TFTP servers if they can accept responses from DHCP servers.

Note:

When there are multiple TFTP servers available on the network the system administrator must ensure that all the TFTP servers will provide the same revision level of phone software. This will guarantee that the user will receive the same phone feature support regardless of which TFTP server provided the phone software.

Use of Third-party DHCP or TFTP Servers for IP Phones

These can be used. If an external DHCP server is used, it is recommended that the internal DHCP server be disabled to avoid conflicts. If an external DHCP server is used, Option 67 will need to be defined for the E2T card in the DHCP Options form.

If an external TFTP server is used, then this needs to be specified in the appropriate DHCP Options Form.

Default Gateway IP Address

This is the IP address of the network router that the phone will utilize to access other subnets.

TFTP Server IP Address

This is the IP address that the phone will access to obtain an application software load.

- In a network that has multiple DHCP servers the system administrator should provide an alternative TFTP address on the secondary DHCP server. See also 5.1.6 DHCP Servers.
- Once a phone obtains an IP address it attempts to download some running code. This code is normally present on the local ICP; however, if the ICP is in fail state, then the phone is unable to download code from the ICP. The next best thing is to give another address, for example a backup server. If the IP address is obtained from a backup DHCP server this might be given at this point. An alternative is a local TFTP server. This can be used in normal operation to help speed up registration by sharing the download task, as well as reducing the potential bandwidth block on a WAN link when a system dies.

RTC IP Addresses

This is in fact the basic requirement in order for (call control) resiliency to work! The phone needs more than one RTC IP address if it is to find an alternative ICP when the primary ICP is out of service.

- Up to 4 ICP addresses can be provided in DHCP
- The first IP address **MUST** always be the primary for that device
- Subsequent secondary addresses can change order, for example some phones in one subnet might be given ICP B then C as backup, and others given ICP C then B as backup. This will reduce network loading during initial start up and registration.

When the phone registers with an ICP, that ICP will update the phone's primary and secondary ICP/RTC addresses. Currently this is restricted to one primary and one secondary. These values can be dynamically updated depending upon which controller they are currently registered with. ICPs in a cluster are aware of other ICPs in the same cluster.

IP Phone Address Range Per Subnet

The DHCP server is responsible for assigning IP addresses to IP phones. To ensure resilient operation there needs to be more than one DHCP server so that if a DHCP server fails or is unreachable the phones can obtain IP addresses from an alternate DHCP server.

The options available to the system administrator are:

- Use overlapping IP address ranges on multiple DHCP servers, relying on Ping and ARP to resolve potential IP address conflicts.
- Provide 50% backup for IP addresses on one server, and 50% on another server. This means that a single server failure would result in losing only 50% of the devices. An additional benefit is that you can still run the maximum number of phones on a subnet.
- If overlapping is not an option and 50% backup is not an option, then more subnets will be required since each subnet will be split in numbers to allow 100% on one server and 100% backup on secondary server.

8.6.6 Controller Registry Configuration

The Controller Registry form is accessed via the ESM System Configuration Tool. This screen provides the System Administrator with a method of changing numerous system parameters that affect resilient operation.

Controller Registry Form shows the default values and provides recommendations that the system administrator may want to consider adopting depending on the particular application.

**CAUTION:**

Do not change the default settings of the Controller Registry Configuration parameters without consulting a Mitel representative.

Table 17: Controller Registry Form

Parameter	Default Value	Recommended Value
Link Maintenance - Heartbeat Rx Link Failure Time-out (1/10 seconds)	300	Default value should be used.
Link Maintenance - IP Trunk Heartbeat Rx Link Failure Time-out (1/10 seconds)	200	Default value should be used.
Link Maintenance - IP Trunk Inactivity Tx Timer (1/10 seconds)	200	Default value should be used.
Link Maintenance - Key Ack Early Warning Timer (1/10 seconds)	50	Default value should be used.
Link Maintenance - Key Ack Link Failure Time-out (1/10 seconds)	300	Default value should be used.
Link Maintenance - Phone Inactivity Tx Timer (1/10 seconds)	300	Default value should be used.
Link Maintenance - Special Key Time-out (1/10 seconds)	100	Default value should be used.

Parameter	Default Value	Recommended Value
MiTAI/ MitaiCallCtrlHeartbeatInterval	60	Default value should be used.
MiTAI/ MitaiClientHeartbeatInterval	3600	Default value should be used.
Resiliency - Allow Return to Primary ICP (Yes/No)	Yes	Default value should be used.
Resiliency - Secondary ICP Health Check Interval (seconds)	60	Default value should be used.
Resiliency - Successful Health Checks Prior to Hand Off (1 - 10)	5	Default value should be used.

8.6.7 DHCP Servers

Dynamic Host Configuration Protocol (DHCP) servers are responsible for passing configuration information to devices such as IP phones and PCs on a TCP/IP network.

The DHCP protocol consists of two components:

- A protocol for delivering device-specific configuration parameters from a DHCP server to a device.
- A mechanism for allocation of IP addresses to devices.

DHCP supports three mechanisms for IP address allocation:

- Automatic allocation assigns a permanent IP address to a device.
- Dynamic allocation assigns an IP address to a device for a limited period of time, called the lease time.
- Manual allocation allows the network administrator to assign an IP address to a device and DHCP is used to convey the assigned address to the client.

A particular network will use one or more of these mechanisms, depending on the policies of the network administrator.

Dynamic allocation allows automatic reuse of an IP address that is no longer needed by the device to which it was assigned. Dynamic allocation is particularly useful for assigning an address to a client that will be connected to the network only temporarily or for sharing a limited pool of IP addresses among a group of devices. Dynamic allocation is suitable for environments where there are frequent user moves, adds or changes, since it automates this process and reduces the chance of errors.

Manual allocation allows the DHCP server to be used to eliminate the error-prone process of manually configuring devices with IP addresses in networks where the system administrator wishes to assign IP addresses without the help of DHCP mechanisms.

MiVoice IP Phones can work with 3rd party (external) DHCP servers or the internal DHCP server provided with the ICP.

Under some circumstances such as remote office applications or Teleworker applications a DHCP server might not be available, in these situations the IP phones will need to be statically programmed with network addresses. For details, see [IP Phone Behavior During Resilient Mode Operation](#) on page 286.

Internal DHCP Server Configuration

The ESM System Administration Tool is the tool that is used to configure the Internal DHCP server. Within the System Administration Tool is a directory called DHCP, this directory contains a number of forms that are used to configure the Internal DHCP server, these forms are discussed in the following sections.

DHCP Server Form

The DHCP Server Form is used to enable or disable the Internal DHCP server. For details refer to the System Administration Tool Help.

DHCP Subnet Form

The DHCP Subnet Form is used to configure the IP address and bit mask for the DHCP Subnet. The MiVoice Business DHCP Server provides for a maximum of 10 subnets.

For details on programming this form refer to the System Administration Tool Help.

Additional information regarding the use of subnets and network design can be found via Mitel On Line, under White Papers, refer to the *3300 ICP Engineering Guidelines*.

DHCP Static IP Form

The DHCP Static IP Address Form is used to assign IP addresses to devices that require a permanent IP address. The MiVoice Business DHCP Server provides for a maximum of 50 static IP addresses. For details on programming this form refer to the System Administration Tool Help.

DHCP E2T Configuration

DHCP E2T configuration applies to programming Mitel DHCP servers when there is more than one E2T card or there are more than one DHCP server residing on the same subnet.

The 250- and 700-user 3300 Controller contain E2T cards that need to obtain information from a DHCP server. The 100-user 3300 Controller does not contain an E2T card; however, this type of ICP can act as a DHCP server. In situations where a 100-user controller is acting as an alternate DHCP server for a given subnet, it must be programmed to satisfy the requirements of any E2T cards that are residing on this subnet.

When using DHCP servers it is important to ensure that the E2T obtains its software load from the RTC that resides in the same ICP and not from an RTC that is installed in a different ICP. This is achieved by filtering the E2T DHCP request against the E2T MAC address. This address should be made permanent (static).

 Note:

When using an external DHCP server, you must define option 67 in the DHCP Options form.

The ESM System Administration Tool is used to make Static IP Address Assignments. Within the System Administration Tool the DHCP Static IP form is used to program a static IP address and other parameters for the E2T. For details on programming the internal DHCP server refer to the Technician's Handbook and the System Administration Tool Help.

When entering information into the DHCP Static IP form for the E2T, the correct MAC address to use is the MAC address of the E2T, this value can be determined by referring to a MAC Address label on the rear panel of the ICP. There are three MAC Address labels on the rear panel of the ICP, the correct label refers to "E2T MAC" "Slot 1".

When there are multiple DHCP servers that can be reached by the E2T then it is important that the same reservations and information be stored in each server. This ensures that the E2T card gets the same information independent of which server responds.

Program E2T Information for DHCP

To complete this procedure, you must use the DHCP Static IP form and the DHCP Options form. For more information about the DHCP Options form, see [DHCP Options Form](#).

To program E2T information into the MiVoice Business system’s internal DHCP server

- 1. In the **Selection** drop-down menu in the System Administration Tool, click **System Administration**, click **DHCP Static IP**, and then assign a static IP address to the E2T card. Ensure that you use the MAC address of the E2T card.

Figure 85: Programming DHCP Static IP Address for E2T Card

AddChangeChange PageChange AllDelete

PreviousPage 1 of 1NextGo to: value:

DHCP Static IP

Name	IP Address	Subnet	Client ID
7100MAP_NPL_1	10.35.12.19	P59WAN Data (10.35.12.0)	11020068
Scott's PC	10.35.29.19	P59 Data (10.35.29.0)	000874d446e5
E2t	10.35.30.12	P59 Voice (10.35.30.0)	08000f119f47

DHCP Static IP

Name:

Subnet:

IP Address:

Protocol:

7100MAP_NPL_1

P59WAN Data (10.35.12.0)

10.35.12.19

BOOTP or DHCP

Hardware Address

Type:

Other - Type:

Address:

Other - Address Length:

MAC Address

00:02:c6:00:07:fa

Client ID:

11020068

IP0755

- 2. In the **DHCP Options** form, click **Add**.

Figure 86: DHCP Options View

P59 - Mitel Networks® 3300 Integrated Communications Platform (ICP) - Microsoft Internet Explorer

FileEditViewFavoritesToolsHelp

Alarm Status: No Alarm 2003-Jun-11 10:21:58Data RefreshHelpExit

P59 - System Message:

Selection: System Administration

System Administration

User Authorization Profiles

System Options

Automatic Route Selection (AR)

Automatic Call Distribution (AC)

Call Handling

Telephone Management

Telephone Directory Management

Property Management System

Voice Mail

DHCP

DHCP Server

DHCP Subnet

DHCP Static IP

DHCP IP Address Range

DHCP Options

AddChangeDelete

PreviousPage 1 of 1NextGo to: value:Go

DHCP Options

ID	Name	Format	Value	Scope
3	Router	IP Address	10.35.12.1	Subnet: P59WAN Data (10.35.0)
3	Router	IP Address	10.35.13.1	Subnet: P59WAN Voice (10.35.0)
3	Router	IP Address	10.35.29.1	Subnet: P59 Data (10.35.029.0)
3	Router	IP Address	10.35.30.1	Subnet: P59 Voice (10.35.030.0)
6	DNS Server	IP Address	134.198.27.52	Global
44	NetBIOS Name Server	IP Address	134.198.84.34	Global
66	TFTP Server Name	ASCII String	10.35.30.11	Static: E2t (10.35.030.012)
67	Boot File Name	ASCII String	sysone218269	Static: E2t (10.35.030.012)
120	User Defined	IP Address	10.35.30.11	Global
129	User Defined	IP Address	10.35.30.11, 10.35.6.11, 10.35.6.116	Global
130	User Defined	ASCII String	MITEL IP PHONE	Global
131	User Defined	IP Address	10.35.5.19	Global
132	User Defined	Numeric	430	Global
133	User Defined	Numeric	6	Global

MITEL 3300 ICP

IP0756

3. If you want to configure multiple resilient ICPs on the same subnet, complete steps 4 to 8 of this procedure to program options 66 and 67.

If you do not want to configure multiple resilient ICPs on the same subnet, do not complete steps 4 to 8. In this case, you do not need to program options 66 and 67.

4. In the **DHCP Options** view, enter the following information to add option 66 (TFTP Server Name):

- ID: **66**
- Format: **ASCII String**
- Value: <Type the IP address of the RTC card IP address without any leading zeros. For example, instead of "010.012.001.003", type "10.12.1.3">
- Scope: <Select the Static scope for the E2T>

Figure 87: DHCP Options Form (Programming Option 66)

-- Web Page Dialog

DHCP Options

ID: 66

Name:

Format: ASCII String

Value: 10.35.30.11

Scope: Global

Options Help

- Global
- Subnet: P59WAN Data (010.035.012.000)
- Subnet: P59WAN Voice (010.035.013.000)
- Subnet: P59 Data (010.035.029.000)
- Subnet: P59 Voice (010.035.030.000)
- Static: CORP-EMANAGER (010.035.012.018)
- Static: 7100MAP (010.035.012.019)
- Static: Scott's PC (010.035.029.019)
- Static: E2t (010.035.030.012)**
- Range: P59WAN Data (010.035.012.100-010.035.012.200)
- Range: P59WAN Voice (010.035.013.100-010.035.013.200)

on parameter.

Save Cancel

IP0757

5. Click **Save**.

6. In the **DHCP Options** view, click **Add**.

7. In the **DHCP Options** form, enter the following information to add option 67 (Boot File Name):

- ID: **67**
- Format: **ASCII String**
- Value: **/sysro/e2t8260**
- Scope: <Select the Static scope of the E2T>

Figure 88: DHCP Options Form (Programming Option 67)

The screenshot shows a 'Web Page Dialog' window titled 'DHCP Options'. It contains the following fields and values:

- ID:** 67
- Name:** (empty)
- Format:** ASCII String
- Value:** /sysro/e2t8260
- Scope:** Global

Below these fields is a list of DHCP options. The option 'Static: E2t (010.035.030.012)' is selected. Other options in the list include:

- Global
- Subnet: P59WAN Data (010.035.012.000)
- Subnet: P59WAN Voice (010.035.013.000)
- Subnet: P59 Data (010.035.029.000)
- Subnet: P59 Voice (010.035.030.000)
- Static: CORP-EMANAGER (010.035.012.018)
- Static: 7100MAP (010.035.012.019)
- Static: Scott's PC (010.035.029.019)
- Static: E2t (010.035.030.012)
- Range: P59WAN Data (010.035.012.100-010.035.012.200)
- Range: P59WAN Voice (010.035.013.100-010.035.013.200)

At the bottom right, there are 'Save' and 'Cancel' buttons. A status bar at the bottom left shows 'IP0758'.

8. Click **Save**.

DHCP IP Address Range Form

The DHCP IP Address Range form is used to program the range of DHCP addresses. This form is also used to program the lease time and to select the protocol to be used for an address range. For details on programming the DHCP IP Address Range Form refer to the *System Administration Tool Help*.

DHCP Options Form

The DHCP Options Form allows for the creation of a DHCP option and specifies where the option should be applied. For details on programming the DHCP Options Form refer to the *System Administration Tool Help*.

For information on how to program options 66 and 67, see [Program E2T Information for DHCP](#).

Note:

The RTC entry (Option 129), is able to accept up to 4 IP addresses which the user enters into the **Value** field of the DHCP Options Form. The correct syntax for entering multiple IP addresses is the first IP address followed by a comma followed by a space which is followed by the next IP address. For example: xxx.xxx.xxx.xxx, xxx.xxx.xxx.xxx, xxx.xxx.xxx.xxx, xxx.xxx.xxx.xxx

Available DHCP Options

Since the DHCP server options can differ from release to release, consult the MiVoice Business *System Administration Tool Help*. The following table shows the options currently used.

Table 18: DHCP server options

DHCP option	Information
003 – router address	IP address (e.g. 192.167.22.251)
066 – FTP server for gateway/E2T	IP address (e.g. 192.167.22.10)
067 – boot file name for gateway	See the <i>Technician's Handbook</i> .
128 – (specific) TFTP server	IP address (e.g. 192.167.22.10)
129 – (specific) RTC	IP address (e.g. 192.167.22.10)
130 – (specific) IP identifier	"MITEL IP PHONE"

DHCP option	Information
131 – debug	IP address (e.g. 192.167.22.100)
132 – (specific) VLAN ID (32 bit)	0x2
133 – (specific) Priority (32 bit)	0x6
134 – (specific) DSCP value 0 to 63 (32 bit)	0x2e (this is 46 decimal in hexadecimal format)

Pay special attention to the configuration of option 54 in external DHCP servers. Incorrect configuration can misdirect responses and affect startup of IP devices.

Options 66 and 67 are used when an external DHCP server boots the internal gateway on the 3300 ICP LX (700-user) and 250-user platforms. Be careful using these options in DHCP as certain workstations (without built in operating systems) also use these options to boot. Make sure that these options are visible only on the voice VLAN to which the MiVoice Business system is connected. Data devices should be on a separate VLAN with a separate DHCP server, or 'scope' setting.

The option 130 field identifies the group of options from 128 to 134 as being valid for IP phones. If any of these options is used for different purposes, other than stated above, then the phones will read this information and may give unpredictable results. Unused fields should be left blank

Option 131 is the IP address of a PC that is running the IP Analyzer software. Use this software to check normal operation and to obtain device and call statistics.

The VLAN (132) and Priority (133) option are present (or not) depending upon whether VLAN is used within the network. If not present, the IP phone does go through the second sequence of determining an IP address at startup and continues to operate in an untagged mode.

The DSCP option (134) will allow the Layer 3 priority value to be defined. This is entered in hexadecimal format (base 16) and can accept decimal values from 0 to 63, or 0x00 to 0x3F. Multiplying this value by 4 results in the effective value in the TOS field.

Note that the use of DHCP options 128 to 224 may no longer be unique to Mitel. Conflict may arise where a number of different devices exist within the same subnet, or DHCP scope. It may be necessary to redefine options, or place some equipment in different scopes, or select options based on device MAC address. It is known that Microsoft Server 2003 also uses options 132 and 133.

DHCP Lease Viewer Form

The DHCP Lease Viewer form displays all programmed IP addresses and the lease time for these addresses. This form can also be used to delete the lease of an IP address. For details regarding this form, refer to the *System Administration Tool Help*.

Known Condition

The following known condition exists: If the phone's original ICP has permanently failed, or there is a permanent link failure to it, the phone's designated DHCP service is probably gone as well (usually the ICP). The phone's IP address lease will eventually expire. Unless a network action has been taken to reconfigure a reachable DHCP for the phone, to allow renewing its lease to the same IP address on the different DHCP server, the phone will reset.

One solution would be to make the IP address lease times very long (order of many hours to days) to help with this case. However, this may not help in a situation where there are frequent Moves/Add/Changes to end devices.

The following guidelines will highlight other solutions:

- The lease time can be extended in the hope that the failed ICP or link will recover in time
- There needs to be access to multiple DHCP servers, be they two standalone servers (e.g. PC), servers within different ICPs, or a combination of the two. To achieve this DHCP forwarding in the routers must be supported and enabled. The DHCP servers will need to verify address selection as per RFC2131(Echo/Echo Reply and ARP), otherwise address conflicts will ensue. See also the section on configuration of routers.
- The information pertinent to a particular subnet must be copied equally into both DHCP servers, otherwise there is likely to be conflict and potentially undetermined operation. The one option that might change is the selection of the TFTP server where a secondary server is specified, for example, if the primary ICP has failed, then the TFTP server local to that device may also have failed. Choosing a backup will help with initialization of devices.
- Sufficient IP addresses must be available on the secondary DHCP server, equal to the number on the primary DHCP server, even if some of the phones never register with the backup ICP, otherwise the first few devices will empty the pool of addresses (distributed on a first come, first served basis).
- Consider conflict resolution. The phones currently take the first DHCP response to arrive and work with that. If the response came from a remote DHCP server first, then that will be used. Address conflict and repetition will be dealt with according to RFC2131.
- Consider conflict resolution between DHCP servers. For example, if ICP1 fails and loses information on used IP addresses, and a phone renews its lease on a secondary DHCP server, and the first server returns, how will this server know not to release the

IP address already assigned. Phones use ARP and DHCP servers use ICMP Echo. All being well there should not be a conflict.

- The DHCP servers work independently with little knowledge of each other, unless they use a common database held in a third location. If the lease renewal fails at the original DHCP server the end device will continue to operate. The secondary DHCP server will have this address marked as already in use. Later (at 87% lease time) the end device will send out a broadcast attempting to renew the same IP address. If this fails, it will attempt to get any address. At this point the secondary DHCP server will issue another IP address and renew the lease time. The original address will be marked as 'in-use' in the secondary DHCP server for a period of time, in the in th this is currently 3 days, in an external server this may be for the lease duration. Once this time expires, this address will become available. When the primary DHCP recovers, it will recognize when the lease time has expired on the original IP address and make it available again.

RTC IP Address Assignment

The RTC does not obtain an IP address or gateway IP address from the DHCP server. This information is programmed into the RTC via the VxWorks prompt. For details refer to *3300 IP Communications Platform Technician's Handbook*.



Note:

This address should be omitted from the list of available addresses in the DHCP servers.

DHCP support for Multiple E2Ts

The DHCP server can handle DHCP requests from multiple E2T devices through either subnet and scope assignment as well as MAC address filtering.

DHCP support for Multiple TFTP Servers

Since the phones will only accept a single TFTP server address a single DHCP server cannot provide phones with multiple TFTP Server IP addresses.

However, if a network makes use of multiple DHCP servers, each DHCP server can, if desired provide unique TFTP Server IP addresses. If the system administrator wishes to use this strategy the various TFTP servers must be loaded with the same revision of phone application software so that phone users obtain the same feature support no matter which TFTP server provided the phone application software.

DHCP Support for Multiple Routers

Phones and E2T cards can only accept one router IP address.

External DHCP Server Configuration

Much of the configuration requirements for the DHCP server internal to the ICP apply to an external DHCP server. The key differences are the requirements to program Options 66 and 67.

If an external DHCP server is being used, it will need to be programmed with the same parameters that are programmed into an internal DHCP server. For details on programming an external DHCP server, the administrator should refer to the documentation provided with the 3rd party DHCP server.

Some 3rd party vendors of DHCP servers have implemented proprietary protocols or mechanisms that allow their primary and backup DHCP servers to exchange relevant DHCP information with each other.

If a particular network employs these types of DHCP servers the system administrator should be aware that there might be alternative methods to provision IP address lease times and IP address ranges across multiple DHCP servers then has been described in previous sections of this document, the system administrator should refer to the 3rd party documentation for implementation details.

Using a Secondary ICP as a Backup DHCP Server

A secondary ICP could, if desired, serve as a backup DHCP server. As mentioned previously in the section, "[DHCP Servers](#)", it is important that the primary and secondary DHCP servers provide network devices with the same DHCP information; otherwise, device behavior could become unpredictable.

Note:

If a second ICP is used to provide backup DHCP server functionality, the system administrator should ensure that the backup DHCP server does not include information for the remote E2T card. This ensures that if the remote E2T is alive but the local RTC is not, the remote E2T will not try to operate.

8.6.8 Connection Scenario for Routers

The routers will play a key part in the operation of resiliency. Typically routers will be used to isolate broadcast areas and isolate ICPs from one another. With resiliency these broadcast domains need to be opened in a controlled manner, so that ICPs can provide backup to one another.

Physical Connection

Ideally this is a 100BaseT FULL DUPLEX connection back into the Core part of the network. With virtual port configuration, this will be a single connection capable of handling VLANs.

In the case where virtual-ports are not available, then an individual connection per subnet and/or VLAN will be required. Again, the fastest connection possible is desirable.

Router Requirements and Guidelines

The system administrator needs need to ensure that the network routers have the ability to:

- Allow DHCP broadcasts/forwarding between defined subnets
- Allow ICMP Echo and Echo Reply to transition between sub-nets
- Handle multiple virtual connections on a single port, i.e. virtual port and handle multiple VLAN (A router without this can still be used but this will require a router with a physical port per subnet and/or VLAN. In a remote location with minimum VLANs this may not be a big issue.)
- Queue traffic based on TOS/Precedence or Diffserve
- Provide TOS/Diffserve to COS (Layer 2 802.1p) mapping
- Provide MTU shaping (needed potentially on lower speed WAN links)
- Provide ICMP Redirect message pending failure of a link

The following guidelines also pertain to routers:

- Where multiple physical paths are possible between routers, switching within this WAN connection should be transparent to the Layer2 LAN connection.
- Where a secondary DHCP server is on the far side of a router, then the information in the appropriate subnet, or scope, must be identical to that entered into the main primary DHCP server, otherwise excessive hunting, or even failure to find the correct ICP could ensue.
- Where DHCP forwarding is used, then ICMP Echo and ICMP Echo Reply must also be available. This is the main mechanism used by DHCP servers to determine if an offered IP address is already in use, and if it's on the other side of a router, it's difficult to see that unless ICMP Echo/Echo Reply is available.

- The system administrator needs to ensure that network routers are aware of each other so that in the event of a link or router failure the surviving routers can take corrective action.

It should be noted that a failed router does not issue an ICMP redirect to any hosts on the LAN (such as ICPs and the connected IP phones). However, if configured to do so, a router that has detected an inactive link will issue an ICMP redirect command to the connected devices, which cause the ICPs or IP phones to be redirected to a good router provided the routers are "aware" of the pathway.

Devices that do not support IRDP will be unable to locate an alternate router using this protocol.

In the event that the primary router completely fails the router will be unable to issue an ICMP redirect to end devices.

However, if configured to do so, a router that has detected an inactive link will issue an ICMP redirect command to the end device, redirecting it to a known good router providing the routers are aware of each other.

The following points should be considered:

- Routers can be made aware of each other with the use protocols such as GRP and RIP, however these protocols have slow convergence times, i.e. in the order of 3 to 10 minutes.
- The use of routing protocols such as the Enhanced Interior Gateway Routing Protocol (Enhanced IGRP) and Open Shortest Path First (OSPF) would be a better choice since these protocols have faster convergence times, i.e. in the order of seconds.
- Some router vendors may provide proprietary protocols/mechanisms to keep communication paths open for devices that do not support IRDP when a router becomes unavailable. An example of such a protocol is Cisco's Hot Standby Routing Protocol (HSRP).
- If a particular network employs Cisco routers that support HSRP the system administrator may wish to make use of capabilities provided by HSRP. HSRP allows multiple HSRP-configured routers to use the MAC address and IP network address of a virtual router. The virtual router does not physically exist; instead, it represents the common target for routers that are configured to provide backup to each other. Each physical router is configured with the MAC address and the IP network address of the virtual router. IP devices would be configured with the IP address of the virtual router as their default router.

CXi and Mx Integral Router

Note:

The CXi and Mx controllers provide an integral router. For more details on configuring and using the routing capabilities of these controllers, refer to the MiVoice Business *System Administration Tool Help* or the 3300 ICP Engineering Guidelines.

The CXi and Mx controllers support the router portion of the ICMP Router Discovery Protocol (IRDP) as described in RFC-1256. When IRDP is enabled on these controllers, they transmit router advertisements through the IP network (i.e. LAN). These router advertisements announce to any hosts on the LAN that the controller is a router and provide them with its IP address.

The CXi and Mx transmit “Router Advertisements” on a periodic basis. They also transmit a “Router Advertisement” if a “Router Solicitation” is received from a host.

The CXi and Mx use the following IRDP parameters:

- Advertising Rate: At startup 3 times 16 seconds apart, then every 7.5 minutes to 10 minutes
- Preference Level: 0
- Lifetime Field: 30 Minutes

Note:

1. Only hosts that support IRDP can make use of Router Advertisements, at the present time MiVoice IP Phones are not IRDP capable.
2. If there is an external router present on the LAN, then Router Discovery should be disabled on the CXi.

Avoiding DHCP Broadcast Loops

⚠ CAUTION:

Avoid DHCP broadcast loops. When enabling DHCP forwarding on routers, ensure that a loop does not exist. In the Layer3 environment there is no protection mechanism to prevent broadcast loops, such as in the Layer2 protection offered by the Spanning Tree Protocol.

What is a DHCP Broadcast Loop?

Layer 3 devices are used to prevent broadcasts flooding the network, but in order to pass DHCP broadcasts, this protection is bypassed.

The following information pertains to using protocol-based mechanisms to prevent DHCP Loops.

Proprietary Broadcast to Unicast Converters

Some vendors provide intelligent, protocol-aware features in their Layer 3 switches and routers. These features contain broadcasts such as DHCP broadcasts by converting them into directed unicasts.

i Note:

If you choose to use such a protocol to prevent DHCP broadcast loops, the protocol must be enabled throughout the network, on all Layer 3 devices. For further information refer to the vendor's documentation.

Routing Information Protocol

The Routing Information Protocol (RIP) is a distance-vector intra-domain routing protocol. RIP works well in small, homogeneous networks; however, in larger, more complex internetworks, it has many limitations, such as a maximum hop count of 15, lack of support for variable-length subnet masks (VLSMs), inefficient use of bandwidth, and slow convergence.

By implementing a limit on the number of hops allowed in a path from the source to a destination, RIP prevents routing loops from continuing indefinitely. The maximum number of hops in a path is 15. If a router receives a routing update that contains a new or changed entry, and if increasing the metric value by 1 causes the metric to be infinity

(that is, 16), the network destination is considered unreachable. The downside of this stability feature is that it limits the maximum diameter of a RIP network to less than 16 hops.

If the RIP is employed as a method of limiting the duration of DHCP broadcast loops, you must ensure that RIP is enabled throughout the network on all Layer 3 devices. For further information refer to the vendor's documentation.

Using Protocols to Prevent DHCP Broadcast Loops

Some proprietary and open standard routing protocols could be used to prevent DHCP broadcast loops from occurring, but in complex networks the correct configuration of these protocols may not be intuitive and may not be understood until a loop actually occurs.

These protocols are not recommended and should be used as safety nets rather than the prime method of preventing DHCP broadcast loops.

A properly designed network uses topographic design rather than routing protocols as the prime method (recommended) to prevent DHCP broadcast loops.

Using Structured Network Design to Prevent DHCP Broadcast Loops

Where a single gateway connects to multiple WAN links, there is less likelihood of a DHCP broadcast loop being created.

The DHCP broadcast loop condition is more likely to occur where there are multiple independent routers connecting different subnets from a single subnet. In this situation it is difficult to limit broadcast paths.

The following network topology example in [Figure 89: DHCP Broadcast Loop \(To Avoid\)](#) on page 330 might allow DHCP broadcast loops to occur:


- System A is backed by System B.
- System B is backed by System C.
- System C is backed by System A.

A DHCP broadcast in system A will result in a broadcast into System B, and then on to System C, and finally back to System A.

CAUTION:

Avoid the type of topology illustrated in [Figure 89: DHCP Broadcast Loop \(To Avoid\)](#) on page 330 . This topology would allow a DHCP Broadcast Loop to occur.


Figure 89: DHCP Broadcast Loop (To Avoid)



IP0741

The following figure illustrates a topology that will prevent DHCP broadcast loops but provides the connectivity that was required in the previous figure, this type of topology is recommended.

Figure 90: How to Avoid a DHCP Broadcast Loop



Verifying that Routers Support ICMP Echo Relay

It is recommended that the System Administrator verify that routers within the network support ICMP Echo Relay.

This functionality is verified by pinging an IP address on the far side of a router to verify that the router relays the ping request. This is simply to make sure that ICMP Echo and Echo Return are allowed to pass through the particular router ports as many router devices have this disabled as default

Routers must support ICMP Echo Relay so that remote DHCP servers will be able to identify if IP address is already given out and in use.

RFC 792 ICMP Redirect

The System Administrator must enable the Internet Control Message Protocol on all network routers.

ICMP (Internet Control Message Protocol) is a protocol that is used by gateways to communicate with end devices. The redirect message is one of a number of messages that can be sent from a gateway to an end device under ICMP.

The gateway sends a redirect message to an end device in the following situation:

- A gateway, G1, receives an internet datagram from an end device on a network to which the gateway is attached. The gateway, G1, checks its routing table and obtains the address of the next gateway, G2, on the route to the datagram's internet destination network, X. If G2 and the end device identified by the internet source address of the datagram are on the same network, a redirect message is sent to the end device. The redirect message advises the end device to send its traffic for network X directly to gateway G2 as this is a shorter path to the destination. The gateway forwards the original datagram's data to its internet destination.

Note:

1. ICMP redirects only address the scenario where the link on the “far side” of a router is broken. An inactive router cannot issue an ICMP redirect.
2. When the E2T and the RTC are operating under an ICMP redirect they will attempt to make connections with their respective destination addresses via the initial default gateway (the frequency is based on a time-out period) to determine the status of this link. If the link is still broken, then a further ICMP redirect will ensue and the timer will restart.
3. When operating under an ICMP redirect the IP phones do not attempt to determine the status of the link at the initial default gateway, the phones will remain locked with the gateway that they have been re-directed to.

ARP, Proxy ARP and Phone Mobility

There are two versions of the Address Resolution Protocol (ARP) in existence, RFC 826, the Address Resolution Protocol and RFC 925, the Multi-LAN Address Resolution Protocol, also known as Proxy ARP. These protocols are used by routers and hosts to map IP addresses to hardware MAC addresses.

Most router vendors will support both protocols and will provide a mechanism through the user interface to enable or disable these protocols.

Additionally the user interface should provide a mechanism that allows the system administrator to set the aging or time out period for the ARP cache. The ARP cache is a table maintained in memory on the router, this table is used to map or correlate IP addresses to hardware MAC addresses.

If the aging or time out period is too frequent the network will possibly become congested with ARP traffic as routers attempt to learn addressing information so that they can

update their ARP caches. If the aging period is too infrequent, there will be the possibility that the ARP cache will not reflect recent changes in the network.

One particular network change requires special attention. This is the situation where a phone is physically moved to another subnet. ARP and Proxy ARP will behave differently under these circumstances, this behavior is described below.

When a Phone Moves on a Network that is Running ARP

When a phone physically moves to a different subnet a new IP address will be issued to the phone by the DHCP server.

Connections initiated by this phone will be successful. However, connections initiated to this phone will not initially work since the router's ARP cache will still contain the phone's previous IP address. The ARP cache will not get updated to reflect the phone's new IP address until the ARP cache table ages out or times out, hence the importance of setting the ARP cache aging or time out period to a suitable value.

Note that some vendors may have implemented additional methods of minimizing the amount of time that incorrect addresses can exist within their ARP cache, in these cases fine tuning of the ARP cache time out period will not be as critical.

When a Phone Moves on a Network that is Running Proxy ARP

Proxy ARP can be run under two different addressing schemes, the explicit subnet scheme and the extended ARP scheme.

- Explicit subnet scheme
 - In the explicit subnet scheme, some IP address bits are devoted to identifying the subnet (i.e., the LAN). The address is broken up into network, subnet, and host fields.
 - In the explicit subnet scheme, when a phone is unplugged from one LAN and plugged into another LAN its IP address must change. This behavior is exactly the same as a phone that is moved on a network running ARP, and potentially the same requirement for fine tuning the ARP cache time out period exists.
- Extended subnet scheme
 - In the extended ARP scheme, the address fields employed are the network and host fields. The extended ARP scheme may be used with any class of IP address.
 - In the extended ARP scheme, when a phone is unplugged from one LAN and plugged into another LAN, it can be reissued the same IP address. In this case, the ARP cache remains in sync with the network change since from an addressing perspective nothing in the network actually changed.

What is an Appropriate ARP Cache Time-out Period?

Fifteen minutes is the recommended time interval that should be used for timing out ARP Cache entries. This gives a user enough time to unplug a phone and move the phone to another desk top location.

The ICP maintains an ARP Cache, this ARP Cache has a time out period.

When Should Proxy ARP be Used?

The system administrator will need to make this decision, the following should be considered.

Advantages of Proxy ARP

The main advantage of using proxy ARP is that it can be added to a single router on a network without disturbing the routing tables of the other routers on the network.

Proxy ARP should be used on the network where IP hosts are not configured with default gateway addresses or they do not have any routing intelligence.

Note:

MiVoice IP Phones and devices are configured with default gateway addresses.

Disadvantages of Proxy ARP

Hosts are not aware of the physical details of their network and assume it to be a flat network in which they can reach any destination simply by sending an ARP request; however, using ARP for everything has the following disadvantages. Note that this list is not comprehensive:

- It increases the amount of ARP traffic on your segment.
- Hosts need larger ARP tables to handle IP-to-MAC address mappings.
- Security may be undermined. A machine can claim to be another in order to intercept packets, an act called “spoofing”.
- It does not work for networks that do not use ARP for address resolution.
- It does not generalize to all network topologies (for example, more than one router connecting two physical networks).

Local Area Mobility

Local area mobility is a proprietary feature, offered by some router vendors, that provides the ability to relocate IP hosts or phones within a limited area without reassigning host IP addresses and without changes to the host software.

To create larger mobility areas, the system administrator must first redistribute the mobile routes into the Interior Gateway Protocol (IGP). The IGP must support host routes.

The mobile area must consist of a contiguous set of subnets.

Hosts or phones that roam within a mobile area should rely on a configured default router for their routing.

For implementation details refer to the vendor's documentation.

Using Local Area Mobility features has the following limitations:

- All mobile devices utilize the same router, this becomes a single point of failure in the network.
- Since all mobile devices must use the same router, this router must provide connections to all of the subnets that are supporting mobile devices and must, therefore, have the capacity to handle the concentration of traffic. As a result, there are potential bandwidth-limitation issues with the interfaces to this router.

8.6.9 Spanning Tree Protocol STP and Rapid Spanning Tree Protocol RSTP

As of 3300 Release 6.0 the Rapid Reconfiguration of Spanning Tree Protocol (RSTP) is supported on the CXi Controller. It is also supported on the MXe controller which was introduced in 3300 Release 7.0.

As of Release 5.0 the Spanning Tree Protocol (STP) portion of IEEE 802.1D MAC Bridges is supported on the LX Controller, 700-user ICP (if upgraded to Release 5.0), and 250-user ICP (if upgraded to Release 5.0). For information on how to enable Layer 2 STP, refer to the MiVoice Business *System Administration Tool Help*.


STP is not supported on pre-release 5.0 ICPs, the 200-user ICP, or the 100-user ICP. You must exercise care when integrating an ICP that does not support STP or has STP disabled, into a network that is running STP. See [Connecting MiVoice Business Systems to L2 Switches](#) on page 302. Also refer to the *3300 ICP Engineering Guidelines*.

About STP RSTP

In an Ethernet network that is not using STP/RSTP, multiple active paths between devices are not allowed since multiple paths will cause network loops. See the following Figure.

Network loops are unacceptable because a broadcast or multicast packet sent from Station “A” to Station “B” will be forwarded by Switch “B” to Station “B” and also back to Switch “A”, when Switch “A” receives the packet it will then forward the packet back to Switch “B” and the cycle will repeat for infinity causing a broadcast storm.

Figure 91: Network Loop




The Rapid Reconfiguration of Spanning Tree Protocol (RSTP) is also a Layer 2 Link Level protocol specified by the IEEE (802.1w) that runs on bridges and switches.

STP and RSTP serve the same purpose. The difference between the two protocols has to do with how quickly the algorithms can converge on a network. RSTP "reconverges" networks faster than STP.

The guidelines in this section are applicable to both STP and RSTP; any differences between the two protocols will be highlighted. STP allows for physical path redundancy by placing the redundant paths in standby mode. It does this by blocking traffic on redundant ports.

Should a currently active network path fail due to a Bridge/Switch failure or a network cabling failure, STP will enable the network path that was previously held in a standby mode and network connectivity will be restored. The following Figure depicts how STP breaks a potential network loop by blocking traffic on one of the ports on Switch “B”.

Figure 92: Network Loop Broken by STP



STP, Network Topology, and Terminology

- **Bridged LAN versus Switched LAN**

Bridges and switches are Layer 2 devices that are used to forward packets between different network segments.

Switches offer better data throughput and higher port density than bridges. As a result, Switches have displaced bridges as the preferred internetworking solution.

A bridged LAN is composed of two or more LANs that are interconnected with bridges. In a bridged LAN that is running STP, the logical centre of the network is called the root bridge.

A switched LAN is composed of two or more LANs that are interconnected with switches. In a switched LAN that is running STP, the logical centre of the network is called the root switch.

This document uses the terms bridge and switch interchangeably.

- **Port States**

STP places a port into one of the states listed below. The state of each port dictates how the port handles received packets and whether or not it will forward packets. A special packet called a Bridge Protocol Data Unit (BPDU) is passed between bridges and switches to communicate information about the bridge/switch to other bridges/switches in the network. STP relies on the information carried in these BPDUs to learn about the network topology.

- **Blocking**—A port in the blocking state does not perform frame forwarding. As a result, a port in this state prevents packet duplication by blocking transmission on a duplicate link. The blocking state is used to prevent network loops. Received frames will be discarded and frames will not be transmitted out of this port. BPDUs received will be processed but BPDUs are not transmitted on a port in the blocking state.
- **Listening/Learning**—A port in the listening/learning state is preparing to participate in frame forwarding. Frame forwarding is temporarily disabled to prevent network loops. Received frames will be discarded and frames will not be transmitted out of this port. BPDUs received will be processed and BPDUs will be transmitted on a port in the listening state.
- **Forwarding**—A port in this state is participating in frame forwarding. Received frames can be forwarded and forwarded frames can be submitted for transmission. BPDUs received will be processed and BPDUs will be transmitted on a port in the forwarding state.
- **Disabled**—A port in this state does not participate in frame forwarding, nor does it participate in the Spanning Tree Protocol. Received frames will be discarded and frames will not be transmitted out of this port. BPDUs received will not be processed and BPDUs will not be transmitted on a port in the disabled state.

- **Root Switch**

Within a STP enabled LAN one Switch is elected to be the Root Switch. The Root Switch becomes the logical centre of the network. All ports on the Root Switch become Designated Ports. All decisions made by STP, such as which ports to block and which ports to put in forwarding mode are made with respect to the Root Switch.

- **Root Port**

All Switches (except the Root Switch) in the network must select one of their ports to become the Root Port. The Root Port is a port that leads back to the Root Switch and has the lowest path cost. Path cost for all ports in the network are determined by STP, the value of a particular path cost is based on the interface speed and how far away from the Root Switch this port is.

- **Designated Port**

Any Switch ports in the network that are responsible for connecting a LAN segment will become Designated Ports. Designated Ports are responsible for forwarding packets on behalf of that particular LAN segment.

- **Bridge Protocol Data Unit (BPDUs)**

STP defines a specialized packet called a Bridge Protocol Data Unit (BPDU). All Switches in the network transmit BPDUs to their neighboring Switches.

A BPDU contains parameters specific to the Switch that generated the BPDU, such as Bridge Priority, Path Cost, Port Priority and various STP timer values.

In Release 5.0 of the 3300 these parameters are set to default values and are not user configurable. (See [3300 CXi Release 6 0 and MXe Release 7 0 RSTP Default Parameters](#)).


STP uses the information contained in the BPDUs to create a mental picture of the network, elect a Root Switch, select Root Ports and select Designated Ports. When this process has completed the ports that are not Root Ports or Designated Ports

will be placed into a Blocked state. STP then will start the final stage of convergence whereby the Root and Designated Ports are prepared to enter the Forwarding state.

Converged STP Network, depicts an STP network that has converged or stabilized. Using this Figure as a reference it can be seen that

- L2 Switch 1 has been elected the Root Switch and all of it's ports have become Designated Ports
- L2 Switch 2 has selected it's port connecting to LAN B to be its Root Port (since this port has the lowest Path Cost to the Root Switch) and its port connecting to LAN C to be its designated port.
- L2 Switch 3 has selected its port connecting to LAN B to be its Root Port (since this port has the lowest Path Cost to the Root Switch) and its port connecting to LAN C is put into a Blocking state to prevent a network loop from being formed.
- L2 Switch 4 has selected its port connecting to LAN C to be its Root Port (since this port has the lowest Path Cost to the Root Switch) and its port connecting to LAN D to be its designated port.

Figure 93: Converged STP Network



Note:

A Designated port is the side of the switch that is opposite to the Root port, the designated port can provide connectivity to a LAN or a single network element.

STP Network Design Guidelines

This section is intended as a design guide for the System Administrator or Installer.

CAUTION:

Enabling STP in a live network will cause service disruptions to the end users while the network is converging. To avoid impacting users, enabling of STP should be conducted outside of core hours or during a scheduled maintenance period.

- Create an accurate diagram of the network. This diagram should include interface speed information.
- The System Administrator should establish which switch will be the Root Switch (the Root Switch should be a powerful switch).
- The MiVoice Business system should not be used as the Root Switch. To safeguard against this occurring MiVoice Business is factory programmed with the highest possible value for Bridge Priority.
- The Root Switch should be located in a position that will minimize the average distance between the Root Switch and all other network elements. Typically the Root Switch should be in the core network. Note that STP has a 7 hop limit, see [Topology Restrictions](#).
- If possible, high usage servers and routers should be directly connected to the Root Switch.
- Once you have decided which Switch should be the Root Switch, reduce the value of the Bridge Priority on this Switch to the lowest possible value. Since a lower value equates to a higher priority this ensures that STP recognizes this switch as the Root Switch.

- Understand where the redundant links are located and where blocking might occur. If the Root Switch is optimally located and optimally connected in the network, tuning of individual port costs to control which ports get blocked should not be necessary.
- Be aware that under some circumstances due to poor location of the Root Switch, poor interconnection of Switches to the Root Switch or non standard values for port costs, STP might utilize a non-optimum link and block the optimum link. For example, given a choice between a 10 Mbp/s link and a 100 Mbp/s, STP should choose the 100 Mbp/s link for making the active connection and the 10 Mbp/s link should be put into a blocked state (see [Efficient Usage of Inter-Switch Connections](#)).
- Try to keep the number of blocked ports in the network to a minimum. A blocked port is all that prevents a network loop, by minimizing the number of blocked ports in the network there is less risk of network problems due to a blocked port being erroneously moved into the forwarding state.
- It is recommended that you use the factory default STP parameters for your L2 switches. Only the following parameters might require changing:
 - Bridge Priority, used to select the Root Switch.
 - Port Cost, used to select link redundancy or control traffic load balancing and bandwidth optimization.

Efficient Usage of Inter-Switch Connections

As mentioned in the previous section the System Administrator should not allow STP to blindly choose which network paths will be active, instead the System Administrator should have a good idea of what connections will provide optimum bandwidth usage so that he can guide STP through appropriate network design to provide optimum connectivity.

This section provides an example of one scenario where STP does exactly what it is supposed to do, however, if the System Administrator had configured the network differently a more efficient usage of available connections would have been selected by STP.


In the following diagram there is a 1G Ethernet connection available that could have been used to connect L2 Switch A to L2 Switch B, however it was not used.

When STP converged the network the following decisions were made:

- L2 Switch A was elected as the Root Switch and placed both Port 1 and Port 2 into a Forwarding state.
- L2 Switch B placed both Port 1 and Port 2 into a Forwarding state.
- 3300 ICP B placed Port 1 into a Blocked state to prevent a network loop and placed Port 2 into a Forwarding state to provide connectivity to L2 Switch B.
- 3300 ICP A placed both of it's Ports into a Forwarding state to provide connectivity to L2 Switch A and L2 Switch B.

STP did what it was supposed to do, it removed any network loops and provided connectivity to all devices. But it should be noted that all of the traffic intended to move between L2 Switch A and L2 Switch B now has to flow through 3300 ICP A. This is not a problem to the ICP but all of this traffic is forced into flowing over a 100M Full Duplex link when the 1G Full Duplex link available on the L2 Switches could have been utilized.

Figure 94: Inefficient Use of Connections




The following diagram depicts how STP would have converged the network had the 1G Full Duplex link been available.

When STP converged the network the following decisions would have been made:

- L2 Switch A was elected as the Root Switch and placed all of it's Ports into a Forwarding state.
- L2 Switch B placed all of it's Ports into a Forwarding state.
- 3300 ICP B placed Port 1 into a Blocked state to prevent a network loop and placed Port 2 into a Forwarding state to provide connectivity to L2 Switch B.
- 3300 ICP A placed Port 2 into a Blocked state to prevent a network loop and placed Port 1 into a Forwarding state to provide connectivity to L2 Switch A.

Now all traffic intended to move between L2 Switch A and L2 Switch B is transported over the 1G Full Duplex link.

Figure 95: Efficient Use of Connections



Topology Restrictions

STP/RSTP is by default set to a 7-hop limit, which is a result of the default STP/RSTP aging parameters. Basically, this means that frames should not pass through more than 7 bridges as this limits the size or diameter of the bridged STP/RSTP network to a maximum of 7 hops.

3300 CXi Release 6.0 and MXe Release 7.0 RSTP Default Parameters

Under 3300 Release 6.0 or later for the CXi and 3300 Release 7.0 for the MXe, the System Administrator has the ability to enable or disable RSTP. The System Administrator cannot alter any of the RSTP parameters. The 3300 ICP STP algorithm is based on the IEEE 802.1w standard. Below are the RSTP parameters and their default values. Note that bridge priority is set to a high value which in turn means a low priority. This value was chosen to minimize the possibility of the CXi or MXe becoming the Root Switch.

It is the Root Switch parameters that control overall network RSTP operation.

Bridge Parameters

- STP ENABLE>false
- BRIDGE PRIORITY=61440
- BRIDGE MAX AGE=20 seconds
- BRIDGE HELLO TIME=2 seconds
- BRIDGE FWD DELAY=15 seconds

Port Parameters (Ports 1 to 16)

- PORT ENABLE=true
- PORT FAST=false
- PATH COST-10Mb/s = 2,000,000
- PATH COST-100Mb/s = 200,000

- PATH COST-1Gb/s = 20,000
- PORT PRIORITY-128

3300 LX, 700 User and 250 User 5.0 STP Default Parameters

Under Release 5.0 the System Administrator has the ability to enable or disable STP. The System Administrator cannot alter any of the STP parameters.

The 3300 ICP STP algorithm is based on the IEEE 802.1D standard of 1998. Below are the STP parameters and their default values. Note that bridge priority is the highest possible value which in turn means the lowest priority. This value was chosen to minimize the possibility that the 3300 ICP becoming the Root Switch.

It is the Root Switch parameters that control overall network STP operation.

Bridge Parameters

- STP ENABLE—false
- BRIDGE PRIORITY—65535
- BRIDGE MAX AGE—20 seconds
- BRIDGE HELLO TIME—2 seconds
- BRIDGE FWD DELAY—15 seconds

Port Parameters (Ports 1 to 4)

- PORT ENABLE—true
- PORT FAST—false
- PATH COST—10
- PORT PRIORITY—255

STP RSTP Performance and Convergence

When a topology change takes place in a well designed network running STP, the network does not reconverge for approximately 50 seconds; a network running RSTP takes about 3 seconds to reconverge.

When a topology change takes place in a well designed network running STP the network will not reconverge for about 50 seconds. The protocol uses this guard time interval to ensure that all ports that should be blocked do get put into the blocking state before any ports get put into the forwarding state. This ensures that loops do not get created during a topology change.

Convergence time can be optimized with the correct usage of Port Fast, Uplink Fast and Back Bone Fast when configuring the L2 switching infrastructure. For details refer to the L2 Switch vendor's documentation.

Convergence time can also be optimized when the number of Bridge or Switch hops is kept to a minimum, for details see [Minimizing Bridge Switch Hops](#).

Altering L2 Switch STP/RSTP timer values is not a recommended method of obtaining optimal convergence times unless the System Administrator has a thorough understanding of STP/RSTP and understands the risks involved.

Using non-default values for STP/RSTP timer values will decrease the guard time interval referred to above which in turn erodes any reconvergence safety margin in the network and increases the possibility of loops being formed while STP/RSTP is reconverging.

Minimizing Bridge Switch Hops

One method of achieving optimal STP/RSTP convergence times involves minimizing the number of Bridge or Switch hops in the network.

The following Figure provides an example of a resilient network that will provide optimal STP/RSTP convergence times by keeping the number of bridge hops to a minimum. The root L2 switch would be one of the Core Network L2/L3 switches.

If additional access network L2 switches are required for connecting phones, the switches should be added at the same level as the existing access layer switches so that additional switch hops are not added, e.g. the phones are directly connected to the access layer network.

Figure 96: Network Optimized for Convergence Time





Figure 97: Network Not Optimized for Convergence Time on page 345 provides an example of a resilient network that will provide less than optimal STP convergence times since the number of bridge hops is not kept to a minimum, for example, the phones need to transit two L2 switches before they can connect to the access layer switches. The root L2 switch would be one of the core network L2/L3 switches.

Figure 97: Network Not Optimized for Convergence Time



Enabling STP RSTP

The MiVoice Business systems are shipped from the factory with STP/RSTP disabled. The L2 switch built into the systems will start running during the 3300 ICP initialization/boot phase. For instructions on how to enable STP/RSTP, refer to the *MiVoice Business System Administration Tool Help*.

The following precautions should be observed to prevent temporary network loops from forming and to minimize network disruptions:

- If multiple cables are connected between the MiVoice Business system and the L2 switch(es) during the initialization/boot phase a network loop will be created by the system, this loop will impact the network in a detrimental way unless STP has been previously enabled on the L2 switch(es).
- The safest way to proceed is to configure the L2 switch(es), enable STP/RSTP on the MiVoice Business system and then connect the cables.

**CAUTION:**

Enabling STP/RSTP in a live network will cause service disruptions to the end users while the network is converging. To avoid impacting users, enabling of STP/RSTP should be conducted outside of core hours.

Networks Running Different Versions of STP

There are many of versions of Spanning Tree Algorithm (STA) protocols in existence, a number of these are proprietary and there are three main open standard protocols as listed below. This document does not discuss proprietary STAs; the three open standard STAs are described below.

IEEE Standard 802.1D, also known as Spanning Tree Protocol (STP), provides convergence times of about 50 seconds.

IEEE Standard 802.1w also known as Rapid Reconfiguration of Spanning Tree Protocol (RSTP) provides convergence times in the order of 1 to 3 seconds.

IEEE Standard 802.1s also known as Multiple Spanning Trees Protocol (MSTP) provides support for multiple VLANs and provides convergence times similar to RSTP.

3300 LX, 700 User and 250 User ICPs

The 3300 LX, 700 User and 250 User ICPs support the IEEE Standard 802.1d Spanning Tree Protocol. Both RSTP and MSTP are backwards compatible with STP which means that the LX, 700 User or 250 User 3300 ICPs can be deployed in data networks that are running RSTP or MSTP.

STP convergence times of 50 seconds will apply to the network segment running STP and RSTP/MSTP convergence times of 1 to 3 seconds will apply to the network segments running RSTP/MSTP.

When installing the LX, 700 User or 250 User 3300 ICP into a RSTP or MSTP based network the System Administrator might need to adjust certain parameters in the L2 infrastructure and/or comply with specific L2 topology recommendations. These details are outside of the scope of this document since they are typically vendor specific and can also be specific to a certain model of L2 switch.

The System Administrator should pay particular attention to how Port Cost is programmed on the L2 infrastructure. Port Cost is used to select link redundancy or control traffic load balancing and bandwidth optimization.

The System Administrator should also refer to the L2 switch Vendor's documentation for planning the integration of these 3300 ICPs into a RSTP or MSTP environment.

3300 CXi and MXe ICP

The 3300 CXi and MXe ICP support the IEEE Standard 802.1w Rapid Reconfiguration of Spanning Tree Protocol.

RSTP is compatible with MSTP which means that the CXi and MXe can be deployed in a data network that is running MSTP. RSTP/MSTP convergence times of 1 to 3 seconds will apply to the network segments running RSTP/MSTP.

RSTP is backwards compatible with STP which means that the CXi and MXe can be deployed in a data network that is running STP.

RSTP convergence times of 1 to 3 seconds will apply to the network segments running RSTP and STP convergence times of 50 seconds will apply to network segments running STP. When deploying the CXi and MXe in a network running STP, the System Administrator should pay particular attention to how Port Cost is programmed on the L2 infrastructure. Port Cost is used to select link redundancy or control traffic load balancing and bandwidth optimization.

When installing the CXi or MXe into a RSTP or MSTP based network, the System Administrator might need to adjust certain parameters in the L2 infrastructure and/or comply with specific L2 topology recommendations. These details are outside of the scope of this document since they are typically vendor specific and can also be specific to a certain model of L2 switch.

The System Administrator should refer to the L2 switch Vendor's documentation for planning the integration of the 3300 ICP into a RSTP or MSTP environment.

STP and Resilient Heart Beat Interactions

This section is intended to describe the potential interactions between STP reconvergence times and resilient phone behavior during a network element failure or a network topology change. The topics discussed in this section only apply to 3300 ICPs that are running STP in a resilient configuration. This Section does not apply to 3300 ICPs that are running RSTP.

As described below there will be some instances in a resilient/STP enabled network when phones might rehome to their secondary ICP only to be re-directed back to their primary ICP once STP has reconverged the network.

This is normal system behavior, the System Administrator should consider the following recommendations if movement of phones from the primary ICP to the secondary ICP and back again during a network failure is perceived as a problem:

- Tuning of the ICP heartbeat timer is not recommended since the 30 second default value has been determined by Mitel to be the optimum value.
- Minimizing the STP reconvergence time is where the most value will be realized, this can be accomplished by:

- Adhering to the network guidelines presented in this document.
- If the System Administrator has a thorough knowledge of STP and understands the risks involved, the STP default values can be tuned to enhance reconvergence times.

Behavior

The ICP issues heartbeat messages to the phones, this is the mechanism that the resilient phones use to determine if the ICP is still functional and if the ICP can be contacted over the network. When a resilient phone fails to receive two consecutive heartbeat message, it will rehome to its secondary (backup) ICP. Phones that are not resilient will become non-functional if the ICP cannot be reached.

Resiliency heartbeats between the ICP and phones occur every 30 seconds which is the default value. Time of day message updates sent between the ICP and the phones are also used as a heartbeat mechanism, and will at times supersede the need for the ICP to issue a heartbeat to the phones.

When a network link is broken or a network topology change occurs the Spanning Tree Protocol issues a Topology Change Notice (TCN). In an optimally designed network STP will take about 50 seconds to reconverge the network once a TCN has been issued.

The 3300 ICP heartbeats and STP topology change notices are running asynchronously to each other.

Depending on what point in time a link is broken and where the broken link is located, either of the following scenarios could occur:

- STP could heal the network before the phones miss two consecutive heartbeats since detection of heartbeat absence could take from 30 to 60 seconds depending on when the last valid heartbeat was detected.
- The phones could detect the absence of two consecutive heartbeats before STP heals the network, in which case the phones will either
 - If the secondary ICP is reachable, the phones will rehome to their secondary ICP, only to be sent back to their primary ICP once STP reconverges the network.
 - If the secondary ICP is unreachable the phones will unsuccessfully attempt to rehome to their secondary ICP.

Note:

Phones that are resilient capable but not configured as resilient will attempt to rehome with their primary ICP if they miss two consecutive heartbeats.

Actual behavior will typically lie somewhere between the two conditions mentioned above since when a number of phones are involved some phones will not detect missing heartbeats and other phones will detect missing heartbeats before STP heals the network.

The following Section discusses two network design approaches and how these different approaches can affect phone re-homing behavior.

Network Design and Effects on Phone Rehoming

There are two different approaches to network design to consider that can affect phone Fail-over or rehoming behavior.

One approach is used to optimize the time it takes for phones to rehome to their secondary ICP when a network failure occurs, however to accomplish this there may be certain network links that are not redundant.

The second approach provides a high level of network link redundancy but phones that are rehoming due to a network failure may take bit longer than phones in the first approach.

Which approach is used needs to be based on the System Administrator's resiliency and phone rehoming requirements.

Network Design for Optimal Phone Rehoming

Network Description

The following diagram shows a network that will allow fast Fail-over of phones in the event of an ICP failure. When a network failure occurs this network design will reduce the chance of phones re-homing to their secondary ICP only to be re-directed back to their primary ICP once the network has reconverged.

The ICPs are connected to their respective distribution layer switches with a single ethernet connection, the distribution layer switch ports used to connect ICPs are set to PortFast.

The phones are connected to their respective access layer switches with a single ethernet connection, the access layer switch ports used to connect phones and PCs are set to PortFast.


The access layer switches and distribution layer switches are connected with dual physical connections and are running RSTP.

The distribution layer switches and routers are connected with dual physical connections and are running RSTP.

Behavior During Failure

- If 3300 Controller A fails the phones connected to Access Switch A will detect the absence of heartbeats from this ICP within 60 seconds and rehome to 3300 Controller B.
- If the network connection between Access Switch A and Distribution Switch A fails, RSTP will reconverge the network within about 3 seconds. The reconvergence should take place before the phones connected to Access Switch A detect a missing heartbeat.
- If the network connection between Distribution Switch A and 3300 Controller A fails, the phones connected to Access Switch A will detect the absence of heartbeats from this ICP within 60 seconds and rehome to 3300 Controller B.
- If Distribution Switch A fails completely the phones that had been registered with 3300 Controller A will rehome to 3300 Controller B once they have detected the absence of heartbeats.
- If Access Switch A fails completely the phones connected to Access Switch A will become unusable.

Figure 98: Optimal Phone Rehoming Configuration



Network Design for Optimal Physical Layer Resiliency

Network Description

The following diagram shows a network that is optimized to provide physical layer redundancy. However, under certain failure scenarios network reconvergence may take longer than in the network described in the previous section and the phones may initially move to the secondary ICP and then move back to the primary ICP once the network reconverges. The actual phone behavior during a network failure will depend on the specifics of the network failure.

Each 3300 Controller is connected to both distribution layer switches. The connection between a particular controller and a particular switch is provided by a single physical connection. Since STP is running on both of the 3300 Controllers and both of the Distribution Switches connection redundancy is provided.

To ensure that both 3300 Controllers participate correctly in STP the Distribution Switch ports that are used to connect to the 3300 Controllers must be configured so that PortFast is disabled.

Two physical links are normally considered sufficient for physical layer redundancy. However, if a higher level of physical layer redundancy is desired multiple connections could be used to connect 3300 Controller A to Distribution Switch A and Distribution Switch B and multiple connections could be used to connect 3300 Controller B to Distribution Switch B and Distribution Switch A.

The phones are connected to their respective access layer switches with a single ethernet connection, the access layer switch ports used to connect phones and PCs are set to PortFast.

The access layer switches and distribution layer switches are connected with dual physical connections and are running RSTP.

The distribution layer switches and routers are connected with dual physical connections and are running RSTP.

Behavior During Failure


- If 3300 Controller A fails STP will detect the failure and reconverge the network, this process might take 50 seconds to complete. Once the network has reconverged the phones connected to Access Switch A (which have already detected an absence of heartbeats from Controller A) will now be able to rehome to 3300 Controller B.
- If the network connection between Access Switch A and Distribution Switch A fails, RSTP will reconverge the network within about 3 seconds. The reconvergence should take place before the phones connected to Access Switch A detect a missing heartbeat.
- If the network connection between Distribution Switch A and 3300 Controller A fails, STP will detect the failure and start to reconverge this network segment, it may take about 50 seconds for this process to complete.

Some of the phones connected to Access Switch A will detect the absence of heartbeats from this ICP before STP has reconverged this network segment and rehome to 3300

Controller B. These phones will be re-directed back to 3300 Controller A once STP has reconverged the network.

Some of the phones connected to Access Switch A will not detect the absence of heartbeats from this ICP before STP has reconverged this network segment and, therefore, will not rehome to 3300 Controller B. These phones will remain homed to 300 Controller A.

Figure 99: Optimal Physical Layer Redundancy



STP and VLANS

The 802.1d and 802.1w MAC Bridges Spanning Tree Protocol does not sense VLANs. Therefore, it does not directly affect the portion of the network utilized by the 3300 ICP and the IP phones. If this protocol is being used throughout the network, the System Administrator needs to be aware that STP and RSTP are insensitive to VLANs.

The following diagram shows two 802.1Q compliant L2 Switches. VLAN 1 is connected via the 802.1Q trunk and the regular link.


VLAN 2 is only connected via the 802.1Q trunk.

VLAN 1 presents a problem in that it forms a network loop. When STP is enabled on L2 Switch A and L2 Switch B STP will break the loop formed by VLAN 1 by blocking one of the ports on one of the L2 Switches. If STP chooses to block one of the ports used for implementing the 802.1Q trunk there will be two outcomes:

- The network loop will be broken, as it should be.
- Connectivity for VLAN 2 may be broken, this is not desirable.

To find the preferred solution to this problem the System Administrator should consult the L2 vendor's documentation.

Figure 100: Asymmetrical VLANs



Known Issues

Failure to Forward BPDU Packets Causes Network Loops. Some third party switches do not forward Bridge Protocol Data Unit (BPDU) packets when STP/RSTP is disabled on the switch. This will prevent switches that are running STP/RSTP from detecting network loops since STP/RSTP relies on BPDUs to determine if a network loop exists. This situation can cause broadcast storms and connectivity issues. To correct this situation

- Replace the offending switch with a switch that is fully STP/RSTP compliant.
- Enable STP/RSTP on the offending switch.

A unidirectional link that remains in an enabled state will fail to forward BPDUs in one direction. This situation can be caused by defective cabling or defective interface hardware. For example if a link between switch A and switch B is provided by a fibre cable and the A to B transmit fibre is defective but the transmission path from B to A is functional the following will take place:

- BPDUs transmitted by A will not be received by B.
- If B is supposed to be in the Blocking state it can only remain in the Blocking state if it is receiving BPDUs from a switch with a higher priority.
- Since B is not receiving BPDUs from a switch with a higher priority, B will start to forward traffic to A and this will create a network loop that STP is unable to detect.

To try and prevent this situation from occurring:

- Some L2 switch vendors provide proprietary schemes for detecting unidirectional links. If this feature is available the System Administrator should enable the feature.

Ethernet Duplex Mismatch Causes Network Loops

Ethernet duplex mismatch is one of the leading culprits for STP/RSTP failures and problems. The problem is created when one device has been hard coded into full duplex mode and the other device is set for auto-negotiation. This results in the device that is set for auto-negotiation moving into the half duplex mode because the hard coded device is unable to auto-negotiate.

If a switch sends BPDUs out of a port that is set to half duplex to a switch port that is configured for full duplex operation the duplex mismatch can cause a network loop to occur.

The reason is that the full duplex switch will transmit packets even if the half duplex switch is transmitting. If there is enough traffic originating from the full duplex switch the half duplex switch will be unable to successfully transmit packets due to collisions and the collision/backoff algorithm, this includes BPDU packets.

Eventually the full duplex switch will conclude that it has stopped receiving BPDUs from the half duplex switch and conclude that this link or switch is faulty, the full duplex switch will then unblock it's previously blocked port and a loop is created.

To correct this situation:

- Ensure there are no duplex mismatches between switch ports.
- If the switch vendor has provided a mechanism for detecting duplex mismatches it should be enabled.

ICP Operation with STP/RSTP Enabled Networks

This section provides a summary of how ICPs should be integrated into networks that are running STP/RSTP.

Using the ICP Without STP/RSTP in a Network Running Spanning Tree

This section describes using an ICP that either does not support STP/RSTP or an ICP that has STP/RSTP disabled in a network that is running STP/RSTP. These recommendations pertain to the following situations:


- Using a Pre-release 5.0 ICP in a network running STP/RSTP. (Prior to Release 5.0 STP was not supported).
- Using a Release 5.0 ICP with STP/RSTP disabled in a network running STP/RSTP.
- Using a 3300-100 or 3300-200 ICP a network running STP. (These products do not support STP/RSTP).
- Using a 3300 CXi or MXe ICP with RSTP disabled in a network running STP/RSTP.

Since there is a real possibility of creating network loops care must be exercised when integrating an ICP that does not support STP (or has STP disabled) into a network that is running STP. The following rules must be observed:

- Ideally, there should be only one physical Ethernet connection between the ICP and the Distribution Network L2 Switch. However, if multiple physical connections are used, the guidelines under [Physical Connection, Layer 2 Switch to ICP](#) and [VLAN and Priority Configuration, Layer 2 Switch to ICP](#) must be followed.
- The system administrator must ensure that the above rule is not violated at a later date by network technicians.
- The Ethernet port on the L2 switch that is used for connecting to the ICP should be configured for PortFast operation, this will ensure that this port skips the first stages of the STP/RSTP Algorithm and directly transitions to Forwarding mode.

The following diagram depicts how an ICP that does not support STP/RSTP should be connected to the Distribution Network.

Figure 101: ICP (Without STP/RSTP) to L2 Connection



Using the ICP with STP/RSTP in a Network Running Spanning Tree

This section describes using an ICP that supports STP in a network that is running STP/RSTP. These recommendations pertain to the following situations:


- Using an LX Controller (or 700-user 3300 ICP upgraded to Release 5.0) with STP enabled in a network running STP/RSTP.
- Using a 250-user 3300 ICP upgraded to Release 5.0 with STP enabled in a network running STP/RSTP.

The following rules must be observed:

- Multiple connections are allowed between the ICP and the L2 Switch(es). However, two connections is considered a sufficient number for resiliency.
- The L2 Switch port(s) should NOT be configured for PortFast operation (PortFast should be disabled), this will ensure that this port fully participates in the STP/RSTP algorithm.

The following diagram depicts how an ICP that supports STP/RSTP should be connected to the Distribution Network.

Figure 102: ICP (with STP/RSTP) to L2 Connection



8.6.10 External Applications

Voice Mail

Embedded and external (either can be centralized) voice mail is resilient. For information about resilient voice mail, see ["Voice Mail", Planning Voice Mail and Other Options](#) on page 223, and [Configuring T1 E1 Trunk Resiliency](#) on page 232.

Other Applications

For information on external applications in a resilient environment, see [Interactions Between System and External Applications](#) on page 104.

8.7 3300 ICP

Configurations and Sizes

The 3300 ICP is available in several configurations to meet different customer requirements.

To determine which configuration is the most appropriate for a particular application, refer to the 3300 ICP Engineering Guidelines on Mitel Online and to [Planning License Requirements](#) on page 214 of this Resiliency guide.

If you plan to have resilient devices in your system or if you plan to cluster 3300 ICPs, see [De-rating ICP Capability to Support Resiliency and Clustering](#) on page 359.

The capacity of a 3300 ICP controller depends on the

- Level of traffic
- Number of physical trunks
- Number of E2T channels,
- Number of echo channels provided.

These values are all interconnected. If you change one value, the change will affect the value of one of the other resources. If you require a configuration that is different from the configurations that are outlined in the Engineering Guidelines, consult a Mitel Customer Engineering Services representative.

8.7.1 De-rating ICP Capability to Support Resiliency and Clustering

Clustering and Resiliency introduces additional burdens to the system. This has the effect of reducing the capability of an individual system, depending upon the configuration and use of the system.

Table 28 shows some typical network configurations with user de-ratings for a 3300 LX 700 User Controller. The reduction factors provided in [Derating the Number of Users on an LX Controller](#) can also be used for Controllers that support less than 700 users, but these reduction factors cannot be used for a 3300 LX 1400 User Controller. To support 1400 users the 3300 LX 1400 User Controller must use IP trunks to connect to PSTN gateways. As a result of functionality being distributed Table 28 is not applicable to the 1400 User ICP.

Table 19: Derating the Number of Users on an LX Controller

Network Configuration	IP Phone Capacity	Reduction (%)
Standalone	700	0%
Networked (Geographically isolated)	580	17%
Clustered (with PSTN trunk sharing)	410	41%
Clustered (without PSTN trunk sharing)	650	7%

The configurations shown above are:

- **Standalone**—This is an individual unit not connected by any form of networking. An example might be a SME business
- **Networked**—This is the case where a number of locations are interconnected, but the level of interoffice traffic is not particularly high. This might be an example of a business with multiple corporate offices, say with an office in Vancouver and another in New York.
- **Clustered (with PSTN trunk sharing)**—This would be where a number of systems inter-operate in order to create a bigger system. This might be a larger office where

there are a number of trunk connections to the PSTN, but where the PSTN can present a call to any of the trunks. Therefore an incoming call could arrive on any system, and likewise on outgoing call could also go through any system.

- **Clustered (without PSTN trunk sharing)**—This would be a business, perhaps dispersed across a city connected via a MAN, where each business is connected to the PSTN via local trunks, but where internal traffic can flow freely from office to office. This might be a campus environment, a large department chain or a government establishment.

In a resilient environment, the above configurations need to be considered as they have affect on the inter-unit traffic and loading under normal conditions. From a numbering plan the systems are considered to be clustered.

In terms of the number of resources to be assigned to each unit, for resiliency, the main consideration is the expected operation following a Fail-over condition. If the level of blocking and delay to call progression tones, e.g. dial tone, are to remain within normal limits then the units need to be dimensioned for the worst case, i.e. Fail-over. In most cases this may simply be an additional DSP card. The alternative is that the box will increasingly block in terms of available resources, affecting all users connected to that box. In practice, it's traffic that generates resource requirements. So, if a number of users were to Fail-over, but then make few phone calls, the loading on the system would be reduced from that normally expected. The number of users to be handled in a Fail-over situation needs to be considered and pre-assigned to each unit to determine which resources are needed.

Where additions are required on a system, then it is recommended that Mitel Customer Engineering Services be contacted. They can then recommend the best course of action. An example might be more DSP for more VM capability, or DSP for additional compression on IP trunks.

Note:

An LX Controller with a Real Time Controller (RTC) that has 256 MB of RAM supports up to 700 users. An LX Controller with a RTC that has 512 MB of RAM supports up to 1400 users.

8.7.2 Provisioning MiVoice Business Resources

When a MiVoice Business system is being deployed in a resilient network, the system administrator must bear in mind that the ICP needs to be provisioned to support both the primary users and the secondary users. For example, an LX controller with 256 MB of RAM can support a total of 700 IP Phones.

If this ICP is acting as a primary ICP for 400 IP Phones, then this ICP will only be able to provide secondary ICP support to 300 IP Phones in the event that the primary ICP for the 300 Phones has failed. In this case it is the maximum number user licenses that limits the total number of supported phones.

There are other system limits that must be taken under consideration when provisioning ICP resources.

For details regarding ICP performance, maximum limits and guidelines on provisioning an ICP refer to the *3300 ICP Engineering Guidelines and if necessary consult with Mitel Customer Engineering Services*.

8.8 Software Version Control

Ensuring ICP-to-ICP Software Version Consistency

3300 ICP Releases 4.0 and later, do not provide a mechanism for ensuring that ICPs are running the same version of software. The following guidelines pertain to ICP software versions in a resilient network:

- All 3300 ICP system within a resilient cluster must be running a minimum of 3300 ICP Release 4.0 software.
- Any ICP running Pre-4.0 software must be deployed outside of the resilient cluster, if a Pre-4.0 ICP wishes to communicate with an ICP that is within the resilient cluster, the Pre-4.0 ICP must use a boundary node ICP as an intermediary or gateway.

8.8.1 IP Device Interface Protocol Version Control

An IP device (IP phone or other IP endpoint) may be required to change its protocol interface software version, due to the ICP recognizing an interface version incompatibility. This can occur in any of the three following scenarios:

- The IP device is registering for the first time with its primary or a secondary system, as at start up or after a configuration change.
- The IP device is failing over to a different ICP as a result of failure handling.
- The IP device is being handed off to a different ICP as a result of Fail-back, maintenance, or for other reasons.

In each of these circumstances, the ICP's expectations for the IP device protocol interface version may be higher or lower than the protocol interface currently supported by the IP device, so in those cases the IP device may need to be updated with a new software load and/or a different IP device protocol interface level negotiated. As a last resort, backward compatibility mechanisms are invoked at the ICP.

If it is determined that the ICP and device interface protocols do not match, the device-side interface protocol will be altered to match the ICP interface protocol by one of three different methods. The method employed depends on whether the IP device supports software re-loading or not, and the vintage of the IP device (does it support the negotiation mechanism or not). The three different methods that an ICP uses for dealing with an IP device interface protocol mismatch are described in the following sections.

Dynamically Loadable Devices (Forced Software Update Method)

5xxx-series IP Phones that have been loaded at initial start up with Release 4.0 compliant software or higher are considered to be dynamically software reloadable, these devices also support the software version negotiation mechanism.

Note:

1. The TFTP server address supplied to the IP phone is normally the IP address of the ICP itself, i.e. the ICP uses its own internal TFTP server as the repository for compatible software loads.
2. The ICP's required version of interface protocol software is administratively configurable, to allow deployment of "higher" but known to be backwards compatible set software into the field (for example to allow uniform set software across a mixed cluster, or for bug-fix/special load purposes), and to allow for use of external TFTP server.

Non-loadable Devices (Version Negotiation Method)

Non-5xxx-series IP devices (or applications) such as Unified Communicator Advanced which have been upgraded to support interface protocol version reporting and negotiation messaging will be able to report their interface protocol version to the ICP, they will also support the version negotiation mechanism. However, these devices are not dynamically software reloadable.

Note:

1. If the device's interface protocol version is compatible to the ICP's required interface protocol version, the device is accepted for control and the registration completes normally.
2. If the device's interface protocol version is not compatible a downward negotiation is initiated by the ICP.
3. Once an interface protocol match is established the device registration proceeds and the device is put in service at the ICP.
4. If an interface protocol match cannot be established, the device registration process is terminated and the ICP closes its communication socket with the device.

Special Case Handling for Backwards Compatibility to Version 3.x Interfaces

For devices that are unable to report their interface protocol version, the ICP will be unable to determine the version of the device's interface protocol. In this case, the ICP will assume that the device is running a pre-Release 4.0 version of the interface protocol.

In this situation either the device is a 5xxx-series IP Phone that has not been loaded with Release 4.0 or later compatible software (possibly a configuration error caused the phone to access an inappropriate TFTP server at start up and as a result the phone received the wrong version of software) or the phone is a non-5xxx series IP device such as a Unified Communicator phone that has not been upgraded with Release 4.0 or later compatible software.

Regardless, no negotiation or reloading will be possible – the interface protocol version supported remains “unknown”.

However, it is not adequate to simply refuse registration in these cases, so these devices must be accepted for registration, configuration can only ensure the device type *can* support resiliency, not that the software it is running actually does.

Note:

1. 5xxx-series IP devices with an “unknown” interface protocol version (i.e. Pre-4.0) are blocked from being configured on a secondary ICP.
2. In release 4.0, the specific list of IP devices and applications (non-5xxx series IP devices) known to support resiliency has been effectively hard-coded into the ICP via the data base.
3. IP Devices, both 5xxx-series devices and external applications not in the SX2Kdata base list of resiliency capable devices are blocked from being configured on a secondary ICP.
4. No mechanism is provided to update the list of resiliency capable devices through administrative control in the field. As a result, when external applications become resiliency-capable, the ICPs to which they are configured will need software upgrades so that the ICPs recognize the new resiliency capable applications.

8.8.2 Ensuring Phone Application Software Version Consistency

If multiple TFTP servers exist in the network, the system administrator must ensure that the phone application software available on these TFTP servers is at the same software revision level. If phone application software consistency is not assured on multiple TFTP servers, there is no guarantee that users will be provided with consistent phone features.

The following describes how the 500x series of IP sets obtain software:

- The phone uses a boot load that is resident in on-board Flash memory to enter into a DHCP cycle.
- The DHCP server provides the phone with the IP address of a TFTP server, the phone downloads main executable code from the TFTP server.
- The main executable code in the phone behaves as if it is ROM code, so the phone will now enter into another DHCP cycle.
- Once the DHCP cycle is complete the phone starts running telephony functions.

Note:

1. The IP-DECT Phones do not obtain their application software from the TFTP server. Application software for these Phones is downloaded from a PC via a USB interface. For details see the IP DECT Wireless Solution Technical Manual

MiVoice IP Phone s will upgrade their software on resilient Fail-over if the software version on the secondary ICP is higher than the version that is on the phone. If the phone's software version is higher than that of the secondary ICP, then no upgrade is initiated.

Phone Software Storage Location

For the internal TFTP server, phone software is TARed into the ICP software builds and placed in the Sysro/TFTP directory.

For external TFTP servers the file location of phone software is under the control of the system administrator.

- The Teleworker Gateway software is delivered on a CD.
- The IP Phone Analyzer software is being bundled with the MiVoice Business system software.
- The IrDA application software is delivered on a CD, and is installed into the Palm Pilot device.

Minimum Software Version Levels

For correct operation in a resilient Release 5.0 network, the following minimum software version levels must be met or exceeded:

- To support resilient operation, 3300 ICPs must be running 3300 ICP Release 4.0 or later software.
- 3300 ICPs running pre-Release 4.0 software can function as transit nodes in a resilient cluster but cannot participate in resilient operation.
- The following minimum software version levels are required for resilient operation of IP phones:
 - 5002 IP Phone requires a minimum of version 6.3.1.9 software
 - 5215 and 5220 Dual Ethernet Port IP Phones; 5540 IP Console; and the 5485 IP Paging Unit require a minimum of version 8.3.1.11 software
 - 5320 and 5360 IP Phones require a minimum version 1.x software

The following products are not supported as part of the resilient solution:

- SX-2000 NT

