MiCloud Business Solution

ENGINEERING GUIDELINES 3.3
April 2017

Mitel®

NOTICE

The information contained in this document is believed to be accurate in all respects but is not warranted by Mitel Networks™ 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.

Trademarks

MiCloud Business Solution Engineering Guidelines
Release 3.3
April 2017
Document Version 2.0

CHAPTER 1: MICLOUD BUSINESS SOLUTION OVERVIEW

CHAPTER 2: AVAILABILITY AND RESILIENCY

CHAPTER 3: TRAFFIC AND SCALING CONSIDERATIONS

CHAPTER 4: EXTERNAL CONNECTIVITY

CHAPTER 5: NETWORK AND NETWORKING CONSIDERATIONS

APPENDIX A: GLOSSARY

Chapter 1

MICLOUD BUSINESS SOLUTION OVERVIEW

Introduction

MiCloud Business Solution is a Mitel unified communications (UC) solution built on Mitel's MiCollab and MiVoice Business capabilities. UC solutions are easily adapted or scaled to match the changing demands and needs of the hosted market, the enterprise, and the end-user.

This release of the MiCloud Engineering Guidelines applies to MiCloud 3.3. Table 1 shows the product lineup for the MiCloud releases.

Table 1: MiCloud Business platforms and applications by MiCloud release

COMPONENT / MICLOUD VERSION	3.0	3.1	3.2	3.3
VOICE PLATFORMS				
MiVoice Business	7.2 SP1	7.2 SP1	8.0 PR2	8.0 PR3
MiVoice Business Multi Instance	2.0 SP1	2.0 SP1	2.0 SP1	2.0 SP1
MiVoice Business Express	7.1	7.2.1	7.2	7.3 PR1
UNIFIED COMMUNICATIONS (UC) APPLICATIONS				
MiCollab	7.1 PR1	7.2.1	7.2.2	7.3 PR1
MiCollab Client Multi-tenant	7.1 PR1	7.2.1	7.2.2	7.3 PR1
CONTACT CENTER APPLICATIONS				
MiContact Center Business	8.0 SP1	8.1	8.1 SP1	8.1 SP2
MiVoice Business Reporter	8.0 SP1	8.1	8.1 SP1	8.1 SP2
MiVoice Integration for Salesforce		2.0	2.1	2.1
MiVoice Integration for Google		1.1	1.1	1.1
MiVoice Call Recording	9.0 SP2	9.0 SP3	9.1	9.1 SP1
MiVoice Border Gateway Secure Recording Connector	9.2	9.3.1	9.4 PR2	9.4 PR2
MANAGEMENT APPLICATIONS				
Oria	5.0	5.1	5.2	5.3 SP1
Mitel MarWatch	5.1	5.1	--	--
Mitel Performance Analytics	--	--	2.1	2.1
MiCloud Business Analytics	9.0	9.0	3.2	3.3
NETWORKING APPLICATIONS				
MiVoice Border Gateway	9.2	9.3.1	9.4 PR2	9.4 PR2

Table 1: MiCloud Business platforms and applications by MiCloud release

COMPONENT / MICLOUD VERSION	3.0	3.1	3.2	3.3
MiCloud Management Gateway	5.0	5.0	5.0	5.0
Open Integration Gateway	3.0	3.0	4.0	4.0
PRODUCTIVITY APPLICATIONS				
Vidyo	--	Cloud-based	Cloud-based	Cloud-based
MiCloud CRM Integrations	--	--	--	3.3
OTHER				
Redirection and Configuration Service (RCS)	1.1	1.1	1.1	1.1

Unified Communications documentation

Mitel documentation is available on-line through Mitel OnLine.

The *MiCloud Business Solution Blueprint* describes the best practice architectures for deploying MiCloud offerings for various customer sizes. The Blueprint does not include deployment instructions, pricing, order codes, or release note details.

The deployment guides provide step-by-step instructions for implementing a successful MiCloud solution deployment.

- *MiCloud Business Multi-Instance Deployment Guide*
- *MiCloud Business Virtual Deployment Guide*
- Product documents are also needed to help with designing and using your UC system. Product-specific documents are available on Mitel OnLine.

  The MiCloud Business Deployment guides do not include most product-specific details, instead referring to the appropriate guide and section name. Figure 1 shows the hierarchy of products whose documents are referenced in the MiCloud documentation.

Figure 1: MiCloud product hierarchy

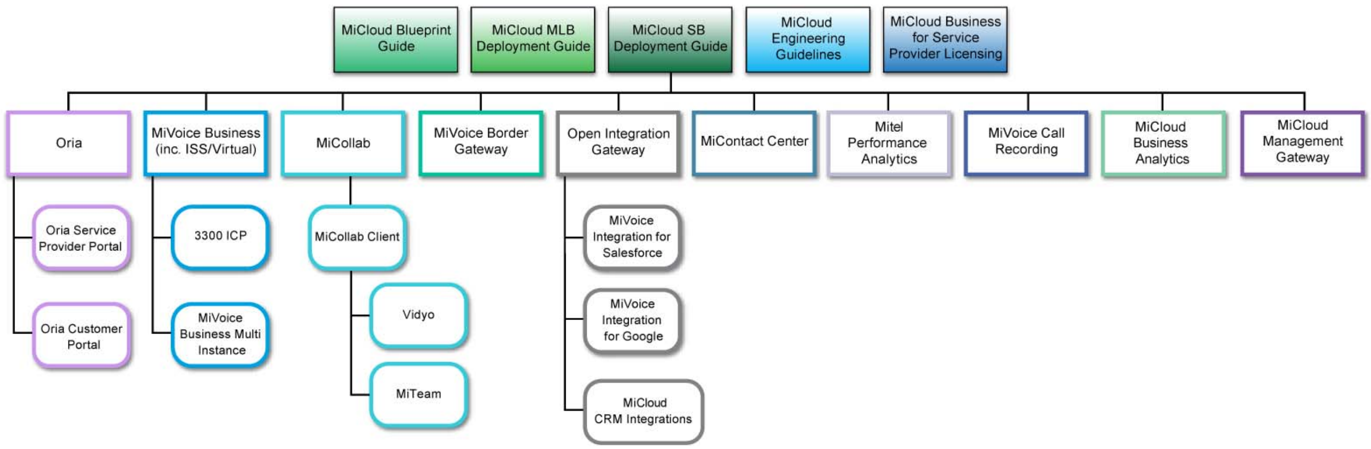
MiCloud Blueprint Guide
MiCloud MLB Deployment Guide
MiCloud SB Deployment Guide
MiCloud Engineering Guidelines
MiCloud Business for Service Provider Licensing
Oria
MiVoice Business (inc. ISS/Virtual)
MiCollab
MiVoice Border Gateway
Open Integration Gateway
MiContact Center
Mitel Performance Analytics
MiVoice Call Recording
MiCloud Business Analytics
MiCloud Management Gateway
Oria Service Provider Portal
Oria Customer Portal
3300 ICP
MiVoice Business Multi Instance
MiCollab Client
Vidyo
MiTeam
MiVoice Integration for Salesforce
MiVoice Integration for Google
MiCloud CRM Integrations

IP1824

Solution topologies

MiCloud Business Solution topologies are a number of different Mitel Unified Communications reference designs to meet service provider and customer requirements, network connectivity requirements, and system Unified Communications scaling. The solutions cover scaling from a few users up to and beyond 10,000 users. Each user may also be associated with multiple devices.

The underlying computing resource infrastructure may affect the choice of reference architecture. The two main architecture types described in this guide are:

- Multi-tenant applications offer most advantages when deployed on large servers, affording higher performance and supporting larger pools of tenants. With multiple customers on a single platform, the ability to configure affinity rules and high availability are important to achieving higher reliability. These requirements tend to align with a bare metal or vCenter IaaS offer.
- Virtual applications require optimization of the computing resources allocated to a single customer, and efficiency in the operational processes for turning up customers. These requirements tend to align with an IaaS offer at the orchestration layer that abstracts the underlying virtual machines.

The topologies include different virtualization technologies, and in many cases a mix of technologies, to take advantage of different features. The applications and call server platforms range from on-premise servers and appliances, through Mitel-optimized virtual call servers, to virtualized application packages running on VMware virtual platforms.

Key architecture components

The MiCloud Business Solution topologies are based on a common set of solution components.

The key MiCloud components include:

Table 2: New architecture components for MiCloud 3.3

NEW COMPONENT	DESCRIPTION
MICLOUD CRM INTEGRATIONS	
	MiCloud CRM Integrations is a new “connector” application that allows seamless integration of many popular Customer Relationship Management (CRM) applications to the MiVoice Business for much improved call center operation.

- **MiVoice Business**

  The feature-rich Mitel communications system provides IP telephony with seamless IP networking and SIP trunking. MiVoice Business offers over 500 telephony features, provided to users through easy-to-use phones and web-based user desktop interfaces. MiVoice Business is tailored for different platforms, including custom hardware (3300 ICP), Industry Standard Servers (x86-based), and virtualized environments (VMware and Microsoft Hypervisor).

- **MiCollab**

  MiCollab unifies Mitel applications into an easy to use, cost effective communications solution. Users have a single point of access to all their Mitel applications through the My Unified Communications portal, a web-based interface offered by Mitel. MiCollab may be deployed with multiple applications per server, or with a single application per server, when increased capacity is required.

  New for MiCollab: Flow Through Provisioning allows SDS data sharing between the MiVoice Business cluster and MiCollab. Flow Through Provisioning replaces the Single Point Provisioning feature.

  The services co-resident with MiCollab include:

  - **MiCollab NuPoint Unified Messaging (MiCollab-UM)** - Provides voice messaging, unified messaging, paging support and speech auto attendant; Users can access their voice mails remotely, listen to their voice mail messages via integrated e-mail clients, or through the telephony user interface.
  - **MiCollab Speech Auto Attendant (MiCollab-SAA)** - Provides speech-enabled auto attendant.
  - **MiCollab Client** - Provides contact management, dynamic status, instant messaging and audio conferencing.

  Note: MiCollab User and Services Provisioning (USP) is not supported in MiCollab Client Multi Tenant.

  - **MiCollab Audio, Web and Video Conferencing (MiCollab-AWV)** - Provides web conferencing supporting audio, video, chat (text), and presentations.
  - **MiVoice Border Gateway** - See "Mitel Gateways" below.
  - Integral management tools including MiCollab Suite Application Services, which provides centralized management of shared system resources, provisioning interfaces and license management, and **MiCollab Client Deployment**, which supports simplified deployment and configuration of MiCollab for Mobile.

- **Management Applications -** Management of the communications solution is provided by:
  - **Embedded Native Management** - Provides web-based access for configuring the full capabilities of the MiCloud platforms. Native management tools are included in all MiCloud platforms.
  - **Oria** - Automates the otherwise manual process of provisioning multiple MiVoice Business and MiCollab instances, and allows customer self-service through the new Customer Administration Portal.
  - **Mitel Performance Analytics** - Provides fault and performance monitoring for Mitel components and selected network devices. These capabilities are enabled with MarProbe, a diagnostic tool deployed in the end-customer network. MarProbe may be deployed in the data center network and/or in the end-customer's on-premise network, depending on the specific monitoring required.

- **Mitel Gateways**
  - **MiCloud Management Gateway** - Provides specialized 1:1 NAT capabilities to support management access from the Service Provider domain to multiple Customer domains.
  - **MiVoice Border Gateway** - Provides a specialized application proxy supporting SIP, MiNet, and web protocols. Some of the key functionality provided by MiVoice Border Gateway includes:
    - Teleworker service for connection to remote SIP and MiNet end-points
    - SIP trunk proxy for connection to external third-party SIP providers
    - Secure Recording Connector to facilitate recording of voice streams by call recording equipment
    - Web proxy to allow UC web clients to connect securely to LAN-based applications
  - **Open Integration Gateway** - Provides a proxy for connection to business applications, such as SalesForce.com and Google.
- **MiContact Center Business**

  A powerful Contact Center solution providing both basic voice call center capabilities and rich multi-media contact capabilities, and supported by a comprehensive real-time and historical reporting engine.
- **MiVoice Call Recording**

  The Mitel call recording solution provides core voice call recording and quality monitoring that can be deployed for general business users and contact center agents.
- **MiVoice Business Reporter**

  MiVoice Business Reporter provides Customers with rich business communications usage reporting and optional cost analysis for departmental attribution.
- **MiVoice for Skype for Business**

  MiVoice for Skype for Business enables Mitel call control, natively launched from Microsoft Skype for Business.
- **Mitel end-points**

  Mitel offers a range of desk phones, soft phones, wireless phones, and web-based applications to suit user communication preferences.

The solution components are available in various packages and configurations optimized for different network and system requirements. These variations are reflected in the component selection and network design for the topologies described in the following sections.

Cloud services

MiCloud Release 3.3 includes the following cloud services:

- MiCollab MiTeam Collaboration
  - Vidyo - Immersive video offering from Mitel Cloud Services
  - Mitel Performance Analytics - Performance analytics to monitor network and platform performance in order to proactively detect and solve issues.

  - MiCloud Business Analytics - Business analytics for call metrics to improve communications management and reporting
- Management services
  - Application Management Center (AMC) - Provides services for software downloads and license management
  - Redirection and Configuration Service (RCS) - Used to simplify configuration of Mitel IP phones
  - Redirect Server - Used to simplify installation and configuration of MiCollab for Mobile

The *MiCloud Solution Blueprint* describes the components used in the topologies and their purpose and functionality.

MiCloud topologies

Engineering guidelines in this book apply to the following MiCloud topologies, divided between multi-tenant architectures and single-tenant architectures.

MiCloud Business Multi-Instance:

- “Small Business architecture” on page 10
- “Small Business Virtual architecture” on page 11

MiCloud Business Virtual:

- “Small Medium Business (SMB) architecture” on page 13
- “Small Medium Business, Low Density (SMB-LD) architecture” on page 14
- “Medium Large Business (MLB) architecture” on page 16
- “Scalable architecture” on page 18

MiCloud Business Multi-Instance

MiCloud Business Multi-Instance architectures use the MiCollab Client and MiContact Center multi-tenant capabilities. They are ideal for service providers who cater to smaller businesses, such as those with fewer than 50 users.

Small Business architecture

The MiCloud Business Solution Small Business (SB) architecture is a straight-forward solution for service providers offering voice-centric deployments, with a light level of UC attachment, to end-customers. The SB architecture is based on the MiVoice Business Multi Instance platform, and easily scales to multiple instances and customers, ideal for a small business where there are typically fewer than 50 users per customer.

The SB architecture uses a common address space for the service provider, and is essentially a flow through model, from SIP trunks to end-users. This type of deployment is commonly referred to as an Over-the-Top (OTT) network deployment (deployed over the public Internet), although dedicated trunking, such as MPLS, can also be used.

MiVoice Border Gateway functionality is an integral part of network security in this solution. A group of MiVoice Border Gateways isolates the SIP trunk connection point. End-user isolation from the public network, or Internet, is provided by an additional group of MiVoice Border Gateways. Customer isolation from the public network is accomplished through standard Internet access firewalls and NAT.

Customer management is provided by a publicly-accessible management portal, or directly through the service provider.

Figure 2: Small Business architecture

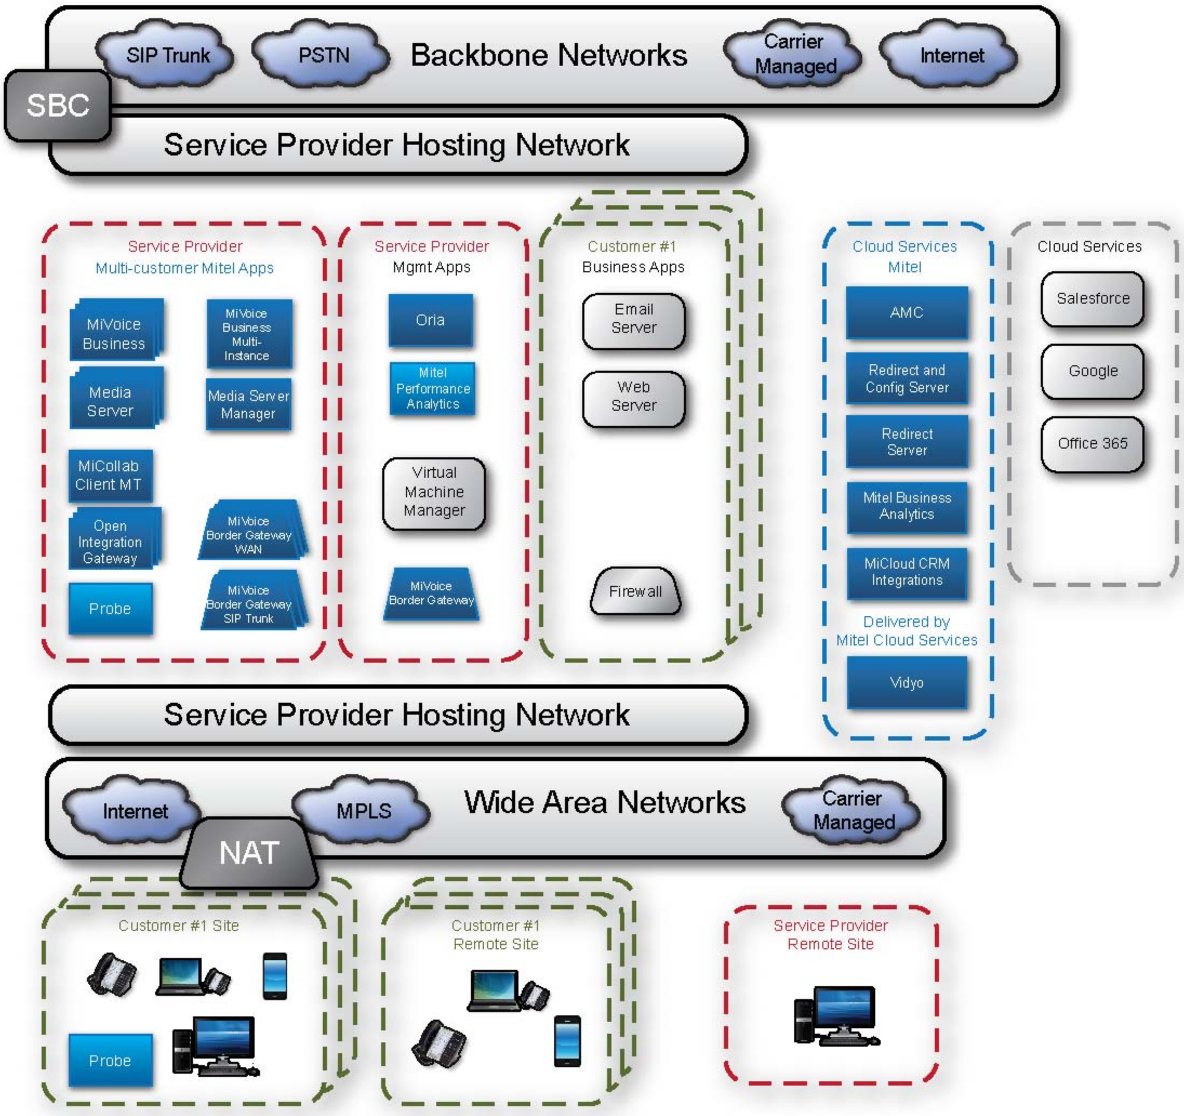

Small Business Virtual architecture

The SB Virtual architecture is very similar to the SB architecture except that, rather than using the MiVoice Business Multi Instance platform, it uses the VMware platform.

In SB Virtual, multiple instances of MiVoice Business Virtual are deployed in place of the MiVoice Business instances in MiVoice Business Multi Instance (that is, the MiVoice Business Call Manager and Media Server). Aligned with deployment in an IaaS virtualized infrastructure, MiVoice Border Gateway Virtual is substituted for the MiVoice Border Gateway on ISS. Component constraints are the same as described for the SB architecture.

Compared to the SB architecture, there are no significant differences in the service or management capabilities. However, for service provider operations staff, the MiVoice Business Multi Instance management interface is no longer available to manage the MiVoice Business instances. In the SB Virtual case, operations staff access the MiVoice Business Virtual instances directly for any configuration or customizations not handled by Oria.

Figure 3: SB Virtual architecture

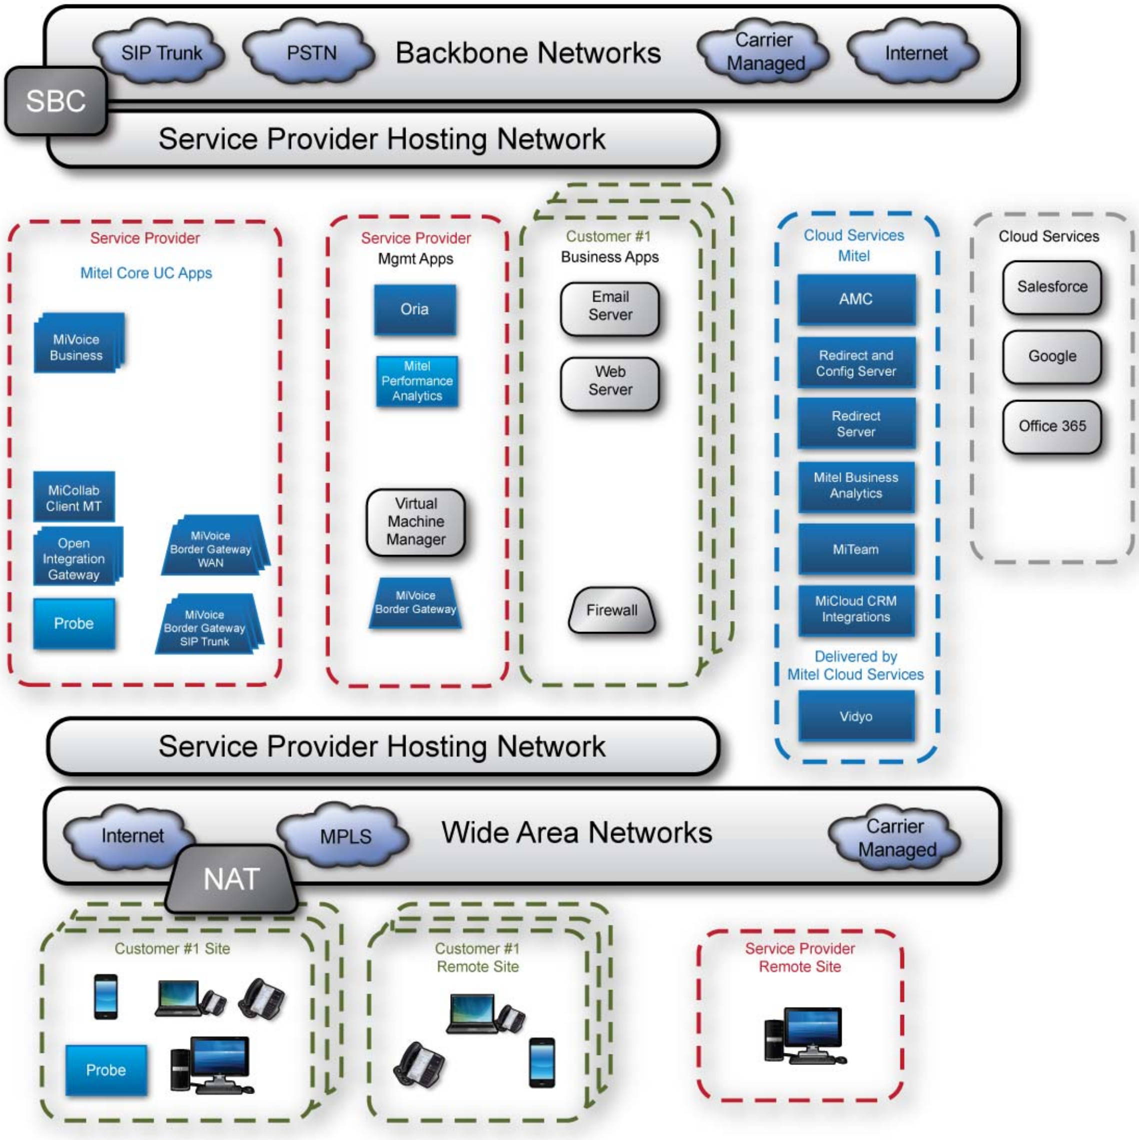

MiCloud Business Virtual

MiCloud Business Virtual architectures are based on MiVoice Business Virtual and MiCollab Virtual, on a VMware infrastructure.

Small Medium Business (SMB) architecture

The SMB architecture is designed to offer hosted UC services to end-customers in the small to medium-size business market segment. The small to medium business market segment covers customers with 30 to 500 users located at one site. The SMB architecture is suitable for end-customers with non-business-critical UC requirements.

Note: The SMB architecture is one of four MiCloud Business Solution topologies designed with a common set of components, the others being the SMB-LD, MLB, and Scalable topologies. Each of these solution architectures is capable of offering essentially the same UC capabilities.

The SMB architecture is based on the MiVoice Business Express platform. MiVoice Business Express integrates MiVoice Business, MiCollab, and MiVoice Border Gateway to provide call control, applications, and gateway functionality in a single virtualized unit. Mitel Open Integration Gateway is also available in this architecture to provide integration to related business applications.

MiVoice Business Express provides a configuration wizard for initial system set up. The wizard allows configuration of the basic settings such as configuring trunk main business numbers, SIP trunk proxy, and administration e-mails. Configuration for advanced features is achieved through use of the MiVoice Business System Administration Tool.

Oria provides end-user service provisioning. MarProbe provides fault and performance information. Billing is done through the third-party SIP trunk provider based on customer DID numbers.

Figure 4: Small Medium Business architecture

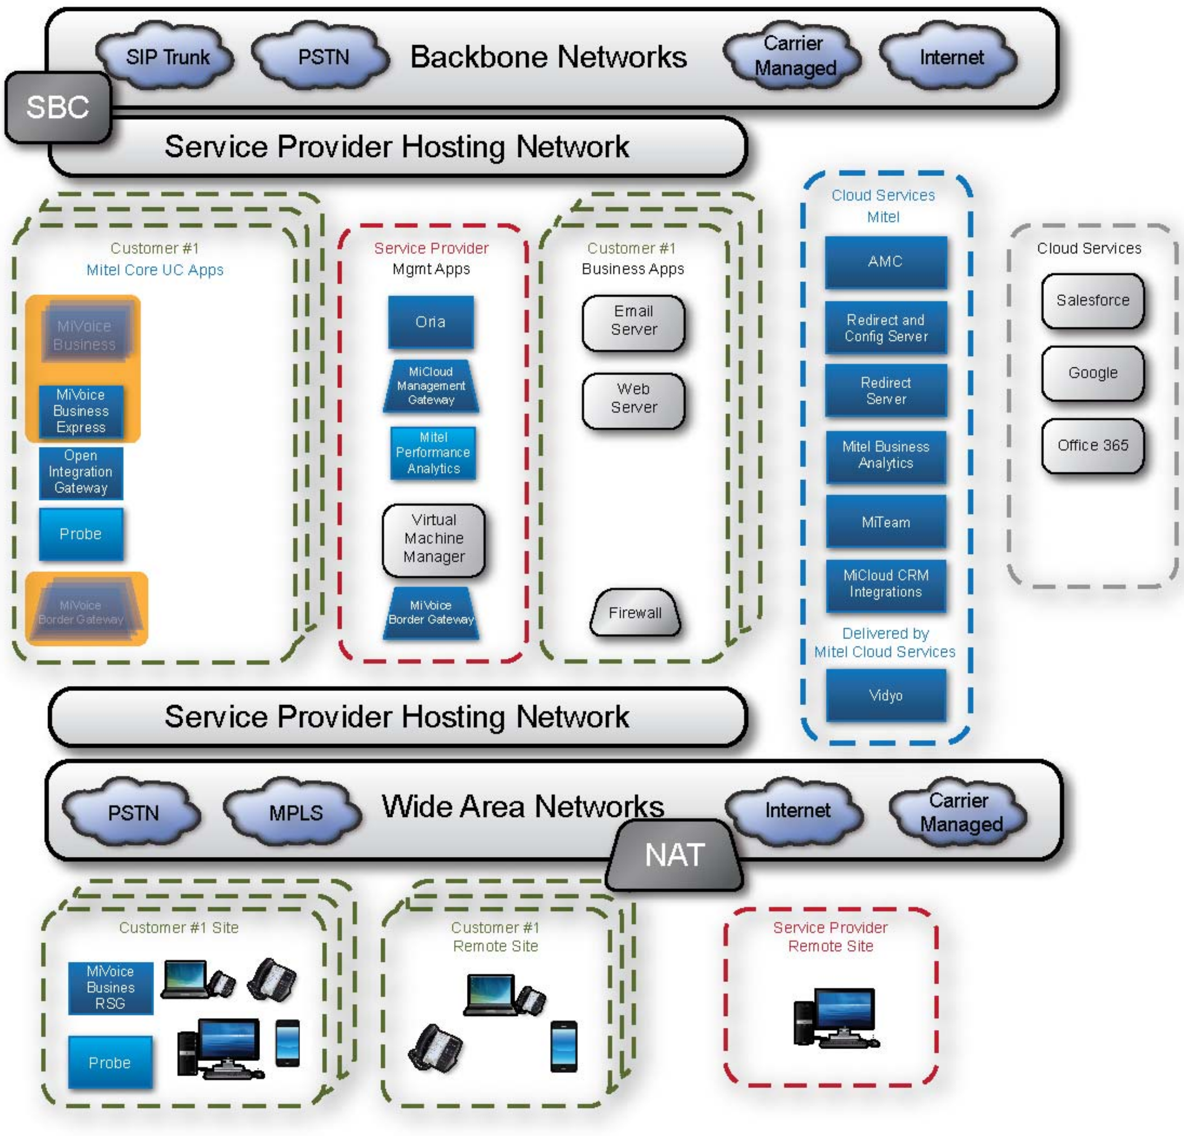

Note: The SIP trunks and Public Network (Teleworker) are terminated on the same MiVoice Border Gateway public IP address.

Small Medium Business, Low Density (SMB-LD) architecture

The SMB-LD deployment is based on the MiVoice Business Multi Instance platform, and uses a single flat address space for the service provider, with local isolation of the customer networks. A MiVoice Border Gateway or Session Border Controller (SBC) provides network isolation at the edge of the customer network. This MiVoice Border Gateway, or SBC, provides connectivity to external SIP trunk providers and customer Teleworker end-users.

Connectivity to the customer's remote networks is provided over a dedicated MPLS circuit. Extending the customer network into the service provider environment using MPLS connections removes the need to add a border gateway to isolate the end-customer networks from the hosted network. Extending the customer network into the hosted site allows increased UC functionality, compared to connections via a gateway. Extending the customer network also allows tighter integration with existing customer applications and other hosted services, and better connections to on-site equipment, such as a local trunk breakout.

Customer segregation is achieved through use of VLAN technology and a common VRF router onto the MPLS carrier network.

Because the SB architecture and the SMB-LD architecture both use the MiVoice Business Multi Instance platform, it is possible to combine both topologies onto the same hardware, while keeping the topologies logically separated. This separation is achieved by using unique VLANs per customer on the SMB-LD deployment and a common VLAN for the Over-the-Top (OTT) deployment of the SB architecture. Combing the physical infrastructure of the two topologies can provide cost savings, especially when using common management components, such as Oria and Mitel Performance Analytics servers and Probes.

Rich UCC functionality is provided by a dedicated MiCollab Virtual instance per customer. The MiCollab is dedicated to one customer and is not shared. The use of MPLS allows direct connection to the MiCollab Virtual without resorting to Over-the-Top (OTT) techniques, although a per-customer MiVoice Border Gateway is present, and can provide Teleworker and mobile device access on a per-customer basis.

Each customer also has its own dedicated MiVoice Border Gateway, which can be used to provide access to SIP Trunks, Teleworker users and SIP/UC users. Additional MiVoice Border Gateways can be used for scaling and resiliency.

Management access is provided by a common Oria portal in the service provider space, and through a dedicated per-customer management portal.

Figure 5: Small Medium Business, Low Density architecture

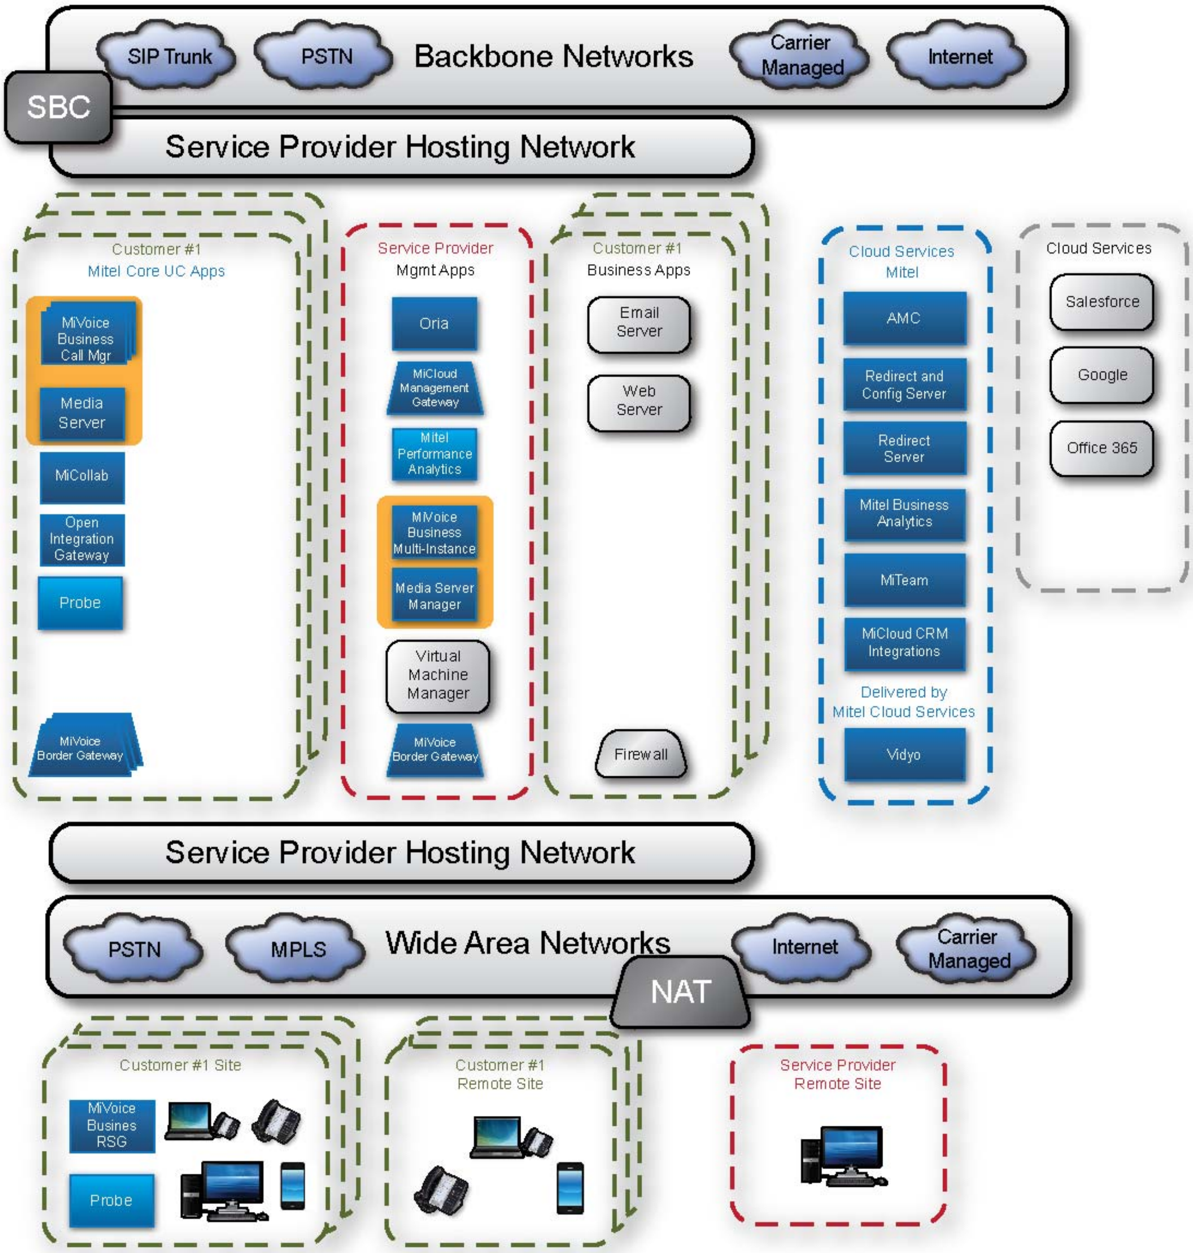

Medium Large Business (MLB) architecture

The MLB architecture is designed to offer hosted UC services to end-customers in the medium to large size business market segment. The medium to large business market covers customers with 250 to 1500 users, but the MLB architecture has the capability to scale to up to 5000 end-users.

The MLB architecture is also suitable for customers with 500 or fewer users, and whose business-critical communications needs justify the additional cost of a fully resilient system. The MLB architecture supports a number of service availability mechanisms, including MiVoice Business resiliency, MiVoice Border Gateway redundancy, and VMware High Availability.

Note: The MLB architecture is one of four MiCloud Business Solution topologies designed with a common set of components, the others being the SMB-LD, SMB, and Scalable topologies. Each of these solution architectures is capable of offering essentially the same UC capabilities.

The MLB architecture is built on the Mitel virtual appliance portfolio comprised of MiVoice Business Virtual, MiCollab Virtual, MiVoice Border Gateway Virtual, and Oria user service provisioning.

Figure 6: Medium-Large Business architecture

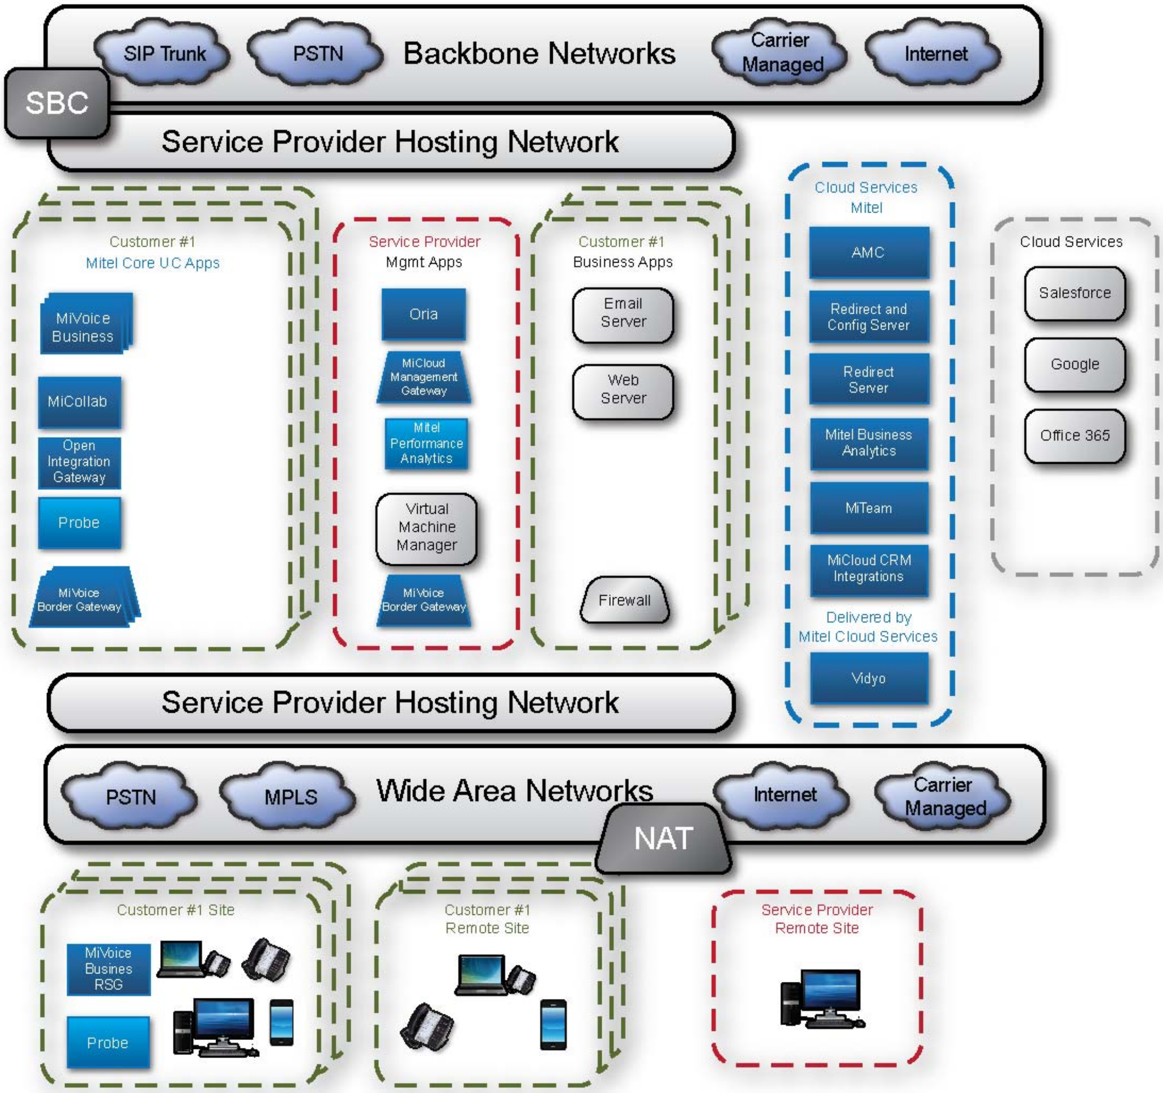

Scalable architecture

A service provider looking to address the requirements of an enterprise that has a large number of employees would typically use the MiCloud Business Solution Scalable architecture. This architecture is designed for a single customer, although the customer may have a number of offices that are geographically distributed; an enterprise with corporate headquarters and a number of regional offices, for example. Hosting the services in one or more data centers and linking the sites with a common network simplifies deployment and reduces costs, compared to a multi-site premise solution.

The Scalable architecture is a large and complex deployment. Deployment is managed, in part, by the enterprise IT staff, who may require training to fully understand the management requirements. Management is performed from inside the customer address space, with external access provided for support by the service provider and/or the reseller.

The Scalable deployment is based on a building block of up to 5000 UC users. It requires the identification of business units within the Enterprise, limited to 5000 users. Multiple business units are then linked together to scale to larger deployment sizes. The linkage requires application and federation connections between the UC components across all the business units, plus the addition of common functionality, such as centralized Auto-Attendant.

Each deployment requires a level of custom design to meet customer needs, and although an initial template is available, the end requirements may change that design. Other configurations or solutions may also be included for these larger deployments, especially to link to existing business processes and equipment; e-mail and messaging services, for example.

Contact Mitel Professional Services when considering a Scalable architecture deployment to ensure that the deployment meets customer requirements, links to existing business equipment, and experiences a timely roll-out of the deployment.

Figure 7 is provided as an example template for the MiCloud Business Solution Scalable architecture.

Figure 7: Scalable topology

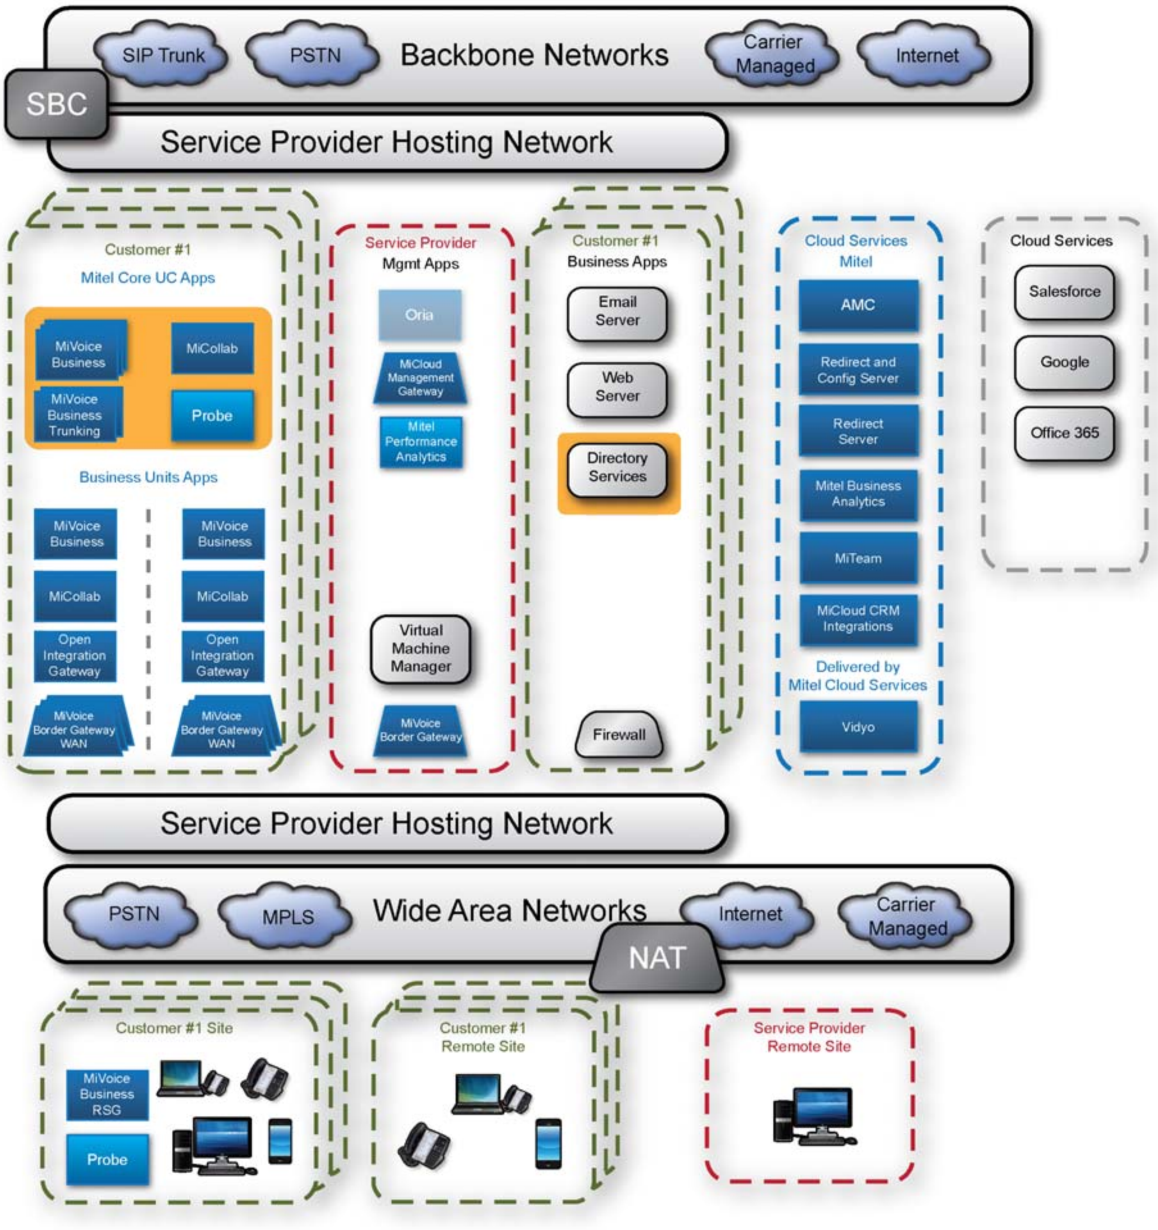

Chapter 2

AVAILABILITY AND RESILIENCY

Availability and resiliency

System reliability, redundancy, and resiliency are all inter-related and they have a direct effect on the availability of services. A non-redundant, non-resilient system provides users with a high level of availability; however, a higher level of service availability can be achieved with redundant and resilient design techniques.

Resilient systems provide higher levels of service availability than non-resilient systems by providing continued availability even when a system component fails.

For virtualized infrastructure, Mitel recommends deployment with VMware High Availability (HA) for improved resiliency, and VMware Distributed Resource Scheduler (DRS) to ensure optimal performance of the virtual machines. It is also important to set anti-affinity rules within the DRS configuration to ensure that primary and secondary elements are never allowed to run on the same physical host.

When access is restricted to VMware orchestration tools, there are typically three commercial offers from IaaS vendors, namely pay-as-you-go, reservation model, and allocation model. The first two are what the name implies. The allocation model has CPU reservations at the resource pool level but not the VM level. The Mitel virtual application portfolio should be deployed with the reservation model. For details, both for anti-affinity rules and VMware orchestration tools, see the *Virtual Appliance Deployment Solutions Guide*.

Additional information about meeting availability requirements and designing networks for availability is explained in a suite of Mitel white papers. For details of this doc suite, see *List of Telephone System Availability Document*. Log in to Mitel OnLine. Click **Solutions** > **Business Continuity Solutions** for these white papers. For details about MiVoice Business and IP telephone resiliency, see the *MiVoice Business Resiliency Engineering Guidelines*.

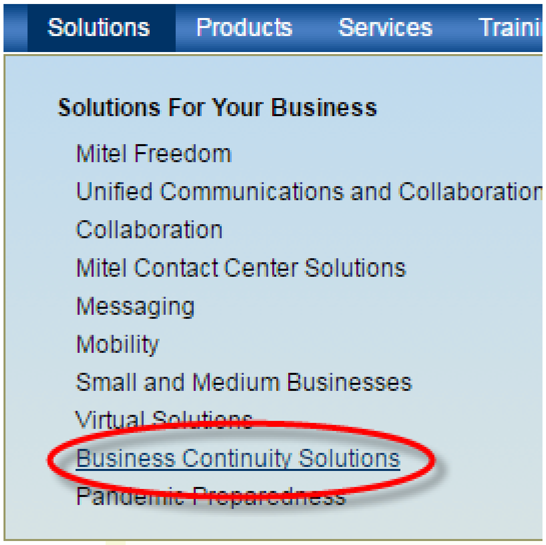
Solutions
Products
Services
Traini
Solutions For Your Business
Mitel Freedom
Unified Communications and Collaboratior
Collaboration
Mitel Contact Center Solutions
Messaging
Mobility
Small and Medium Businesses
Virtual Solutions
Business Continuity Solutions
Pandemic Preparedness

This chapter discuss the availability and resiliency mechanisms as they relate to the various components, applications, and services that are used to build the overall UC solution. This chapter is arranged in the following sections:

- "Mitel IP desk phones" on page 24
- "MiVoice Business" on page 26
- "Applications" on page 30
- "Networking and availability" on page 36
- "VMware and service reliability" on page 41
- "Determining system availability" on page 42

Mitel IP desk phones

Prior to deploying the UC solution, determined whether there are specific users and devices that require resilient operation.

If required, the MiVoice Business solution allows individual IP phones to be configured and licensed for resilient operation to meet the needs of users that require high availability telephony service.

The Administrator may choose to make all users and devices resilient, or only those required for critical services. The Administrator also must determine which MiVoice Business system each user and device will fail over to in the event that the primary MiVoice Business fails.

When a resilient IP phone is on an active call with another IP phone and its primary MiVoice Business system or the link between the phone and the MiVoice Business system fails, the phone operates in the following way:

- The current phone call survives, provided the network media path remains operational.
- The IP phone fails over to its secondary MiVoice Business system when the current phone call terminates.
- When the primary MiVoice Business returns to service, the IP phone recovers to the primary MiVoice Business.

For information about planning for resilient operation, network design, and configuring users for resilient operation, see the *MiVoice Business Engineering Guidelines* and the *MiVoice Business Resiliency Engineering Guidelines*.

MiVoice Business Console

The MiVoice Business Console is an IP console that supports resilient operation.

- Like other resilient IP phones, 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, then 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 has ended and the console is in the idle state.
- When the primary MiVoice Business returns to service, the IP Console recovers to the primary MiVoice Business.

For MiVoice Business Release 7.0+, when an IP console is connecting to the primary and the secondary MiVoice Business controllers/servers through a MiVoice Border Gateway connection, if the MiVoice Border Gateway fails, the console does not failover to the secondary MiVoice Border Gateway.

However, if the communication link between the MiVoice Border Gateway and the primary MiVoice Business fails, or the primary MiVoice Business fails, the MiVoice Border Gateway fails over to the secondary MiVoice Business, which allows the IP console to communicate with the secondary MiVoice Business.

MiVoice 5540 IP Console

The MiVoice 5540 IP Console is an IP console that supports resilient operation. Like other resilient IP phones, if a resilient MiVoice 5540 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, then 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 has ended and the console is in the idle state.

When the primary MiVoice Business returns to service, the IP Console recovers to the primary MiVoice Business.

The MiVoice 5540 IP Console supports Power over Ethernet (IEEE 802.3af) operation, which allows the console to be remotely powered from a PoE capable L2 switch. For enhanced availability, the L2 switch could be powered via a UPS.

MiVoice Business

MiVoice Business availability and resiliency is discussed in the following topics:

- "MiVoice Business hardware platforms" on page 26
- "MiVoice Business resiliency" on page 27
- "MiVoice Business as survivable gateway" on page 28

MiVoice Business hardware platforms

The MiVoice Business call control software may be installed on a number of different hardware platforms, including 3300 ICPs, Industry Standard Servers running Mitel Standard Linux (MSL), or Industry Standard Servers running VMware. For information beyond what is provided here, refer to the MiVoice Business product documentation.

Mitel 3300 ICP

The 3300 ICP platforms are Mitel proprietary hardware platforms used for running the MiVoice Business software. There are a number of 3300 ICP platforms available to meet different user scaling and analog/digital PSTN connectivity requirements. The 3300 ICP can also provide the following functionality:

- IP to PSTN gateway (analog and digital/PRI)
- Analog and TDM phone connectivity
- Voice mail
- Music on Hold
- Ad-hoc conferencing
- Basic call recording

Industry Standard Servers (ISS)

MiVoice Business can be installed on a number of Mitel-approved Industry Standard Servers (ISS). ISS platforms offer higher call processing performance and faster fault recovery times than the 3300 ICP platforms.

When compared to 3300 ICP platforms, ISS platforms can offer improved availability of hardware components that are typically known to exhibit lower levels of reliability. These include use of multiple power supplies, multiple hard disk drives with RAID configurations, RAM with parity and error correction mechanisms, and multiple redundant cooling fans.

For large service provider deployments, there are additional improvements that can be achieved through the use of chassis-based technologies rather than traditional "rack and stack" ISS platforms. Improvements include advanced integrated management, which provides advance warning of a pending failure, and hot swap capabilities for power supplies, fans, and hard drives, which allow a component change to be performed without taking the server out of service.

Chassis-based technologies may support some resilient networking and switching functions, and may include link aggregation capabilities, multiple redundant switching backplanes, and multiple external connections, which can simply the deployment of systems that must be resilient.

To determine whether a particular ISS is supported, refer to the Mitel Qualified Hardware List, available on Mitel OnLine.

MiVoice Business resiliency

While failures and outages in any complex system are unavoidable, the MiVoice Business call control resiliency solution provides the ability to preserve telephony service in the event that a MiVoice Business goes out of service or network connectivity to the MiVoice Business has failed. The MiVoice Business resiliency capability is available on all MiVoice Business hardware platforms.

For additional information related to resilient operation, see the *MiVoice Business Engineering Guidelines*, the *MiVoice Business Resiliency Engineering Guidelines* and the *Virtual Appliance Deployment Solutions Guide*, available on Mitel OnLine.

Resilient operation description

Resilient call routing handles cases in which a phone is in service on its secondary MiVoice Business system or in transition between MiVoice Business systems. When configured in a resilient cluster, all MiVoice Business systems 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 primary or secondary systems.

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. An IP phone does not have to be resilient to experience call survival because maintaining a previously established voice path between two IP phones is not dependent on the availability of the MiVoice Business system. When a call that was in survival mode has ended, a non-resilient device goes out of service because it cannot reach its MiVoice Business system, while a resilient device fails over to its secondary MiVoice Business system.

During a MiVoice Business system or network failure, both non-resilient and resilient devices (that are currently in talk state), experience call survival, but because they have lost the link to their MiVoice Business they cannot access phone features or dialing functionality.

Only calls in talk state are capable of call survival in a failure situation. Calls that are in the process of being set up or are in feature transition, for example, do not survive.

Resilient Call Control features

While in service on a secondary MiVoice Business, an IP phone retains basic call service. Most call features are available while a phone is in service on a secondary system, some with possible behavior differences. For details refer to the *MiVoice Business Resiliency Engineering Guidelines*, specifically the chapter on feature resiliency.

MiVoice Business as survivable gateway

In some cases, the customer may require the ability to access the local PSTN from the customer premises if there has been a service interruption related to the WAN connectivity between the customer site and the data center. PSTN access may be also required for legacy business purposes and for emergency services. These requirements can be satisfied by deploying a 3300 ICP platform running MiVoice Business at the customer site. The 3300 ICP can provide trunking capabilities to connect to the local PSTN, and process incoming and outgoing calls.

This is discussed in the following topics:

- "MiVoice Business for Industry Standard Servers (ISS)" on page 28
- "MiVoice Business Multi-Instance" on page 28
- "MiVoice Business Virtual" on page 29

For more information about using a 3300 ICP as a survivable gateway, see "External Connectivity" on page 69, the MiVoice Business product documentation, and the *Mitel Survivability of Remote Branch Offices Solutions Guide*.

MiVoice Business for Industry Standard Servers (ISS)

MiVoice Business for ISS offers the same resiliency capabilities as 3300 ICP. When compared to 3300 ICP platforms, however, ISS platforms offer improved availability of hardware components that are typically known to exhibit lower levels of reliability.

MiVoice Business for ISS works with SIP trunks and can handle multiple routes. This offers improved trunking availability compared to a 3300 ICP. For instance, if a 3300 ICP loses a physical trunk connection, it may result in the loss of an external trunk connection, but with MiVoice Business on ISS, SIP trunks can be routed to alternative paths.

MiVoice Business Multi-Instance

MiVoice Business Multi-Instance is a product that allows numerous instances of MiVoice Business to be run on a single ISS. The ISS has the same hardware availability characteristics as MiVoice Business on ISS. Each MiVoice Business instance on a MiVoice Business Multi-Instance supports the same resiliency capabilities as MiVoice Business.

When deploying a resilient MiVoice Business Multi-Instance solution, the IP sets are configured with a primary MiVoice Business running on one physical server, and a secondary MiVoice Business running on a different physical server.

While MiVoice Business Multi-Instance allows individual instances of MiVoice Business to be run, a tenanted model is not supported. This improves service availability since it is possible to reset a single instance of MiVoice Business without affecting the other MiVoice Business instances running on the same server platform.

MiVoice Business Virtual

The MiVoice Business Virtual offering is a version of the MiVoice Business software that is packaged in Open Virtualization Archive (OVA) format for deployment in a VMware environment.

MiVoice Business Virtual offers the same resiliency capabilities as MiVoice Business, specifically, primary node failure detection and automated failover to a secondary node.

In addition to the resilient operation capabilities of MiVoice Business, VMware can offer additional resiliency capabilities. This is accomplished with VMware High Availability (HA), used with VMware vCloud Center for automated operation.

VMware HA is a service that automatically detects a physical server failure or a virtual machine failure, and restarts virtual machines on alternative servers when a failure is detected. In the case of MiVoice Business Virtual, VMware HA detects a failure on the server hosting the MiVoice Business Virtual or the failure of the virtual machine running MiVoice Business. Upon host failure, VMware HA starts a new instance of the MiVoice Business Virtual on an alternate server host. Upon failure of a virtual machine, VMware HA starts a new instance of MiVoice Business Virtual.

When combined, the resiliency capabilities of MiVoice Business Virtual and VMware HA capabilities provide a solution that offers very high service availability. For example, should the server that is hosting the primary MiVoice Business Virtual experience a failure, the phones failover to the secondary MiVoice Business Virtual and service is maintained. Meanwhile, HA has also detected that the primary MiVoice Business Virtual has failed, and it starts the primary MiVoice Business Virtual on an operational server. After the primary MiVoice Business Virtual is back in service, the phones recover (return) to the primary MiVoice Business Virtual.

Without HA, the system would have been operating in a non-resilient mode until action was taken to repair the primary MiVoice Business Virtual and restore it to service.

The major benefit provided by VMware HA is that it reduces the recovery time of a failed unit, and automates the recovery process. This greatly reduces the Mean Time To Recovery, which would normally involve a technician call-out and possible travel time charges.

Certain policies must be adhered to when using the resiliency capabilities of MiVoice Business Virtual. For instance, anti-affinity rules should be established to ensure that the primary MiVoice Business Virtual and secondary MiVoice Business Virtual are never deployed on the same physical host server in a VMware cluster.

For more information, see the *MiVoice Business Resiliency Engineering Guidelines*. Information about deploying the Mitel virtual appliances, including MiVoice Business, is available in the *Virtual Appliance Deployment Solutions Guide*.

Applications

This section discusses the resiliency capabilities of the various applications available in the Unified Communications solution. For additional information, refer to the specific product documentation.

- “MiCollab Unified Messaging voice mail” on page 30
- “MiCollab” on page 31
- “Mitel Open Integration Gateway” on page 33
- “MiVoice Business Express” on page 33
- “MiContact Center Business” on page 33
- “MiVoice Call Recording” on page 34
- “MiCloud Business Analytics” on page 34
- “Vidyo” on page 35
- “MiVoice Border Gateway” on page 34
- “MiCloud Management Gateway” on page 35

For information about deploying on VMware, see “VMware and service reliability” on page 41 and the *Virtual Appliance Deployment Solutions Guide*.

MiCollab Unified Messaging voice mail

The MiCollab Unified Messaging is an e-mail application feature of MiCollab. MiCollab Unified Messaging offers users the ability to manage their voice mail, e-mail, and fax messages from their PCs or telephones. MiCollab Unified Messaging also allows inbound callers quickly find the person they need to talk with using a speech-enabled auto attendant and call routing functionality.

Like an IP set, the MiCollab Unified Messaging application supports MiVoice Business resilient operation. Specifically, the MiCollab Unified Messaging system has the ability to fail over to a secondary MiVoice Business when the primary MiVoice Business is non-operational or unreachable, and then to recover to the primary MiVoice Business when it returns to service.

The MiCollab Unified Messaging application does not natively support resiliency of the application itself. When higher levels of service availability are required:

- The use of MiCollab Virtual should be considered because MiCollab Virtual can take advantage of the additional resiliency capabilities provided by the VMware High Availability (HA) feature.
- MiCollab Unified Messaging storage should be located on a different server than the MiCollab server. Typically, storage is located in a SAN, and access to the SAN will use redundant connections.

For information about resiliency, programming for resilient operation, and SAN technology, refer to the following documents.

- *MiCollab Unified Messaging Technician's Handbook*
- *MiVoice Business Resiliency Engineering Guidelines*
- MiVoice Business System Administration Tool Help
- *Virtual Appliance Deployment Solutions Guide*

MiCollab

MiCollab is a comprehensive, integrated solution that unifies business communications applications. MiCollab supports any combination of the following applications:

- MiCollab Unified Messaging
- MiCollab Client Service
- MiCollab Audio, Web, and Video Conferencing
- MiVoice Border Gateway

MiCollab is available on the following hardware platforms:

- Industry Standard Server (non-resilient)
- Virtual Appliance (resilient capable)

Higher service availability is supported when MiCollab is deployed as a virtual appliance in VMware vSphere. This product offering is called MiCollab Virtual.

For information about deploying on VMware, see "VMware and service reliability" on page 41 and the *Virtual Appliance Deployment Solutions Guide*.

MiCollab Client Multi-Tenant

MiCollab Client Multi-Tenant is a product that combines the call control capabilities of Mitel communications platforms with contact management, dynamic status, and presence information to simplify and enhance communications. MiCollab Client Multi-Tenant is available as a component of MiCollab.

If a PC soft phone is deployed on a desktop or workstation, it can connect to the MiCollab Client Multi-Tenant server directly via the LAN. If, however, the PC soft phone is installed on a laptop or a tablet that may go outside of the range of the business network (Wi-Fi or LAN), then the PC soft phone should always register via the MiVoice Border Gateway, even when on the local network. Registering with the MiVoice Border Gateway allows for hand-over when switching between business and external networks. This improves network connectivity and availability.

The mobile soft phone should always connect to the external MiVoice Border Gateway, even if it uses the internal business network to reach the business Internet gateway.

MiCollab Client Multi-tenant Service availability

The MiCollab Client Multi-tenant server must be operational for MiCollab Clients to be able to place calls. The MiCollab Client Multi-tenant server is not resilient, which means that if the MiCollab Client Multi-Tenant server fails, then any device that connects to the MiCollab Client Multi-Tenant server, such as mobile phones and PC Softphones, lose telephony connectivity.

MiCollab for Mobile, the new mobile soft phone is designed for improved resiliency.

- In the case of MiVoice Business server failure, the MiCollab server connection to MiVoice Business fails over to the secondary MiVoice Business, and recovers when the primary is back in service. Mobile clients are supported on the secondary controller until the primary recovers.
- In the case of MiVoice Border Gateway server failure in one or more clustered MiVoice Border Gateways, the MiCollab for Mobile client registers with an alternate MiVoice Border Gateway in the cluster. This resiliency is achieved using DNS Service (SRV) records; the FQDN is mapped to multiple host names. For further information about configuring the DNS record, MiCollab Client, and MiVoice Border Gateway, see the appropriate MiCloud Business Deployment Guide (Multi-Instance or Virtual) and the *MiCollab Client for Mobile Resiliency Guide*.

For the MiCollab UC Client, the legacy mobile soft phone, users who require higher levels of availability should deploy both a mobile soft phone and EHDA on the smart phone device. The users should be deployed with EHDU in the UCC profile. This provides a secondary and alternative carrier connection back to the system, and allows mobility outside the office across different networks. MiVoice Business resiliency is supported using VMware High Availability (HA) functionality when the MiCollab virtual application is deployed on VMware. The MiCollab UC-Client does not support resilient behavior for MiVoice Border Gateway failure.

MiCollab Audio, Web, and Video Conferencing

MiCollab Audio, Web, and Video Conferencing is a comprehensive audio conferencing and web collaboration application that improves collaboration and information sharing among employees and with customers, partners, and suppliers. MiCollab Audio, Web, and Video Conferencing is available as a core component of MiCollab. MiCollab Audio, Web, and Video Conferencing provides no native application resiliency; it relies on MiCollab service availability.

MiVoice Border Gateway

MiCollab includes MiVoice Border Gateway within the application suite. For UC solutions, the MiVoice Border Gateway inside MiCollab is configured as the master instance of the MiVoice Border Gateway cluster. In the case of MiCollab failure, the remaining units in the MiVoice Border Gateway cluster continue to function while the master MiVoice Border Gateway recovers.

See “MiVoice Border Gateway” on page 34 for a detailed description.

Mitel Open Integration Gateway

Mitel Open Integration Gateway (OIG) is a software solution that executes on an industry standard server. Mitel Open Integration Gateway offers web services to applications, and integrates with web-based services such as Google and Salesforce. For detailed information about services, and how an application communicates with Mitel OIG, see the Mitel OIG developer guides.

Mitel OIG does not have a resiliency mechanism that employs a primary and secondary Mitel OIG sever; however, Mitel OIG does support inter-operation with MiVoice Business systems that are configured for resilient operation. For instance, if a Mitel Open Integration Gateway system is configured to operate with a resilient pair of MiVoice Business systems, and the primary MiVoice Business fails or becomes unreachable, the IP phones that were homed to the primary MiVoice Business failover to the secondary MiVoice Business, and Mitel OIG now provides services to the IP phones via the secondary MiVoice Business.

If Mitel OIG application resiliency is required, consider using Mitel OIG Virtual on VMware, and using the VMware High Availability (HA) feature.

MiVoice Business Express

MiVoice Business Express provides a complete communications solution for small to medium size businesses. MiVoice Business Express runs as virtual appliance on a VMware vSphere infrastructure. One level of resiliency can be provided using the VMware High Availability feature.

MiVoice Business Express 7.1+ supports MiVoice Business resiliency when clustering the internal MiVoice Business with an external MiVoice Business controller.

The MiVoice Business Express contains the following components:

- MiVoice Business
- MiCollab: provides the following applications:
  - MiCollab Unified Messaging
  - MiCollab Client Service
  - MiCollab Audio, Web, and Video Conferencing
  - MiVoice Border Gateway

For information about deploying on VMware, see “VMware and service reliability” on page 41 and the *Virtual Appliance Deployment Solutions Guide*.

MiContact Center Business

The MiContact Center Business application, with the exception of IVR, does not support resilient operation. MiContact Center Business can use the VMware High Availability capability to provide cold standby. In addition, MiContact Center Business Release 8.0+ provides internal monitoring for key routines and enables broader and faster auto-recovery to increase application robustness and availability.

Resilient Contact Center IVR operation is available using multiple IVR servers. MiContact Center Business Enterprise server may host one of the IVR instances or multiple remote IVR servers may be deployed. Each instance is configured with active and redundant ports. In the case of one of IVR server being out of service, the redundant ports carry the full traffic load. The IVR servers use a local data cache with a mirror of the Enterprise Server configuration to continue routing calls if the MiContact Center Business server is unavailable.

MiContact Center Business supports MiVoice Business resiliency. In the case of MiVoice Business service interruption, the IVR routing instance fails over to the secondary MiVoice Business and continues routing calls through the secondary controller.

MiVoice Call Recording

The MiVoice Call Recording application does not support resilient operation. MiVoice Call Recording can use the VMware High Availability feature to provide cold standby.

MiVoice Call Recording does support MiVoice Business resiliency and MiVoice Border Gateway-SRC resiliency. In the case of MiVoice Business service interruption, audio for active calls continues to be recorded, although without the associated MiTAI events, so there will be incomplete meta-data.

In the case of MiVoice Border Gateway-SRC service interruption, audio of active calls is terminated. Subsequent calls progress on the resilient MiVoice Business and/or MiVoice Border Gateway-SRC.

MiCloud Business Analytics

Mitel Business Analytics is a cloud-based service. The underlying servers support resilient operation.

The Delivery Controller may be deployed with 1:1 redundancy. However, in most cases, this is not cost justified. In the case of WAN failure, the Delivery Controller buffers SMDR preprocessed records. If the buffer fills, the Delivery Controller shuts down the TCP streaming, triggering the MiVoice Business to buffer SMDR records. In the case of Delivery Controller failure, the MiVoice Business buffers the SMDR records. If Mitel Business Analytics fails to receive SMDR information, an alarm is triggered.

The Delivery Controller supports MiVoice Business resiliency. The Delivery Controller is configured with both the primary and secondary MiVoice Business addresses. If the primary MiVoice Business is out of service, the secondary MiVoice Business controller streams SMDR records. When the primary recovers, the primary controller resumes streaming SMDR records.

MiVoice Border Gateway

MiVoice Business Gateway is available in several variants:

- Standalone, on Industry Standard Server or VMware infrastructure
- With MiCollab, see “MiCollab” on page 31
- With MiVoice Business Express; see “MiVoice Business Express” on page 33

The MiVoice Border Gateway ensures the deployment of secure internal and external work spaces, enabling remote workers, road warriors, and day-extenders seamless access to the voice and data capabilities of the office, wherever they are.

MiVoice Border Gateways can be clustered for license sharing and load balancing for devices that support the MiNet protocol. For SIP device connectivity and SIP-Trunk connectivity Border Gateways are typically clustered as primary/secondary pairs.

Supported MiNet sets can hold IP addresses for two MiVoice Border Gateways. If a set loses its connection to the MiVoice Border Gateway and it cannot re-establish the connection, it tries the second IP address on its list. When MiNet sets are configured for persistent resiliency lists, the resiliency list is retained through a power cycle.

If MiVoice Border Gateway high availability service is required, the MiVoice Border Gateway Virtual offering should be considered.

The VMware High Availability (HA) solution is used to provide application-level resiliency. VMware HA is the recommended deployment mode when used with UC to provide a common IP address that the UC clients can register with. For more information, see the MiVoice Border Gateway product documentation.

MiCollab includes MiVoice Border Gateway within the application suite. For UC solutions, the MiVoice Border Gateway inside MiCollab is configured as the master instance of the MiVoice Border Gateway cluster. In the case of MiCollab failure, the remaining units in the MiVoice Border Gateway cluster continue to function while the master MiVoice Border Gateway recovers.

MiCloud Management Gateway

MiCloud Management Gateway does not support resilient operation, but as a virtual appliance, it supports the VMware High Availability feature.

Vidyo

Vidyo provides a cloud-based service. The underlying servers support hot standby resiliency.

Networking and availability

Network connectivity between the customer premise and the data center, and within the customer premise, is often the weakest link in a UC deployment. Secondary connections and dual routers should be considered for deployments that are sensitive to service outage time. Service Level Agreements (SLA) should be considered when signing with a carrier, because the network link from the customer premise to the data center is critical to the business. Such links are not easily achieved over the Internet, so the carrier should be selected with care, and it should line up with connectivity from the hosted data center.

There are many factors related to network design that must be planned to ensure that the UC solution components can reliably communicate with each other in the event that a network-related failure occurs. This section discusses these factors in the following sections:

- "Trunking Considerations" on page 36
- "WAN and OTT considerations" on page 37
- "LAN considerations" on page 38

Trunking Considerations

IP Trunks are used to connect multiple MiVoice Business controllers together. SIP trunks may be used to connect the customer premises to the data center, and to a service provider who, in turn, passes connections on to the PSTN. TDM trunks may be used to connect MiVoice Business controllers to the PSTN, or to connect multiple MiVoice Business controllers together.

IP trunking, ARS, and Hunt Groups

A well-designed IP network is partially meshed, meaning that there is always more than one IP path between IP networking devices. An IP network that offers multiple connection paths is inherently resilient.

When MiVoice Business controllers are introduced into a partially-meshed IP network, IP trunks in the network are resilient at the physical level. Additional trunk resiliency can be provided by using TDM trunks. Use ARS and secondary routing to configure alternative paths through the network to minimize a single point of connection failure.

SIP Trunking

Service providers offer SIP trunks that provide flexible and cost-effective solutions for connecting the customer premises to the data center. SIP trunking can be used to connect the customer premises to multiple service providers, and the SIP trunking solution can be used to provide trunking resiliency.

For SIP Trunks, a gateway unit or Session Border Controller (SBC) function is required to isolate the customer network from the carrier network. The MiVoice Border Gateway provides this gateway functionality, along with a number of SIP SBC functions. The MiVoice Border Gateways can be deployed in pairs for redundant operation.

Primary and secondary SIP trunk registration with primary and secondary SBCs can be used to a common service provider to provide 100% failover in the event of a lost SIP trunking gateway or connection.

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 increases if these TDM paths are employed as back-up paths for IP trunks.

WAN and OTT considerations

Compared to all of the other components that comprise the UC solution, the external network connections are typically more likely to fail. Keeping this in mind, the availability of all WAN links and Over-the-Top (OTT) connections must be taken into account.

The availability of WAN connections can be improved by:

- Choosing a service level agreement (SLA) that meets the availability requirements.
- Using dual WAN connections to provide WAN redundancy.
- Using routers to implement WAN connections so that routing protocols such as HSRP or VRRP can be employed to detect faults and switch to alternate paths.
- When using dual connections, ensuring that the cabling for the two connections is not sharing the same physical route to eliminate the risk of both links being accidentally severed at the same time.
- When using dual connections, use two different carriers in two different geographic locations, for carrier resiliency.
- Consider using an Over-the-Top connection as a back-up to the WAN connections

For high service availability, consider the installation of additional network connections at the customer premises. For some network deployments, those that do not employ OTT connections, the addition of a survivable gateway may be an alternative option.

Remote site survivability

There are varying degrees of functionality that might be required at a customer premise or at a remote if the connection to the data center fails. Some customer sites may require that certain applications or functions be fully operational, even if the communication path to the head office fails, If this is the case consider:

- Physically locating application infrastructure and application servers at the customer site so that these applications can operate independently of the data center, (DHCP server, for example).
- Locating a 3300 ICP on site to serve as a survivable PSTN gateway.
- Providing local connectivity to external networks such as the PSTN and the Internet, so that these networks can be accessed independently from the customer premise location.

- Consider providing back-up power to critical pieces of equipment at the customer premise.

Direct connections to the PSTN network can be realized with 3300 ICP platforms. The preferred 3300 ICP platforms are:

- MXe III
- CX II

As a minimum, even in cases where remote site survivability is not a requirement, ensure that support for emergency calls is provided using local connections to the PSTN.

For more information, see "Emergency Location (E911)" on page 123 and the *Mitel Survivability of Remote Branch Offices Solutions Guide*.

LAN considerations

If a high availability UC solution is required, it is critical that the LAN be designed to be resilient. LAN resiliency ensures that traffic can be redirected to the desired destination if a networking switch or router fails. The LAN must also be designed for sufficient available bandwidth, even during network equipment failures.

Quality of Service (QoS) mechanisms must be employed when designing the network so that different types of traffic are treated with the appropriate priorities across the network.

For additional information related to network design for high availability, see the following resources:

- Mitel White Paper: *Network Design for Availability*
- *MiVoice Business Engineering Guidelines*
- *MiVoice Business Resiliency Engineering Guidelines*

LAN architecture

There are a number of LAN topologies that may be employed when designing a LAN; however, it is highly recommended that the LAN be based on a hierarchical design.

Hierarchical networks are typically designed with three layers; core, distribution, and access layers.

- The core equipment is located in a server room or data center room. The core connects distribution groups to each other, and provides connectivity to servers and gateways.
- The distribution layer equipment is located in wiring/telecommunications closets, and connects to the core and access layer equipment through high-speed links.
- The access layer is located either in the wiring closets or physically close to the work groups that require connectivity. The access layer connects end-user devices to the distribution layer.

Note: Connections from the access layer to the customer or end-point are not redundant. However, in critical situations, such as an emergency call center, it is possible agents may have dual phones and dual connections to different physical access layer switches.

Hierarchical networks allow similar users and their resources to be grouped together into specific work groups. This allows:

- A work group's traffic to be contained in a local area and to be connected to a common Layer 2 network.
- Resources dedicated to a particular work group, such as printers, to be connected to the Layer 2 access switch for the work group.
- Containment of work group traffic and resources to each work group. This allows key work groups to be provisioned with a higher level of availability than other, less critical, work groups. This strategy provides cost savings where higher availability is not required.

Hierarchical network designs easily overlay onto structured building wiring plants. This allows for:

- Less costly installations due to minimal re-wiring requirements.
- Easier moves, adds, and changes, and more efficient network maintenance.
- Easier scalability, when it is necessary to expand the call center size.
- Faster troubleshooting of network faults.
- Prevention of wide-spread network disruptions when a network fault occurs, by containing the disruption to a localized area.

In some cases. the UC solution may have different availability requirements at different levels in the network hierarchy. A hierarchical network allows the network designer to tailor the level of availability to meet these different requirements.

Redundancy mechanisms in LANs

Some specific examples of redundancy mechanisms that can be used when designing a high availability contact center LAN are:

- Use of Layer 2 and Layer 3 networking equipment that supports full redundancy.
- Duplication of Layer 2 switches and Layer 3 networking devices throughout the LAN.
- Duplication of transmission paths using partial mesh networking to support redundant communication paths.
- Use of Layer 2 protocols that enable the use of duplicate network paths as back-up paths, specifically; STP, RSTP, and MSTP.
- Use of two DHCP servers
- Duplication of storage devices; SAN, NAS, and RAID arrays, for example.
- Use of resilient topologies (implies hierarchical network design)

- Use of Layer 3 protocols that enable the use of duplicate network paths as back-up paths; specifically, OSPF, VRRP, or Cisco HSRP.
- Multiple NICs on servers, and use of LACP (IEEE 802.3ad).

Ideally, a network that has been designed for high availability does not contain any single points of failure, and when a failure is encountered, the network re-establishes functionality as quickly as possible.

Network power provisioning

The critical elements of the UC solution should retain power in the event that the mains power feed fails. This can be achieved through an alternative power source such as a secondary mains supply, local UPS, or local generator.

Primary network devices (MiVoice Business controllers, Layer 2 switches, and servers) should be powered from different branch circuits than the circuits used for powering the secondary network devices.

Critical phones should also be provided with back-up power. If the phones are powered locally, then this is required at each phone. It might be prudent, therefore, to use a Layer 2 switch that has Power over Ethernet (PoE) capabilities. The PoE capable Layer 2 switch could then be powered with a common back-up power mechanism such as a UPS, for use in a power outage.

For additional information related to PoE planning, IP phone power provisioning, and network power provisioning, see the following documents:

- *MiVoice Business Engineering Guidelines*
- *MiVoice Business System Engineering Tool* (SET) - This tool has an embedded IP phone power calculator.
- Mitel Data Sheets for 3300 platforms and associated hardware
- *Mitel 3300 ICP Hardware Technical Reference Manual*
- Manufacturer’s data sheets for third-party servers, Layer 2 switches and routers

VMware and service reliability

VMware offers an extensive suite of software solutions that allow customers to create a virtual IT infrastructure. Mitel virtual products are intended to run on the VMware virtual infrastructure and leverage the capabilities that VMware provides. These include the ability to provide application-level resiliency for customers that require high service availability.

VMware High Availability (HA) and VMware Site Recovery Manager (SRM) are two VMware features that are designed to increase the availability of systems running in a VMware environment. If a Mitel virtual application has no native resiliency capabilities, VMware HA and SRM may be used to provide a higher level of system availability.

For those Mitel applications that offer native resiliency capabilities, VMware HA and SRM work with the application's resiliency capabilities to further increase system availability.

VMware High Availability is a failover protection mechanism that is used to recover a virtualized application during hardware or operating system failures. With VMware HA, the servers are geographically co-located; switching time between servers is approximately 15 minutes. VMware HA provides the following capabilities:

- HA detects operating system and hardware failures
- HA restarts the virtualized application on another physical server in the resource pool without manual intervention when a server failure or operating system failure is detected.

When deploying applications with VMware HA, the application must be configured to use a Storage Area Network (SAN) storage, rather than local hard disks on the host. That way, if a host fails, the virtual machine (VM) can be started on another host and attached to the same virtual hard disk.

VMware SRM (Site Recovery Manager) offers customers a disaster recovery mechanism that can also be used for planned migrations. SRM can manage failover from a production data center to a disaster recovery site based in a geographically different location. When used to perform a complete data center recovery, SRM can provide RPOs (Recovery Point Objectives) of several hours to several days.

For information about deploying in a VMware environment, see the *Virtual Appliance Deployment Solutions Guide*.

Note: To deploy the Site Recovery Manager capability with Mitel virtual applications, Mitel Professional Services must be purchased. Request a quotation for SRM support from Services Solutions at the following URL:

http://domino1.mitel.com/mol/servsol.nsf/ServSolApp?OpenForm

You must be logged in to Mitel OnLine to use this request form.

For more information about VMware and its products, visit the VMware web site.

Determining system availability

It is common practice in the telecommunications industry to refer to the reliability level of a product, or a system, as 5-9s or 5 x 9. What is actually being referred to is the availability of a service expressed as a percentage for a particular product or system. When a product has an availability of 5-9s it is available 99.999% of the time.

Determining the availability of a particular UC solution is a complicated task. Hardware products such as a server can achieve better than 5-9s of availability. But it is important to remember that the overall system availability is defined by the weakest link in the chain.

When determining overall system availability, Mitel uses a seven-layer business continuity model. This business continuity model takes into account all of the factors that could have an influence on system availability.

The seven layer business continuity model is shown in Table 3.

Table 3: The Seven-Layer Business Continuity Model

LAYER NAME	DESCRIPTION
Server hardware	The server hardware forms the foundation of the business continuity model.
Server software	This is software that runs on the server hardware (e.g., the operating system, call control software, application software). It can have a major impact on system availability.
Data Network	The data network includes the data networking hardware and related protocols (e.g., Layer 2 switches, routers, networking protocols). Overall system availability is dependent on the data network availability.
Power Distribution	This fundamental requirement needs to be taken into consideration so that if required, equipment can continue to be powered even under fault conditions (e.g., uninterruptible power systems, generators).
Geography	System availability can be enhanced when the network design takes geographical distribution of equipment and personnel into account, particularly when the business is dispersed across multiple locations and / or cities.
Process	Company processes related to maintenance and repair need to be considered, since these processes can have a direct effect on availability.
People	The availability of maintenance and repair personnel and their impact on system availability needs to be considered; that is, whether or not repair and maintenance personnel are located on site or off-site.

Mitel white papers describing telephone system availability are available on Mitel OnLine under the topic **Business Continuity**. These white papers outline the steps that must be taken to determine the availability of a particular UC solution. Performing such an availability analysis is beyond the scope of this document, and it is recommended that Mitel Professional Services be contracted to assist with an availability analysis.

Chapter 3

TRAFFIC AND SCALING CONSIDERATIONS

Traffic and scaling considerations

This chapter discusses scaling topics and traffic topics in the following sections:

- “Scaling by architecture” on page 45
- “Traffic and scaling considerations” on page 55

Scaling by architecture

Multi-Instance architectures have significant advantages when operational and infrastructure costs are spread across multiple tenants. Mitel multi-tenant applications are designed to support multiple tenants within the maximum total user capacity of the platform. As such, the reference architectures are designed with maximum capacity platforms and is most cost- effectively deployed for smaller customers.

Virtual architectures have advantages when platforms are sized according to customer loads. The reference architectures are designed to address mid-to-large customers. Platform variations within the architecture allow scaling to various target markets, with the variants SMB and SMB-LD architectures designed for small to mid-size customers, and Scalable designed for very large enterprise customers.

Scaling considerations differ by architecture, as described in the following sections:

- “Small Business architecture scaling” on page 45
- “Small Medium Business scaling” on page 52
- “Medium Large Business scaling” on page 53

Small Business architecture scaling

In the SB architecture, most of the dependencies are related to scale, and therefore need to be considered by the service provider deployment, rather than at the end-customer.

Scaling the UC solution is primarily dependent on the number of users and instances that can be handled by the MiCollab server. The number of tenanted MiCollab Client Service instances, the number of associated MiVoice Business controllers, and the number of users are the primary factors for determining the tenanted MiCollab Client Multi-Tenant server scaling requirements for the SB architecture.

The MiCollab Client Service, MiVoice Border Gateway, and MiVoice Business are the main items that interact for scaling the solution.

The scaling for the SB architecture is based on the consolidation of end-user requirements against common resources within the service provider space. Consolidation reduces the number of required resources, as compared to allocating resources on a per customer basis. Some examples of service provider consolidation improvement include reduction of trunks and voice mail storage, compared to providing these on an individual end-user basis.

Additional scaling factors that apply to all installations, including UC and non-UC-based users, include:

- Scaling for SIP trunks and gateways
- Scaling for end-user phones and gateways
- Scaling of call control instances and associated media services and storage

Table 4: MiCloud Multi Tenant scaling by product

PRODUCT	TYPICAL SIZE RANGE	MAX CAPACITY ON ONE INSTANCE	MAXIMUM DENSITY
MiVoice Business Multi Instance	50-100 users	5000 users	Max 250 instances on one server
MiCollab Multi Tenant	50-100 users	• 12,500 users • 25,000 devices	Max 250 tenants
MiContact Center Business Multi-tenant	10-15 agents	• 2 MiVoice Business controllers paired for resiliency • 10 active concurrent agents • 225 calls in the busy hour • 500 agents maximum	• 50 tenants, with a max of 25 tenants on the Enterprise Server and 25 tenants across all Remote Server instances • 500 active concurrent agents • 11250 calls in the busy hour

Each application has limits on how many other applications it can interface to, as shown in the table below. The table should be read as “Row header” interfaces “N column headers”.

Table 5: Application connection limits

	MIVOICE BUSINESS	MIVOICE BUSINESS CLUSTER	MICOLLAB	MIVOICE BORDER GATEWAY	OPEN INTEGRATION GATEWAY	MICONTACT CENTER BUSINESS	EXTERNAL APP
MIVOICE BUSINESS	999	1	1	100’s	6	1	N/A
MIVOICE BUSINESS CLUSTER		0	1	100’s			
MICOLLAB		1	0	multiple	N/A	N/A	
MIVOICE BORDER GATEWAY	multiple	multiple	1	6		1	
OPEN INTEGRATION GATEWAY	250	1			0		1500
MICONTACT CENTER BUSINESS	25			multiple			
EXTERNAL APPLICATION					unlimited		

Scaling considerations for voice-only deployment (without EHDU)

The SB architecture is a voice-only deployment. The phones are limited to the Mitel MiVoice 53xx IP Phones, the 68xx MiNet Phones, and SIP Analog Terminal Adapter (ATA). MiCollab is not included in this configuration. This configuration covers an installation without UCC licenses, and end-devices are limited to one per user. Introduction of multiple devices is covered in the following sections. Additional scaling considerations apply to use with EHDU, and are also covered in the following sections.

The MiVoice 53xx IP Phones can work with an MiVoice Border Gateway cluster and can handle registration re-direction natively. This allows a common public registration IP address to be used, and the MiVoice Border Gateway to provide load sharing and registration redirection information. Additionally the public Mitel Redirection and Configuration Service (RCS) server can be used to simplify end-user phone installation and registration to the correct MiVoice Border Gateway gateway IP address.

The supported SIP Analog Terminal Adapter (ATA) devices do not natively support redirection or connection to the RCS server. The registration IP address must be provided by the service provider and provisioned locally at the device. Information about supported SIP ATA devices can be obtained from the Mitel SIP Center of Excellence (CoE).

The MiVoice Business scales based on the number of users, the devices, and the UC profile as described in “MiVoice Business scaling” on page 58.

Because of the flow-through model of the SB architecture, there are two groups of MiVoice Border Gateways; SIP trunk MiVoice Border Gateways and user MiVoice Border Gateways.

The user MiVoice Border Gateway scaling is determined by media streaming and device registration limits, as defined in the *MiVoice Border Gateway Engineering Guidelines*. In this case, the voice MiVoice Border Gateway software is more likely to be directly installed on server hardware, rather than virtualization, and therefore limited by server performance. Limits on the number of MiVoice Border Gateways in a cluster also limit the number of streams and devices that can be handled within that cluster.

A user MiVoice Border Gateway cluster is normally deployed in a N+1 configuration, where one of the MiVoice Border Gateways is used as a backup, in case another unit fails. The MiVoice Border Gateways should be designed such that required capacity is possible when a single unit is unavailable, i.e. N.

The MiVoice Border Gateways used for SIP Trunks are primarily limited by the number of number of active trunk connections. The same active user connection limits apply for the user MiVoice Border Gateways. In practice, the number of active trunk connections are lower than the active user connections, but for simplicity, it is easier to assume the active user and active trunk connections are the same.

Scale the trunks and user MiVoice Border Gateways according to server capacity and within the limits defined in the *MiVoice Border Gateway Engineering Guidelines*. See the section“Small Business MiVoice Border Gateway scaling” on page 49 for more considerations for scaling in this architecture.

Consider future expansion and whether UC will ever be deployed. If UC capabilities will never be needed, then the AMC configuration does not require use of the ULM construct. However, if the architecture will migrate to using UC at some time in the future, use the ULM construct in AMC, even if MiCollab is not currently being installed.

Scaling considerations with Voice Plus Entry UCC (including EHDU)

The Voice plus Entry UCC deployment includes a level of light Unified Communications including EHDU functions. The scaling requirements in this section cover the situation with Entry IPT and Standard IPT users that only have a single EHDU device, rather than a desk phone, as well as users with multiple devices using the Entry SB UCC licensing profile.

The added Unified Communications is provided by the introduction of MiCollab with the Entry SB UCC license and user configuration. Entry SB UCC functionality provides the capability to include a smart phone with the deployment, for use as an External Hot Desk User (EHDU) device. Users receive calls on their mobile phone as well as their desk phone, with the mobile phone acting as an extension. The Entry SB UCC functionality allows the option to replace the Hot-Desk phone with a PC Softphone. For example, the user could deploy Hot-Desk phone plus EHDU, or PC Softphone plus EHDU. Users that are licensed with the Entry SB UCC profile can be assigned multiple devices. Maximum device limits for the Entry SB UCC profile may exist.

The Voice plus Entry UC business model has the same dependencies as the voice-only model with two additional considerations. More SIP trunking and additional MiVoice Border Gateways for user UC connections are needed.

- For the SIP trunk MiVoice Border Gateways, there is an expected increase of 35% in required capacity. This is based on 50% of users having this added UC capability, with 50% of the user's call traffic handled by the EHDU device.
- For the user MiVoice Border Gateways, units are added for the UC connection. The UC connection scales according to the capacity of the MiCollab Client Multi-Tenant server. The scaling is primarily driven by the number of registered users and the number of instances, or end-customers, that the MiCollab Client Multi-Tenant server can handle. Multiple MiVoice Border Gateways may be needed per MiCollab Client Multi-Tenant.
- Additional devices increase the total number of registered devices across the User and UC groups of MiVoice Border Gateways. The scaling split between the groups is dependent on the ratio of soft phones to total phones, excluding EHDU.

See “Small Business MiVoice Border Gateway scaling” on page 49 for more MiVoice Border Gateway scaling details.

Small Business MiVoice Business scaling

For the SB architecture, the typical end-customer deployment size is expected to be around 50 users. This number of users and devices can be handled by a single MiVoice Business. Two MiVoice Business controllers are typically assigned to a customer, one as primary and one as secondary backup. The MiVoice Business can be scaled to larger sizes; however, depending on the number of UC devices in use, different scaling limits may apply.

Table 6 shows suggested scaling for the different UC profiles. In the table, it is assumed that to obtain 1.5 devices per user, 50% of the users have a single device, and the other 50% have two devices—comprising a selection of internal phone, EHDU phone, and UC soft phone.

The MiVoice Business scaling limits for one device per user assumes that Multi-Device User Group (MDUG) licensing is not configured for these users. However, for users with multiple devices, MDUG is necessary.

Table 6: SB architecture UC Profiles

UC PROFILE	DEVICES PER USER	MIVOICE BUSINESS USER LIMITS (5600 DEVICE LIMIT)	MIVOICE BUSINESS USER LIMITS (700 DEVICE LIMIT)
Voice Only (no MDUG)	1	5600	700
Voice (no MDUG) 50%, Entry SB UC 50%	1.5	2800	350
Entry SB UCC 100%	2	2450	310

Small Business MiVoice Border Gateway scaling

Deploying the flow-through model of the SB architecture requires three groups of MiVoice Border Gateways, with the following functions:

- Connecting SIP trunks
- User connections
- UC connections

The number of MiVoice Border Gateways needed for SIP trunks is primarily driven by the number of active voice connections, and the database redirection entries (example: Which DDI/DIDs are assigned to which MiVoice Business?). Use of common DDI/DIDs reduces this limitation impact.

Typically, the SIP trunk MiVoice Border Gateways are statically assigned to specific MiVoice Businesss controllers. The MiVoice Border Gateways are also deployed on a 1:1 basis. For example, if four MiVoice Border Gateways are needed to handle the media streams, then eight MiVoice Border Gateways are deployed. SIP MiVoice Border Gateways are clustered in pairs to share licenses and database information.

The number of device registrations and the number of voice connections determine the number of MiVoice Border Gateways required for user connections. Typically, the number of voice connections for the user and trunk MiVoice Border Gateways are similar; however, the introduction of EHDU connections may increase the number of SIP trunk MiVoice Border Gateways required. EDHU connections also affect the number of MiVoice Border Gateways needed for UC connectivity if a UC client is used.

Table 7: MiVoice Border Gateway scaling logic

DEVICE	USER MIVOICE BORDER GATEWAY	SIP/UC MIVOICE BORDER GATEWAY	SIP TRUNK MIVOICE BORDER GATEWAY
	REGISTERED DEVICES		TRUNK MEDIA CHANNELS
Desk phone	+1		Trunk rate X Users
PC Softphone		+1	
Mobile Softphone		+1	
EHDU			+EHDU%

The User and SIP/UC MiVoice Border Gateway scaling is primarily driven by the number of registered devices; if there are 800 desk phones, the User MiVoice Border Gateway must scale to 800 devices. If there are 200 PC Softphones, the SIP/UC MiVoice Border Gateway must scale to 200 devices.

The SIP Trunk MiVoice Border Gateway scaling is driven by the number of active media channels. The number of active media channels is driven by the traffic rate and the number of users in the solution. When EHDU devices are deployed, an additional percentage-scaling factor is added. This scaling factor is due to calls being re-routed back out the PSTN/SIP Trunk; for example, an incoming trunk call results in two trunk connections, one incoming and one outgoing to the EHDU device.

The trunk rate is determined by the number of users, the average call rate, and the average call hold time. Erlang B is then applied to determine the peak number of trunks needed to maintain the desired blocking rate. The formula is:

Trunk Rate = Erlang B (Average Call Per Hour x Average Call Hold Time (seconds) / 3600)

For larger deployments, as a quick rule of thumb, the following formula can be used for trunks:

Trunk Rate = (Average Call Per Hour x Average Call Hold Time (seconds) / 3600) x 1.1

If EHDU is deployed, this adds trunks to the SIP Trunk MiVoice Border Gateway. It is assumed in the calculation that, for the SB architecture. the multi-device user count is based on two devices, one of which is EHDU. The additional EHDU% overhead is therefore:

EHDU% = (Number of Users with single devices x 150% / Total number of users)

+ (Number of Users with multiple devices x 75% / Total number of users)

For example: Suppose the solution has 1000 users with the following deployment:

- 400 single device users with 100 using EHDU
- 600 multi-device users with 300 using EHDU

The number of trunks is determined by the number of users and the traffic. Assume a standard six Calls per Hour and an average hold time of 100 seconds. Using the rule of thumb calculation, we need approximately 183 trunks (the Erlang B calculation gives a result of 187 trunks).

Adding the EHDU consideration, we need to add 37.5% more trunks:

= (100 single device EHDU users x 150% + 300 multi-device users x 75%) / 1000 users

For this example, we need 252 trunks (183 x 1.375)

User MiVoice Border Gateways can be deployed as a N+1 cluster for MiVoice Business phones. For example, if four MiVoice Border Gateways are needed for the user connections, five MiVoice Border Gateways would be deployed in an N+1 configuration. See the *MiVoice Border Gateway Engineering Guidelines* for cluster scaling limits.

MiVoice Border Gateway scaling rules are different for SIP devices. MiVoice Border Gateways are deployed in clustered pairs for resilient SIP phones using DNS to locate a secondary access point. The access points for SIP and Mitel proprietary phones are separated due to their different scaling requirements.

The UC soft phones (PC Softphone and UC Mobility Softphone) are SIP devices, and are not resilient between MiVoice Border Gateways. Therefore, MiVoice Border Gateways deployed for UC connections must be virtualized so that they can use VMware High Availability (HA) for resilient operation.

MiVoice Border Gateway supports only a single connection to a single MiCollab Client Multi-Tenant server. MiVoice Border Gateways may not be associated with multiple MiCollab Client Multi-Tenant servers. The use of MiVoice Border Gateway Virtuals and VMware HA is recommended for both SIP phones and UC soft phones. Using smaller MiVoice Border Gateway Virtuals allows multiple MiVoice Border Gateways per server, and reduces server hardware requirements.

In general, the physical server MiVoice Border Gateways provide better performance and scaling. The common SIP Trunk and user voice MiVoice Border Gateways are usually deployed with physical servers. The UC and SIP users connections are typically deployed with MiVoice Border Gateway Virtual units. These deal with UC scaling and provide VMware HA capability for end-devices that are not resilient and cannot deal with multiple IP addresses. For lower scaling, the physical servers can be replaced with virtual servers.

See the *MiVoice Border Gateway Engineering Guidelines* for registration limit, voice connection, and cluster limits for both MiVoice Border Gateway Virtual and MiVoice Border Gateway.

Small Business MiCollab Client Multi-Tenant Service scaling

The UC deployment scaling dependencies include:

- the number of users on a particular MiCollab Client Multi-Tenant server
- the number of customers on a particular MiCollab Client Multi-Tenant server
- the number of MiVoice Border Gateways assigned to a MiCollab Client Multi-Tenant server, including the number of clients and SIP soft phones
- the requirement to logically separate user MiVoice Border Gateways for voice only devices from those used for UC, such as PC Softphones, and also for UC Clients.

The deployment requires Oria to provision the users on MiVoice Business, MiVoice Border Gateways, and also on MiCollab Client Multi-Tenant Service. Some initial provisioning is needed prior to this operation. For details, see the *MiCloud Business Solution Small Business Deployment Guide*.

Small Medium Business scaling

The Small Medium Business architecture is optimized for customers with 50-250 users, with the ability to scale to 500 end-users. End-users are located at one site.

The SMB deployment is a self-contained virtualized OVA intended for UC attachment, and for deployment in virtualized data centers. The scaling of the components in the OVA are already predefined to work correctly with the target UC profile of 2.75 devices up to 500 users.

Since the deployment does not need to consider additional external units or connections, from a performance and scaling view, this solution is complete. It does not need to consider additional scaling factors for MiVoice Business and MiVoice Border Gateway, as these are already included and scaled for this deployment.

Small Medium Business - Low Density scaling

The MiCollab Virtual scales to numbers larger than expected to be deployed with the SMB-LD architecture. The key factors for scaling this architecture are based on the deployment of

MiVoice Business and MiVoice Border Gateway per customer, even though the instances are consolidated into the same service provider infrastructure.

For a voice-only service without Teleworkers, the MiVoice Border Gateway Virtual provides simple SBC functionality for SIP trunks only. An alternative option to MiVoice Border Gateway Virtuals is to deploy a common VLAN-aware SBC and manage the MiVoice Business via the MiVoice Business Multi-Instance management access portals. The list of suitable SBCs is available from the Mitel SIP Center of Excellent (SIP CoE).

When Teleworker services are required, MiVoice Border Gateway Virtuals must be deployed for the phones to function correctly. A third-party SBC cannot provide the necessary Application Level Gateway (ALG) functions that the MiVoice Border Gateway provides. When UC is deployed, more direct management access to the MiCollab Virtual is required and a number of services are proxied via the MiVoice Border Gateway. In this case MiVoice Border Gateway Virtuals must be used.

Medium Large Business scaling

The MLB architecture is optimized for customers that have from 250 to 1500 end-users. The architecture can scale up to 5000 users.

MiVoice Business Virtual scaling

For the MLB UC deployment, the deployment size can range from 250 to 1500 end-users, and can scale up to support 5000 users.

MiVoice Border Gateway Virtual scaling

The MiVoice Border Gateway Virtual supports a number of different logical functions, such as SIP Trunking connectivity, Session Border Controller functionality and Teleworker gateway functionality for both SIP and MiNet IP end-points.

For the MLB architecture deployment:

- One MiVoice Border Gateway Virtual provides connectivity to Teleworkers via the SIP Trunk Service provider, but the MiVoice Border Gateway Virtual is deployed in resilient pairs to support resilient operation.
- One MiVoice Border Gateway Virtual (internal to MiCollab Virtual) provides connectivity to the Internet Service provider, but the MiVoice Border Gateway Virtual is deployed in resilient pairs to support resilient operation.

As a condition of the UC licensing model, all Teleworkers need be in the same cluster, but MiVoice Border Gateway load sharing can be used to effectively provision resilient SIP users with MiVoice Border Gateways in a 1+1 mode, and resilient MiNet users in a N+1 mode.

The number of MiVoice Border Gateways required for SIP trunks is primarily driven by the number of active voice connections, and the database redirection entries (example: Which DDI/DIDs are assigned to which MiVoice Business?). Use of common DDI/DID ranges therefore reduces this limitation impact.

Typically the SIP trunk MiVoice Border Gateways are statically assigned to specific MiVoice Business controllers. For resiliency, the MiVoice Border Gateways are deployed on a 1:1 basis. For example, if two MiVoice Border Gateways are needed to handle the media streams, then four MiVoice Border Gateways are deployed. SIP trunk MiVoice Border Gateways are clustered in pairs to share licenses and database information.

The number of device registrations and the number of voice connections influence how many MiVoice Border Gateways are required for user connections. Typically, the number of voice connections for the user and trunk MiVoice Border Gateways are similar. However, the introduction of EHDU connections may increase the number of SIP trunk MiVoice Border Gateways required.

User MiVoice Border Gateways can be deployed in an N+1 configuration for Mitel proprietary hard phones (MiNet phones). For example, if four MiVoice Border Gateways are needed for the user connections, five MiVoice Border Gateways would be deployed in an N+1 configuration. See the *MiVoice Border Gateway Engineering Guidelines* for cluster scaling limits.

MiVoice Border Gateway scaling rules are different for user SIP devices. MiVoice Border Gateways are deployed in resilient pairs (1 +1) for resilient SIP phones using DNS to locate a secondary access point. The access points for SIP and Mitel proprietary phones are separated because they have different scaling requirements.

MiVoice Border Gateways used for UC connections, with UC Clients and soft phones, are associated with a single MiCollab server. These MiVoice Border Gateways should be clustered with the MiVoice Border Gateway internal to MiCollab for license sharing and configuration.

See the *MiVoice Border Gateway Engineering Guidelines* for registration limit, voice connection, and cluster limits for the MiVoice Border Gateway.

Traffic and scaling considerations

Traffic on the system determines how many resources are used—bandwidth, trunks, and voice mail. UC deployments must also consider the types and number of devices that are assigned to each user UC profile.

This document uses some base assumptions for calculating the traffic and scaling limits for each architecture. Customers wishing to use different traffic rates or assumptions should contact Mitel Professional Services to confirm that their scaling and quantity of unit calculations are correct.

- "UC profiles" on page 55
- "User traffic levels" on page 57
- "MiVoice Business scaling" on page 58
- "MiVoice Border Gateway scaling" on page 60
- "MiCollab scaling" on page 62
- "Mitel Open Integration Gateway (OIG) scaling" on page 63
- "Oria networking and scaling" on page 64
- "MiContact Center Business scaling" on page 65
- "MiVoice Call Recording scaling" on page 66

UC profiles

The number of devices assigned to a user is an important factor for considering performance and configuration limits. Users, including the devices they use, are classified into five profiles shown in Table 8. Three of the profiles are considered UC-specific.

Table 8: User and UC profiles

USER AND UC PROFILE	EXAMPLE
Base IPT	A physically fixed phone with no voice mail. Typically used in lobbies, conference rooms, or work stations.
Standard IPT	A typical office desk phone assigned to a user or hot-desk. Extension includes voice mail for the user
Entry UCC	An extension on the standard IPT with inclusion of an EHDU or forwarding of calls to a specific PSTN number. This is the first multi-device user level, up to two devices assigned to the user **Note:** With the SB architecture the Entry SB UCC license will be used, but is still considered Entry UCC for the scaling calculations
Standard UCC	Extends the Entry UCC profile with the addition of PC soft phone and Teleworker capability. Up to three devices are assigned to the user
Premium UCC	Extends the Standard UCC profile with the addition of a mobile phone soft phone and an additional Teleworker. Up to four devices assigned to the user

For details about licensing, see the *MiCloud Business Solution Blueprint*.

Traffic and scaling calculations in this document employ an average number of devices per user to consider the different profile mixes for the overall deployment.

Table 6 describes the expected and typical UC profiles for standard deployments and corresponds to the average devices per user profiles in Table 8.

Note: Customers wishing to use different traffic rates or assumptions are advised to contact Mitel Professional Services to confirm that their calculations are correct for any changes to scaling.

Table 9: Average Devices per User Target Scenarios

KEY TARGETS	DESCRIPTION
2.75 Devices Target UC	This is the target UC deployment scenario using the standard templates provide with MiCollab.
1.5 Devices Special	This is a special condition for the SB architecture solution where extended mobility with PC soft phones may be used instead of EHDU or simple twinning. This is a non-standard template. It requires UC licensing, with an addition a-la-carte mobility license.
1.5 Devices Target UC	This is the target UC deployment for the SB architecture solution. This is a reduced UC deployment with only Entry UC.
1.0 Devices Standard	This is the standard non-UC deployment with a mix of fixed and hot desk phones.

A customer may wish to use different target profiles, or devices per user, from those shown above. These target profiles are expected to fit most UCC deployments.

It is possible to increase or decrease the number of devices per UCC profile. Table 10 provides some examples.

Table 10: Typical UCC Profiles for standard deployments

	PHONE		UCC		
DEVICES PER USER	BASIC IPT	STANDARD IPT	ENTRY	STANDARD	PREMIUM
4	0%	0%	0%	0%	Internal HD device PC soft phone (TW) Mobile soft phone (TW) EHDU 100%
3.75	0%	0%	0%	Internal HD device PC soft phone (TW) EHDU 25%	Internal HD device PC soft phone (TW) Mobile soft phone (TW) EHDU 75%

Table 10: Typical UCC Profiles for standard deployments

	PHONE		UCC		
DEVICES PER USER	BASIC IPT	STANDARD IPT	ENTRY	STANDARD	PREMIUM
3.5	0%	0%	0%	Internal HD device PC soft phone (TW) EHDU 50%	Internal HD device PC soft phone (TW) Mobile soft phone (TW) EHDU 50%
3.25	0%	0%	Internal HD device EHDU 25%	Internal HD device PC soft phone (TW) EHDU 25%	Internal HD device PC soft phone (TW) Mobile soft phone (TW) EHDU 50%
3	0%	0%	Internal HD device EHDU 25%	Internal HD device PC soft phone (TW) EHDU 50%	Internal HD device PC soft phone (TW) Mobile soft phone (TW) EHDU 25%
2.75 Target UC	0%	0%	Internal HD device EHDU 25%	Internal HD device PC soft phone (TW) EHDU 50%	Internal HD device Mobile soft phone (TW) EHDU 25%
2.5	0%	0%	Internal HD device EHDU 50%	Internal HD device PC soft phone (TW) EHDU 50%	0%
2.25	0%	Internal HD device 25%	Internal HD device EHDU 25%	Internal HD device PC soft phone (TW) EHDU 50%	0%

User traffic levels

User traffic levels are typically defined as “Standard Office” and are defined as:

- Six Calls Per Hour (6 CPH)
- Hold times 100-120 seconds per call

Other traffic levels can also be used, such as lower values for hospitality deployments or increased levels that might be used for a distribution type of business.

Note: Customers wishing to use different traffic rates, or assumptions, are advised to contact Mitel Professional Services to confirm their calculations for different scaling and quantity of unit calculations.

Traffic blocking is based on Erlang B calculation and uses the following standard blocking levels:

- P.01 for external trunks
- P.001 for internal traffic

Erlang B is used to estimate peak traffic conditions and resources required to meet the peak demand using the blocking ratios above. The peak period is considered for key hours during the day, typically in the morning and early afternoon. Traffic and resource occupancy can be expressed in Erlangs or Centum Call Seconds:

- 1 e (Erlang) = 100% occupancy = 3600 CS (Call seconds) = 36 CCS (Centum Call Seconds)
- 36 CCS = 36 calls of 100 seconds duration

Other traffic factors to consider include where the traffic originates and terminates. It is assumed, on average, that calls are answered across all devices in an even fashion for UC users. Changes to this could influence use of certain resources, e.g. trunks.

MiVoice Business scaling

MiVoice Business scaling is especially important for larger deployments where self contained units, such as those used for the SMB, cannot be deployed. See "MiVoice Border Gateway scaling" on page 60 for the self-contained UC deployment solution.

There are a number of limits, internal to MiVoice Business, that need to be considered for the UC deployments on the MiVoice Business, including, but not limited to:

- Performance
- Registered devices
- Monitors
- HCI and application attachment

Each of these factors is considered, along with the number and type of devices per user. From this a scaling value can be obtained, which determines the number of MiVoice Business controllers that are required for the solution. The platform type may provide additional influence, especially when considering virtual MiVoice Business controllers that have potentially reduced capabilities, such as user and traffic limits, compared to other MiVoice Business controllers.

For non-UC deployments, each user has one device for receiving all calls. Users can have multiple devices in a UC deployment. All of a user's devices receive incoming calls, except when presence settings are configured for an alternative. One device is used to answer the call, but multiple ringing connections are established and torn down. The scaling calculations consider this process for determining the available system performance.

MiVoice Business building blocks

For larger deployments that use virtual MiVoice Business, it has been determined that the MiVoice Business Virtual 2500 has sufficient capacity to handle the defined number of UC users and the required media resources.

For the standard MiVoice Business deployments, the MiVoice Business for Industry Standard Servers (ISS) and MiVoice Business Multi-Instance platforms have sufficient capacity to handle the UC deployments. When using MiVoice Business Multi-Instance, it is strongly advised that the individual MiVoice Business controllers are performance-checked with the System Engineering Tool (SET) and with the MiVoice Business Multi-Instance Engineering Tool (MET). Both of these tools are available on Mitel OnLine.

MiVoice Business building blocks:

- MiVoice Business Virtual 250, 1500, or 2500 for virtual deployments
- MiVoice Business for Industry Standard Servers (ISS) for on-premise deployment
- MiVoice Business Multi-Instance for service provider and on-premise deployment

Note: Although the MiVoice Business Virtual (5000 user) can offer more capacity in certain areas, this is not applicable for the defined UC deployments. The MiVoice Business Virtual (5000) can handle more users, but the total device limit is the same as for the MiVoice Business Virtual (2500 user), with an increased OVA reservation requirement.

For a UC deployment, there is no benefit, and a potential cost disadvantage, in using the MiVoice Business Virtual (5000). This is why the MiVoice Business Virtual (2500) is used as the virtual UCC building block.

MiVoice Business scaling details

For the three MiVoice Business building blocks, the same UC limits apply. Figure 8 shows the number of MiVoice Business controllers required when the number of users, including UC, is compared to different numbers of devices per user, as defined in “UC profiles” on page 55.

Figure 8 shows the number of MiVoice Business units needed for an operational UCC configuration. Where primary and secondary resilient MiVoice Business units are required, double the number of MiVoice Business units. Where VMware HA is being used instead of application resiliency, the numbers on this chart can be used directly.

The special mobility case for the SB architecture does not affect the predicted number of required MiVoice Business controllers, since the number of users is typically less than 50, with fewer than 100 devices per MiVoice Business.

Note: Contact Mitel Professional Services for help with calculating MiVoice Business scaling if there are more than 50 users and 100 devices per MiVoice Business controller planned for a SB architecture deployment.

Figure 8: MLB - MiVoice Business scaling

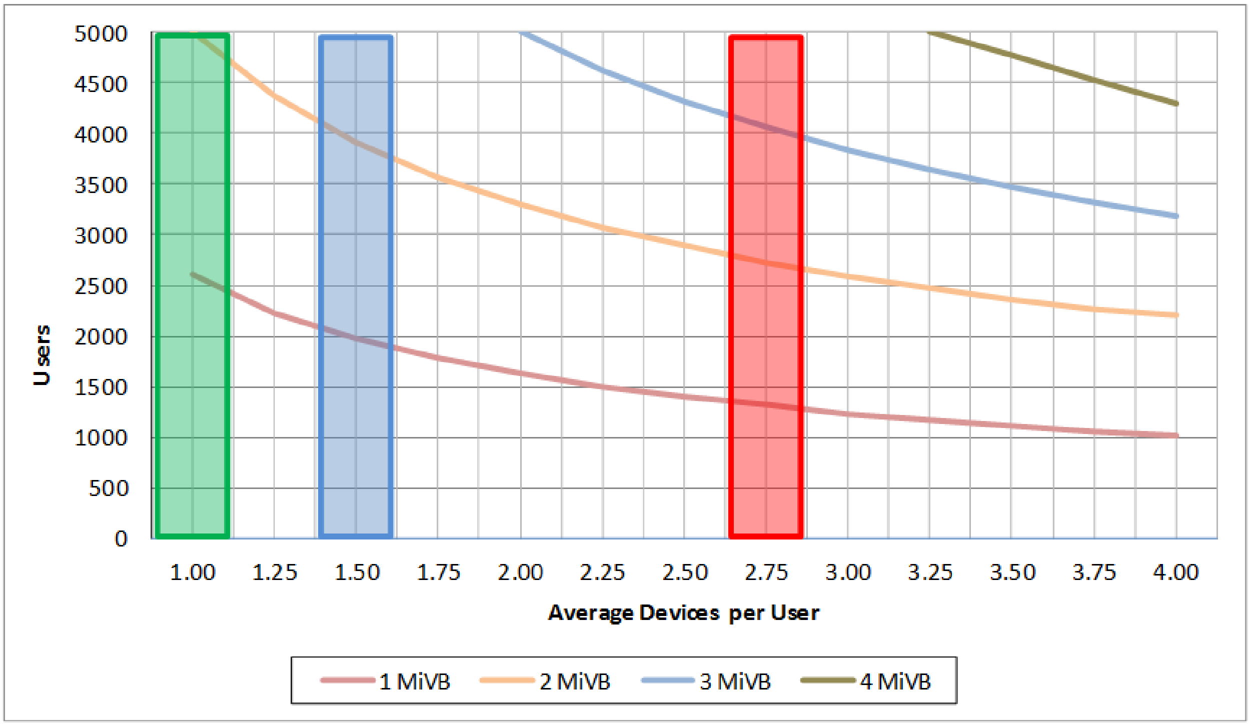

MiVoice Border Gateway scaling

MiVoice Border Gateway scaling in a virtual environment is important for UC deployments requiring a solution composed of individual components. MiVoice Border Gateway Virtual scaling is especially important for larger deployments in which self-contained units, such as those used for the SMB, cannot be deployed.

There are a number of internal limits that need to be considered for UC deployments on MiVoice Border Gateway Virtual, including, but not limited to:

- Performance
- Number of concurrent media and UC connections through the MiVoice Border Gateway
- Number of registered users

Typically MiVoice Border Gateway Virtual is used for two functions:

- SIP Trunk Gateway
- Teleworker Gateway

By combining these functions to a group of MiVoice Border Gateways, it is possible to consolidate the number of instances that need to be created. Figure 9 considers this consolidated case. If it is necessary to split MiVoice Border Gateway Virtuals into separate functions, then you can add two more MiVoice Border Gateway units for deployments up to 2500 users, and a further two units to support up to 5000 users.

Included in the requirement is the possibility that the PC soft phone is deployed on a laptop or tablet, and may therefore be mobile outside of the normal office environment. In this case, the device must connect into the deployment as an external Teleworker, rather than as an internal LAN connected device.

The scaling for the MiVoice Border Gateway Virtual is based on the registration and streaming limits in the *MiVoice Border Gateway Virtual Engineering Guidelines*.

Note: The MiVoice Border Gateway can be associated only with a single MiCollab. The MiVoice Border Gateways should be scaled by the limits of this single associated MiCollab.
Where multiple MiCollab units exist, multiple groupings of MiVoice Border Gateways are required.

Figure 9 shows the number of MiVoice Border Gateway Virtual units needed for a normal deployment. If VMware HA is used for resiliency, or if resiliency is not required, you use numbers shown in the graph. If the MiVoice Border Gateways are clustered for resiliency, increase the number of units by the resiliency configuration; that is, double for a primary/secondary combination and plus one for an N+1 configuration.

The special mobility case for the SB architecture does not affect the required number of MiVoice Border Gateway Virtuals, since the number of users is typically less than 50, with fewer than 100 devices per MiVoice Business.

Scaling of the MiVoice Border Gateway is important for UC deployments that require the solution to be composed of individual components, typically on-premise. This is especially important for larger deployments where self-contained units, such as used for SMB, cannot be deployed.

Note: Contact Mitel Professional Services for help in calculating the MiVoice Border Gateway Virtual scaling if there are more than 125 users and 250 devices per MiVoice Business planned for a SB architecture deployment.

Figure 9: MLB - MiVoice Border Gateway scaling

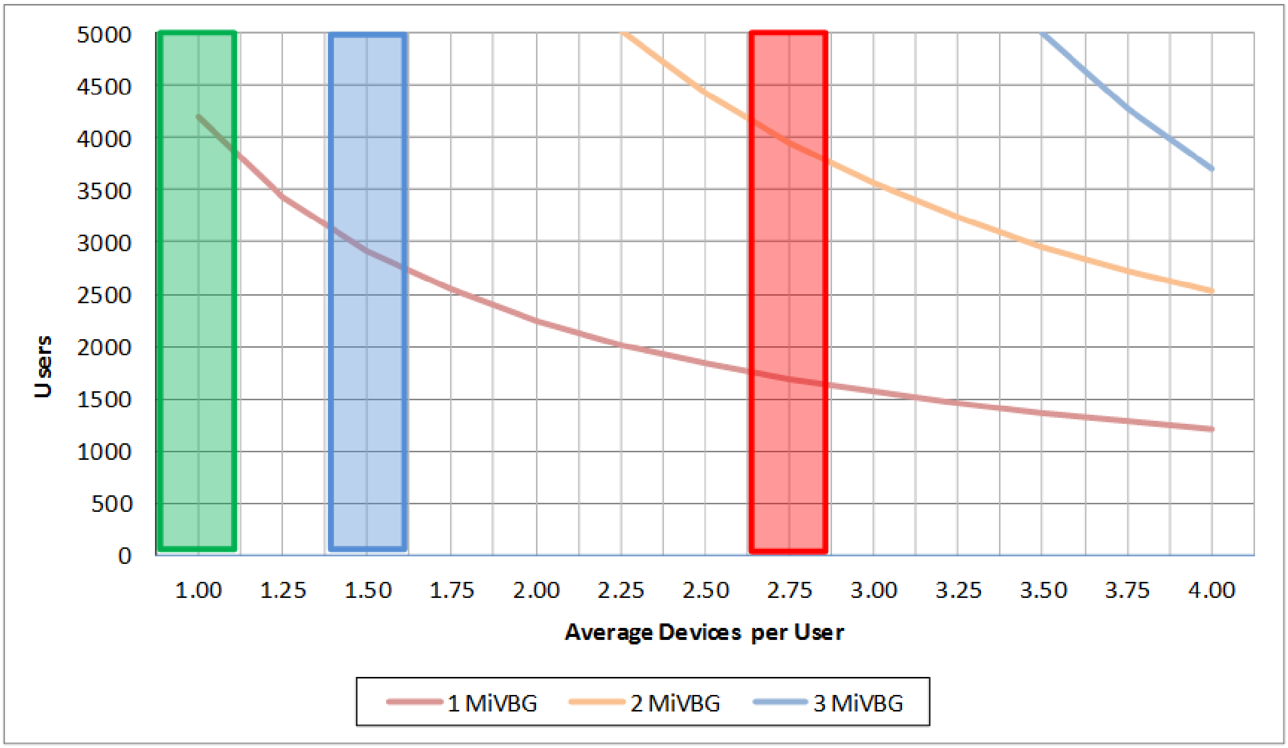

MiCollab scaling

The MiCollab platform is a collection of UCC applications. This is a multi-application deployment for a single customer and may be realized with a dedicated server for on-premise deployment, or as a virtual application for both private and public Cloud deployments.

The MiCollab application can scale up to 5000 users based on 2.75 devices per user, as a single deployment.

Scaling notes:

- MiCollab may be associated with multiple MiVoice Business units
- Multiple MiVoice Border Gateways can be associated with a single MiCollab unit.
- MiVoice Border Gateways cannot be associated with multiple MiCollab units.

MiCollab Client Multi-Tenant scaling

The MiCollab Client Multi-Tenant Service is a specific configuration using the MiCollab UCC platform. This is a single application and can be used across multiple customers to provide UC functionality (telephony presence).

This is not the multi-function MiCollab UCC solution, but a special deployment of one application.The general MiCollab scaling provides the richer UCC solution.

The MiCollab Client Multi-Tenant can scale to the following limits:

- Maximum 250 tenants on a single platform. For example:
  - 250 tenants at 50 users each - 12,500 users total
  - 5 tenants at 5000 users each - 12, 5000 users/25,000 devices (2 devices per user)
- One client per user (12500 clients total)
- Maximum of 5000 users per tenant, regardless of the number of tenants

The MiCollab Client Multi-Tenant solution is associated with multiple MiVoice Business units. Multiple MiVoice Border Gateway units are associated with this single MiCollab Client Multi-Tenant unit. Where multiple MiCollab Client Multi-Tenant units exist in a server provider environment, a similar number of MiVoice Border Gateway groups are also needed.

Mitel Open Integration Gateway (OIG) scaling

The Mitel Open Integration Gateway allows applications to access the UCC solution components and provides connections to Google and Salesforce integrations.

For details of the limits for Mitel Open Integration Gateway, see the *Mitel Open Integration Gateway Engineering Guidelines*.

Some key limits to consider with a UCC deployment include:

- Total number of monitors: 50,000
- Maximum MiVoice Business, or UC application connections (outgoing): 250
- Maximum number of user connected applications (incoming): 1500
- One OIG is deployed to a single customer

These limits mean adding the following considerations into the design:

- For the MiCollab Client Multi-Tenant solution, Mitel Open Integration Gateway must be deployed per customer, or two per customer if resiliency is required without using VMware HA. If OIG is used via a MiVoice Border Gateway, the MiCollab Client Multi-Tenant connections should use the same MiVoice Border Gateway to simplify tenant connections and IP address usage.
- Each MiVoice Integration for Google or Salesforce counts as an individual application attachment. A single OIG is limited to 1500 application attachments, so 1500 MiVoice Integrations can be connected to a single OIG. The number of OIGs must be calculated accordingly, that is, 3000 users on MiVoice Integration for Salesforce would require two Mitel Open Integration Gateway platforms.

Note: If multiple Mitel Open Integration Gateways connect to the same MiVoice Business and they monitor the same device, this adds a multiplier to the HCI count in the MiVoice Business.

OIGs can be deployed as virtual instances alongside existing virtual deployments, or they can be combined on to a single physical server for improved customer to server density. See the

Virtual Appliance Deployment Solutions Guide for further details on OVA and reservation settings.

Oria networking and scaling

Oria networking requirements are divided between access to the Oria platform through the web portal and web service interfaces, and access between Oria and the managed systems.

The web portal and web service interfaces are standards-based interfaces. The web portal server should be assigned a static IP address to avoid issues with DNS updates when dynamic addresses change. Remote access is achieved using the MiVoice Border Gateway web proxy service. In this case, Oria re-uses the available proxy selection for MiVoice Business connections.

Oria access to the managed systems depends on the deployment architecture. When Oria is located in the same local area network as the managed systems, straight-forward local networking can be applied. In the more typical case for service provider deployments, Oria is deployed in the service provider network while managing systems in the customer network.Oria communicates with managed systems using a variety of protocols for which application layer proxies are not readily available—special networking is required for Oria to communicate with managed systems.

For all single-tenant topologies, including the SMB-LD architecture, (all but MiVoice Business Multi-Instance), 1:1 NAT connections must be established between Oria and each managed system. and the NAT link must support bidirectional path initiation; that is, for Oria to initiate a connection to a managed system, and for the managed system to initiate a connection to Oria. This requirement for 1:1 NAT may be implemented with VMware networking or third-party routers. Oria addresses the managed UCC components using fully qualified domain names, and split DNS is used to resolve these domain names based on the local address space. For further information about the recommended network deployment, see “Network and networking considerations” on page 89 and the *MiCloud Business Solution Medium Large Business Deployment Guide*.

MiVoice Business Multi-Instance deployments provide management access independent of the control and signaling access. In this case, the management port may be provisioned within the service provider address space, while the control and signaling port is provisioned in the customer address space. This means that 1:1 NAT is not required for MiVoice Business Multi-Instance, as both Oria and MiVoice Business Multi-Instance management are in the same service provider network.

In cases where Oria is remote from the managed system, a VPN tunnel is required for secure WAN traffic. This is required for remote access between Oria and any managed system, including MiVoice Business Multi-Instance.

Oria Platform Manager and Oria File Server

The Oria Platform Manager and the Oria File Server can be installed on separate servers or together on the same server. The following table describes the server specifications for these options.

Note: Do not install Oria Platform Manager and/or Oria File Server on the server hosting Oria.

Table 11: Oria Platform Manager/Oria File Server: Server specifications

	HARDWARE SERVER	VMWARE
Oria Platform Manager	4 CPU (2 CPU with 2 cores each) 8 GB RAM	4 vCPU (2 CPU with 2 cores each) 8 GB RAM 50 GB disk space
Oria File Server	4 CPU (2 CPU with 2 cores each) 8 GB RAM	4 vCPU (2 CPU with 2 cores each) 8 GB RAM 150 GB disk space
Both Platform Manager and File Server on the same server		6 vCPU 12 GB RAM 200 GB disk space

Notes:

1. For use on VMware, use eager-zeroed thick disks for best performance.
2. For Platform Manager, fast disk configurations are recommended.

MiContact Center Business scaling

All MiContact Center Business functionality except embedded workforce scheduling, is available in the MiCloud offering, but scaling is highly dependent on the architecture deployed. For information about sizing MiContact Center Business servers and deploying resilient IVR ports, see the *MiContact Center Business Engineering Guidelines*.

Required server deployment	8 vCPU 16 GB Reservation: • Multi-Tenant:12.5 GHz • Enterprise: 9.5 GHz	Agent support: • Multi-Tenant: up to 500 agents • Enterprise: up to 1200 agents

Note: See the *Virtual Appliance Deployment Solutions Guide* for the most up-to-date information.

MiContact Center Multi-tenant

MiContact Center Multi-tenant customer configuration data is updated on every MiContact Center client (Ignite and Contact Center Client). With multi-tenant deployment, these updates are more frequent and the update size is larger. Also, it is difficult to control the time of day of updates. For small sites, the bandwidth required for updates may affect the QoS for active agent streams.

- To mitigate the bandwidth consumption, CC clients are supported for supervisors only.
- For sites with minimal bandwidth available, it is recommended to set Max IP packet size to limit QoS effects, and to use a premise-based router with TCP back-pressure to limit bandwidth.

MiVoice Call Recording scaling

MiVoice Call Recording can be installed on virtual servers with capacities ranging from 25 concurrent ports up to 750 concurrent ports, depending on the resources allocated in the virtual machine.

Minimum server deployment	2 vCPU 2 GB 2 GHz reservation	Supports 25 concurrent ports, 100 agents or users, and 50 desktop clients
Maximum server deployment	12 vCPU 8 GB 12 GHz reservation	Supports 750 concurrent ports, 3000 agents or users, and 1500 desktop clients

Note: See the *Virtual Appliance Deployment Solutions Guide* for the most up-to-date information.

Storage capacity and ports supported also vary depending on the SQL server:

- SQL Express supports 200 concurrent ports and 1.5M calls.
- Full SQL on a shared server supports 550 concurrent ports and 3M calls.
- Full SQL on a dedicated server supports 750 concurrent ports and 15M calls.

For larger sites, MiVoice Call Recording nodes may be clustered. For capacity and scaling information, see the *MiVoice Call Recording Engineering Guidelines*.

Chapter 4

EXTERNAL CONNECTIVITY

External Connectivity

This section discusses the various components, protocols and communication paths that are required to support connectivity between the UC solution and external networks, and also between UC components located in a data center or service provider cloud and components located at the customer site or Teleworker site.

While a number of external communication paths are discussed in this section, a specific architecture may not require all of the paths discussed here. At a minimum, all topologies will require external connectivity to support IP telephony and management. More generally, external connectivity is required for voice communications and related applications, such as:

- Telephony signaling and related audio streams
- UC applications and related video and data streams
- Business applications
- Network management
- Troubleshooting
- Product licensing and license renewal
- Software maintenance

In a typical deployment, the call control engine and applications servers are deployed within a private network with a limited number of secure gateways to other networks. The end-points—both UC clients and management clients—are deployed at the customer site and other remote locations. Gateways are required to connect UC servers to clients and UC servers to other service providers. These external-facing gateways must support signaling, audio, video, and data traffic connecting over public or private networks.

In general terms, a gateway is a component that interconnects two different types of networks, both physically and logically; the gateway may also serve as a protocol converter between the two networks. The Unified Communications topologies rely on several gateways for connections to external networks, which include:

- Mitel 3300 ICP for PSTN connectivity
- MiVoice Border Gateway for public data network connectivity, including SIP trunks
- Third-party IP router and firewall for public IP-based data network connectivity
- Third-party label switch router for MPLS-based data network connectivity

Table 12: Additional resources

FOR INFORMATION ABOUT...	SEE...
Mitel 3300 ICP working capacity and using a MiVoice Business 3300 ICP for local breakout	*MiVoice Business Engineering Guidelines*
MiVoice Border Gateway, including clustering, resiliency, and bandwidth requirements	*MiVoice Border Gateway Engineering Guidelines*

Table 12: Additional resources

FOR INFORMATION ABOUT...	SEE...
Details for configuring the MiVoice Border Gateway to support Teleworker services and/or Web proxy	*MiVoice Border Gateway Installation and Maintenance Guide*

Service Provider Gateways

Depending on the external connectivity requirements of a particular architecture, the following service provider gateways may be required.

To connect to the PSTN:

- A PSTN gateway and the associated trunk connections to a Telco; and
- A SIP Aware Proxy and the associated IP-based trunks to a SIP service provider. Typically, the SIP service provider offers connectivity through to the PSTN and to other SIP service providers.

To connect to cloud-based data services:

- A web proxy for Mitel Open Integration Gateway-based services and other web services plus associated public data network connections; and
- A router with firewall for client and management traffic for UC applications, and for other services and associated private or public data network connections.

These gateways are co-located with the UC servers, typically within a data center.

PSTN Connectivity (central and remote)

MiVoice Business running on a 3300 ICP platform can serve as a PSTN gateway in On-Premise deployments, but are not used in hosted deployments,

In some On-Premise deployments, a 3300 ICP may be used to provide remote survivability and local PSTN break out.

SIP Connectivity

SIP trunks are provided by Internet telephony service providers to connect communication platforms to other SIP switches and to the traditional PSTN. SIP trunks offer several benefits compared to traditional PSTN connections, including:

- local phone numbers from any location
- increased resiliency for disaster recovery
- typically less expensive toll-free service and overall cost savings

SIP trunks are established from the MiVoice Business to the SIP trunk provider using MiVoice Border Gateway as a SIP-aware firewall and proxy. The MiVoice Border Gateway SIP proxy implements a full back-to-back user agent. The MiVoice Border Gateway SIP trunk service provides:

- NAT traversal of media and signaling
- Media anchoring for the remote provider
- SIP adaptation and normalization to improve interoperability
- DTMF detection as per RFC 4733, re-ordering of RTP streams, and KPML notifications to support EHDUs
- Protection from malformed and malicious requests, request flooding and various other types of attacks

A “SIP trunk” in the context of a MiVoice Border Gateway is a pair of end-points defined by their IP addresses and signaling ports. One end-point is typically the call control engine (MiVoice Business) and the other end-point is the SIP trunk provider’s firewall or SBC. A trunk can have any number of “channels” each of which corresponds to an active bi-directional media stream.

Outgoing SIP call routing is configured with the MiVoice Business using one or more Automatic Route Selection (ARS) rules. The SIP trunk service provider is configured as a Peer Profile, specifying the MiVoice Border Gateway as the related SIP proxy. Routing rules support specifying the minimum and maximum channels, i.e. active calls per trunk. Multiple routing rules to different SIP Trunk providers are supported. MiVoice Business detects trunk failures and set all trunks to the same peer out of service for both incoming and outgoing calls. In the case of outages causing missing keep alive messages, the MiVoice Business sets the trunks out-of-service to prevent instability in trunk status.

There are several methods to support resilient operations for outgoing calls, including

- SIP primary-secondary: The SIP Trunk provider provides a primary and secondary address. These addresses may be configured as alternate peers in the MiVoice Business (MiVoice Business).
- MiVoice Border Gateway with FQDN: The MiVoice Border Gateway configured as a related SIP proxy may be entered using an FQDN. The DNS server may be configured to resolve the FQDN to multiple MiVoice Border Gateways.
- MiVoice Business as alternate peer: The secondary MiVoice Business may be configured as an alternate peer. In this case, the primary path uses the related primary MiVoice Border Gateway SIP proxy, with the alternate path using the secondary MiVoice Business and its related secondary MiVoice Border Gateway SIP proxy.

These methods of providing resilient trunk routing for outgoing calls are shown in Figure 10.

Figure 10: SIP Trunk routing - outgoing

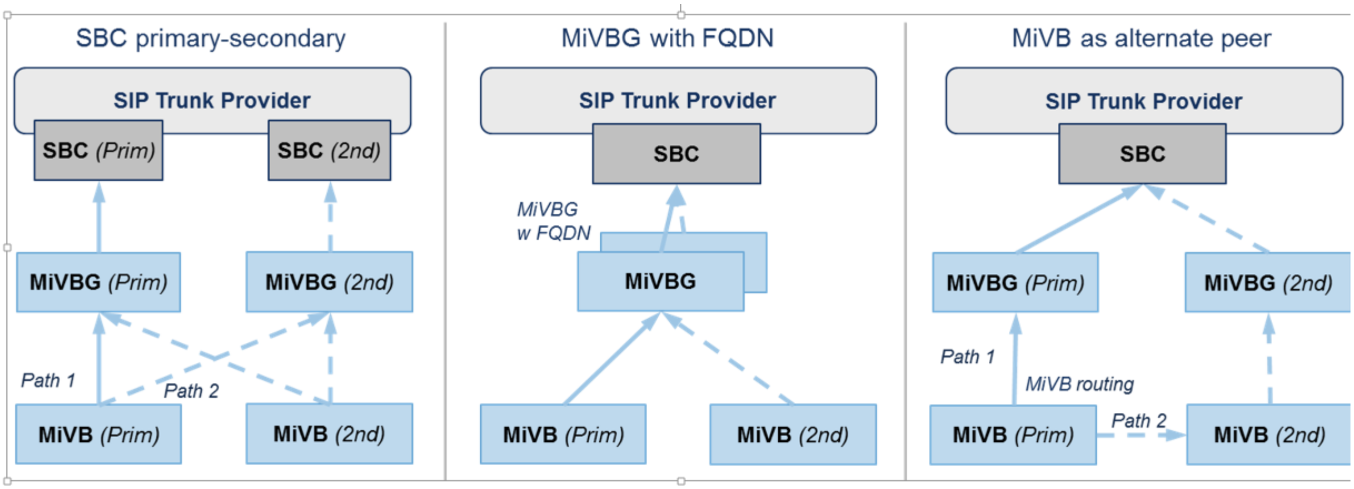

Incoming SIP call routing is determined by the SIP Trunk Service provider routing to MiVoice Border Gateways and by MiVoice Border Gateways routing to MiVoice Business. Typically, third-party session border controllers (SBCs) support primary and secondary routes. The SBC automatically detects trunk failures and re-routes over the secondary path. The MiVoice Border Gateway SIP trunk proxy supports routing rules with match criteria mapped to all or part of the SIP URI within the SIP header Request, From, or To fields, effectively providing routing based on called party, calling party, and original called party. Routing rules can designate both the primary and secondary controller for the route.

A MiVoice Border Gateway automatically detects trunk failures and routes to the secondary controller. Figure 11 shows resilient SIP trunk routing for incoming calls.

Figure 11: SIP Trunk Routing - Incoming Calls

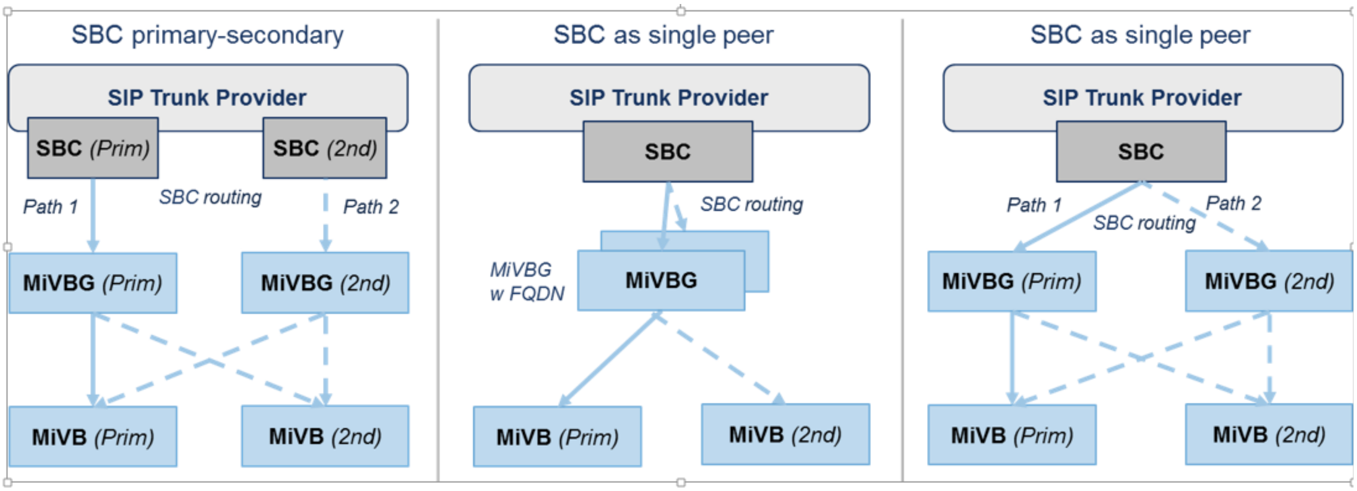

For smaller sites with shared trunking, one MiVoice Border Gateway can act as proxy for multiple MiVoice Business controllers, up to the active trunk capacity limit. For larger sites, multiple MiVoice Border Gateways and MiVoice Business controllers can be deployed for increased capacity. MiVoice Border Gateways support clustering up to six units with license sharing among cluster members. When dedicated to SIP trunking, MiVoice Border Gateway clustering beyond the resilient pair provides no advantage for resiliency or capacity as trunk traffic is determined by routing rules while clustering incurs the cost of increased synchronization communication. As such, MiVoice Border Gateways dedicated to SIP trunking should be deployed in a two-member cluster containing the resilient pair and allowing license sharing between the pair.

For smaller deployments, a resilient pair of MiVoice Border Gateways may have sufficient capacity to act as both the SIP trunk gateway and access gateway for Teleworkers and web proxy, provided that both are reachable with a public address. In the case that SIP trunks are carried over a private network, separate resilient pairs of MiVoice Border Gateways are required, with the SIP trunk gateway configured with a carrier-specific non-public IP address and the access gateway configured with a public IP address.

Table 13: MiVoice Border configurations

CONFIGURATION	SUPPORTS SIP TRUNK SERVICES
MiVoice Border Gateway[1] • DMZ mode • Server-Gateway mode MiVoice Business Express • Deployed with MiVoice Border Gateway in server-gateway mode	√
MiVoice Border Gateway[2] • DMZ mode: deployed on internal LAN configured for LAN mode operations (e.g. Secure Recording Connector) • Server-Gateway mode: deployed on internal LAN configured for LAN mode operations (e.g. Secure Recording Connector) MiCollab • Deployed on an internal LAN with MiVoice Border Gateway in LAN mode	X

Notes:

1. See “Deployment considerations” on page 80.
2. Similar considerations apply when the MiVoice Border Gateway is integrated with other applications.

When deployed in server gateway mode, SIP Service providers connect through the external facing router to the MiVoice Border Gateway’s WAN facing interface.

Licensing

A MiVoice Border Gateway SIP Trunking Channel license is required for each active channel. MiVoice Border Gateway SIP trunk licenses are for concurrent us, so for the MiVoice Border Gateway, the maximum number of active SIP trunk calls equals the number of configured MiVoice Border Gateway SIP trunk licenses, up to the MiVoice Border Gateway capacity limit. With UCC licensing, MiVoice Business SIP-enabled trunks are open licensed, providing sufficient capacity for the licensed users. With à la carte licensing, MiVoice Business SIP trunk licenses are for concurrent use. For MiVoice Business, the maximum number of active SIP trunk calls equals the number of configured MiVoice Business SIP trunk licenses, up to the MiVoice Business capacity limit. Both MiVoice Border Gateway and MiVoice Business support clustering with license sharing among cluster members. This allows the available licenses to be shared across individual MiVoice Border Gateways or MiVoice Business respectively. The total number of licenses required depend on busy hour traffic patterns.

Data Services Connectivity

Several cloud-based services may augment the UC solution. For example,

- The Mitel Application Management Center (AMC) provides software downloads and licensing for Mitel applications. Mitel applications must maintain data connectivity to the AMC to avoid entering license violation mode.
- MarWatch supports remote performance monitoring and management for Unified Communications systems and the associated networking infrastructure. MarWatch is available both as a cloud service and as an application deployed in the customer network.
- Integration with Salesforce.com provides full phone set management features in the hosted Salesforce.com user interface without the need for a locally installed client.
- Integration of MiVoice Business with remotely hosted business applications incorporates voice communications, conferencing, and collaboration capabilities into popular business applications such as personal information managers (PIMs), Microsoft Internet Explorer®, and Microsoft Office.

Integration of telephony features with applications is accomplished by inter-working with Mitel Open Integration Gateway connected to a MiVoice Business system. The Mitel Open Integration Gateway offers a web service, WSDL messages over SOAP, which may be accessed locally or remotely with typical web traffic networking configurations. A brief description of the Mitel Open Integration Gateway capabilities and requirements is provided in the earlier chapter on Applications. Mitel continues to support the legacy MiTAI interface on the MiVoice Business system, providing a Mitel enhanced TAPI interface. However, it is recommended that all new implementations be based on Mitel Open Integration Gateway.

These data services require connectivity from the UC systems to cloud-based servers. Such connectivity is implemented using a router with integral or in-line firewall. The MiVoice Border Gateway in server gateway mode provides static routing and basic firewall capabilities. Alternatively, a third-party router with firewall protection may be used. Firewall protocol/port configuration rules for many Mitel applications and services are described in *MiVoice Border Gateway Engineering Guidelines* and *Mitel Open Integration Gateway Engineering Guidelines*.

For web-based services, the MiVoice Border Gateway implements a web proxy to secure access between UC servers and other cloud service providers. This is recommended for connections from Mitel Open Integration Gateway to hosted business applications. MiVoice Border Gateway used as web proxy is discussed further below.

Access Gateways

Depending on the external connectivity requirements of a particular architecture, the following access gateways may be required.

To connect to remote workers, home-based, mobile, or anywhere off-site:

- A PSTN gateway or SIP-aware proxy providing PSTN connectivity for access by digital and cellular phones as EHDU devices
- A MiVoice Border Gateway providing Teleworker service for access by remote MiNet and SIP end-points, including hard and soft phones. MiVoice Border Gateway also provides a connector for use with MiContact Center desktop clients.
- A web proxy for access by MiCollab and MiCollab Audio, Web, and Video Conferencing clients and management clients
- An IP router with firewall for data traffic, such as associated media streams. A third-party router is used for connection to MiVoice Call Recording clients.

To connect to UC clients located at a customer site:

- An MPLS Edge Router or other LAN extension technology to extend the private LAN between the hosted UC server site and the customer site for LAN-based access by IP phones and other UC application and management clients
- An IP router for Internet-based access by UC application and management clients. The customer edge router uses Network Address Translation (NAT) capabilities to uniquely identify individual end-points.

For remote workers, access gateways are co-located with the UC servers, typically within a data center. For customer sites, access gateways are required, both at the data center edge and the customer LAN edge.

PSTN connectivity (EHDU)

Remote users may connect to UC services using digital or analog phones, with tone dialing, over the PSTN access network or cell phones over the public cellular wireless voice networks, all of which will connect through to the UC servers over PSTN or SIP trunks. The ubiquity of these devices and access networks provides a reliable means of connection from nearly any location.

For connection through the PSTN or cellular networks, the user dials a designated access number and then logs in for service as an External Hot Desk User (EHDU). The MiVoice Business may be configured with trusted trunks for EHDUs with Call Recognition Service enabled, which recognizes the Calling Line ID and automatically log in the user. In other cases, the user will need to log in with their hot desk directory number and Personal Identification Number (PIN). Once logged in, the external user is seen by the system as a local user and has access to extension dialing, voice mail, and other phone system resources. Call handling and call features are available through simple keypad commands.

The PSTN or SIP trunk connections to the MiVoice Business appear similarly to other incoming trunks except that they need in-band DTMF detection for call handling. For PSTN connections, a MiVoice Business running on a 3300 ICP platform can serve as a PSTN gateway, as described above. For SIP trunks, an MiVoice Border Gateway can serve as SIP trunk proxy, also described above. For SIP trunks, the SIP Trunk provider should transmit in-band DTMF tones with RTP packets compliant to RFC 4733, or the earlier RFC 2833. The MiVoice Business uses KPML subscription to the SIP trunk proxy to be notified of key press events used to manage call handling and call features. Termination of SIP trunks for EHDUs is not supported on the MiVoice Business as it does not support detection of in-band DTMF tones. Termination of SIP trunks on third-party SBCs requires interoperability testing to ensure compliance. Use of the MiVoice Border Gateway for SIP trunk termination is recommended.

When determining the required capacity of the PSTN or SIP trunk gateways, the EHDU traffic must be considered in the calculations. A trunk connection is required for every active EHDU connection for the communication path between the MiVoice Business and the external EHDU device. This is in addition to any trunk requirements for routing the call between the MiVoice Business and the other party in the call.

Call recording

Call recording for EHDU calls is supported beginning in MiCloud 3.2. EHDU calls connect via SIP trunks to the MiVoice Business, and for external calls, all call legs are carried on SIP trunks.

SIP connectivity (Teleworker)

Remote users may connect to UC services using IP phones over public or private data networks through the MiVoice Border Gateway Teleworker service. This service implements a SIP back-to-back user agent (B2BUA) or MiNet proxy to connect the remote IP phone to the MiVoice Business IP phone system. The Teleworker phone is registered as a standard extension of the office phone system, providing full access to voice mail, collaboration tools and all features of the system. Most Mitel IP phones, including SIP and MiNet hard desk phones and soft phones, may be configured for Teleworker or normal mode of operation.

An interesting use case is the cellular phone with mobile SIP client which may access the UC services in two modes, namely as an EHDU device connecting via the mobile phone over the cellular wireless voice network and as a Teleworker device connecting via the mobile SIP client over an available wireless data network. The choice of EHDU or Teleworker depends largely on the availability and cost for the access network. When using the MiCollab for Mobile with Premium UC license, the Teleworker mode offers the advantage of handing off between the cellular wireless data network and local Wi-Fi networks. On the mobile client, wireless network availability and signal strength changes trigger enabling the command to hand off between networks.

When using the MiVoice Border Gateway Teleworker service with a Mitel IP phone, the following capabilities are provided:

- Encryption to improve voice path security
- Adaptive jitter buffering and other enhancements to improve voice quality
- G.729 compression to reduce bandwidth requirements

For larger deployments, MiVoice Border Gateways may be clustered to support larger numbers of Teleworker users. Dynamic load balancing across cluster members is supported when both MiVoice Border Gateway and the device support redirection. MiVoice Border Gateway supports redirection of MiNet devices, and most MiNet desk phones support redirection. SIP devices and MiNet soft phones must be manually load-balanced by configuring the device to connect to a specific MiVoice Border Gateway node.

Resiliency available for Teleworker devices depends on the capabilities of MiVoice Business, MiVoice Border Gateway, and the device. For MiNet devices, MiVoice Border Gateway maintains a list of available PBX platforms. On failure of the primary PBX, all sets disconnect and attempt to re-register. For MiNet devices that support redirection, when the devices re-connect, the MiVoice Border Gateway redirect the device to an available PBX. SIP devices and devices that do not support redirection cannot use this mechanism. For devices that support multiple registration configurations, resiliency is provided through manually configuring multiple MiVoice Border Gateway addresses, or via FQDN and DNS redirection. On failure of the path through the primary MiVoice Border Gateway to the primary MiVoice Business, the device attempts to register with a secondary MiVoice Border Gateway. If the device does not support either redirection or multiple registration capability, the device will be unable to make calls until the primary MiVoice Border Gateway and/or primary PBX returns to service.

Table 14: MiVoice Border configurations

CONFIGURATION	SUPPORTS TELEWORKER SERVICE
MiVoice Border Gateway[1] • DMZ mode • Server-Gateway mode MiVoice Business Express • Deployed with MiVoice Border Gateway in server-gateway mode	√
MiVoice Border Gateway[2] • DMZ mode: deployed on internal LAN configured for LAN mode operations (e.g. Secure Recording Connector) • Server-Gateway mode: deployed on internal LAN configured for LAN mode operations (e.g. Secure Recording Connector) MiCollab • Deployed on an internal LAN with MiVoice Border Gateway in LAN mode	X

Notes:

1. See “Deployment considerations” on page 80.
2. Similar considerations apply when the MiVoice Border Gateway is integrated with other applications.

Call recording

For recording Teleworker devices, the gateway MiVoice Border Gateway is configured to route calls to a second LAN-based MiVoice Border Gateway serving as the SRC, which connects to the MiVoice Business controller.

Licensing

MiVoice Border Gateway Teleworker service is licensed per device. If multiple MiVoice Border Gateways are deployed, the licenses may be shared across clustered nodes to support dynamic load balancing of MiNet devices. With UCC licensing, MiVoice Border Gateway Teleworker licenses are included within the Standard and Premium User licenses and also are available as à la carte add-on options.

Data connectivity (Web Client)

The MiVoice Border Gateway Web proxy implements a reverse proxy with URL mapping. The proxy provides a secure method for end-user desktop and web clients to connect with LAN-based applications. The web proxy restricts access to only those URLs that belong to the end-user web interfaces for the recognized applications. The current release supports MiCollab and MiCollab Audio, Web, and Video Conferencing desktop and web clients.

The MiVoice Border Gateway web proxy also supports remote access to native management web pages for both MiVoice Business and MiCollab. For MiVoice Business management access, MiVoice Border Gateway adds additional access control security by requiring user name and password; also management access may be restricted to listed client IP addresses.

To provide a resilient web proxy service, redundant MiVoice Border Gateway servers may be deployed. The client applications can be configured to access the redundant MiVoice Border Gateways, either directly or via multiple DNS entries. For MiCollab Clients, the connection through MiVoice Border Gateway to the MiCollab server is not resilient. Application layer resiliency may be achieved through deploying MiVoice Border Gateway and MiCollab on virtual servers configured with high availability.

The reverse proxy acts between an Internet accessible server and Internet-protected LAN server. It is not required if the application server is deployed in Network Edge mode with direct Internet access or if the web client has direct LAN access to the application server.

MiVoice Border Gateway web proxy is included in the MiVoice Border Gateway base software license. No separate license is required.

Data connectivity (LAN extension)

For UC communications, a commonly encountered scenario is multiple end-points connected on a common voice LAN on the customer site, and MiVoice Business servers on a LAN in a remote data center. Geographically extending the Layer 2 network between these sites allows the UC clients to connect directly to the UC servers without the need to deploy and provision proxies and firewalls. Typically these LAN extension connections provide higher QoS, based on SLAs and/or the underlying transport mechanism.

There are a number of technologies and commercial offers available to provide LAN extension, such as dedicated VPNs or MPLS connections. The choice is based on availability, service level agreements, and cost. Typically, the addressing scheme and required networking equipment is determined by the choice of carrier.

Within the UC solution, these remote end-points appear to be local and are accessible via Layer 2 addressing schemes. UC design considerations include ensuring sufficient bandwidth and QoS mechanisms are available to support the expected traffic.

When LAN access is available, this is the preferred method of connection between fixed UC clients and servers. For mobile clients, there are two possible configurations – normal mode directly accessing the call control engine or Teleworker mode accessing the set-side of the Teleworker gateway. For most Mitel IP phones, changing between normal and Teleworker mode requires manual configuration of the mode. MiCollab Mobile Client requires changing the IP address used for registration (MiCollab Client supports only a single IP address). Dual mode hand-off between Wi-Fi and LTE/4G networks, available for premium UC users, requires that the active call be anchored at the MiVoice Border Gateway, because SIP redirection between MiVoice Business and MiVoice Border Gateway is not supported. For ease of use and to avoid these manual steps and dual-mode hand-off, it is recommended that all mobile clients be configured for Teleworker mode, both within the LAN and externally.

In the customer site, routing rule exceptions may be used to route Teleworker connections over the QoS-enabled LAN extension path. Teleworker devices must connect to the set-side of the MiVoice Border Gateway, which may be either an Internet accessible WAN address or a DMZ address behind a firewall with NAT. In either case, sets register to a public WAN address. Standard routing directs these connections to the customer's Internet access gateway for connection across the public Internet. To avoid this path, manually configured rules may be added to the routing tables in the customer site and data center to ensure that this traffic is directed to the MPLS egress devices such that the MiVoice Border Gateway WAN address routes via the QoS enabled path.

Data connectivity (Other)

UC communications requires access connectivity for various other types of traffic, not described above, such as the TCP connections for presence and collaboration features. Also, the access network requires connectivity for non-UC traffic such as hosted business applications. External-facing firewalls and routers must be configured to allow UC and application traffic, both at the data center and at the customer site.

For Mitel application clients, the MiVoice Border Gateway includes templates and forms that simplify configuration when used in Server-Gateway mode. When MiVoice Border Gateway is deployed in the DMZ or in-line with third-party firewalls, communication paths must be manually provisioned through external-facing firewalls and routers.

Most Mitel management clients gain access through the MiVoice Border Gateway web proxy described above. In a service provider data center, manual provisioning is required for the scenario in which Oria is deployed in the service provider domain and is used to manage UC components in the customer domain. Specifically, Oria requires 1:1 NAT mapping between virtual addresses in the service provider address space for these managed components, which include MiVoice Business, MiCollab, and MiVoice Border Gateways, and the component addresses in the customer space and configuration to allow both web services and THRIFT protocol connections between Oria and UC components. Oria networking is discussed further in "Network and networking considerations" on page 89. Clients access the Oria portal via web proxy through a MiVoice Border Gateway or a third-party proxy.

MarProbe is a data collection agent for MarWatch, available as software or as a custom appliance. It may be installed in the data center or the customer site to monitor device and network performance. MarProbe connects to the MarWatch server using HTTP secured with SSL. MarProbe always initiates IP connections with MarWatch, and this outbound connection passes most firewall rules without specific configuration. Client access to MarWatch portal is available via a third-party proxy or firewall.

At the data center, any firewall in-line with the MiVoice Border Gateway in Server-gateway mode should not implement NAT. At the customer site, the firewall is expected to implement NAT, providing a unique port for identifying individual end-points.

For firewall protocol and port configuration rules for many Mitel applications, see the *MiVoice Border Gateway Engineering Guidelines*. Protocol and port requirements are included in product-specific engineering guidelines.

Deployment considerations

Network design for UC components in the data center and customer site has many options, depending, in part, on the type of access network being used, the service provider peering network, and the network capabilities available from an IaaS provider. This section includes design considerations that are specific to UC components, but is not intended as a comprehensive network analysis.

Access networks

Access between UC clients and UC servers located in the data center typically requires WAN connections. These connections may be over the PSTN, the Internet, or over private data paths. Figure 12 shows a high level view of possible access networks.

Figure 12: UC Access Networking

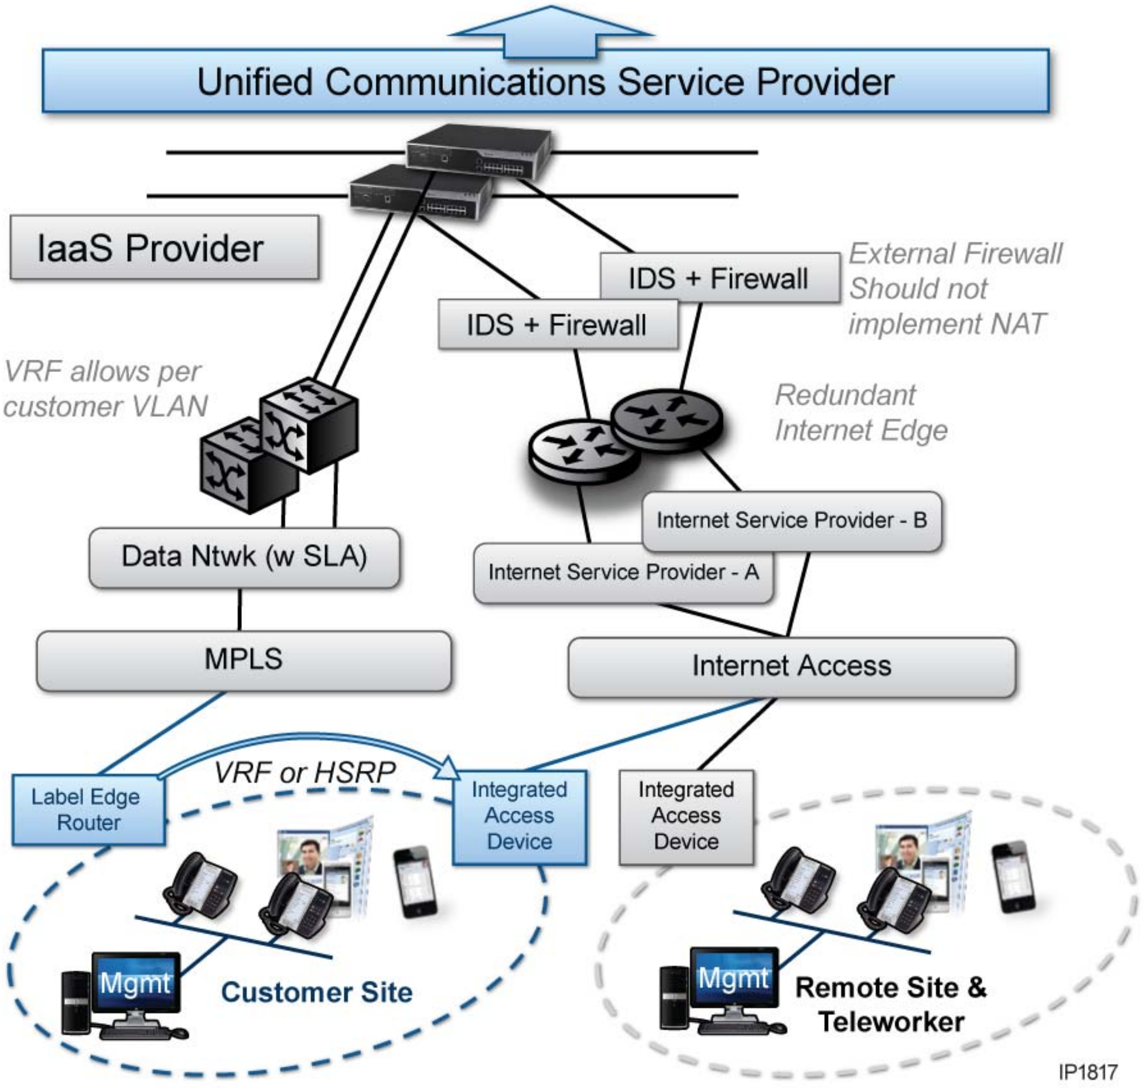

Customer sites connect over MPLS or other LAN extension technologies, shown with the Label Edge Router (LER), or over public data networks, shown with the Integrated Access Device (IAD). Resilient access networking may be achieved with dual connections and redundant networking equipment. In the case that the primary path is made via an MPLS connection with backup via the Internet, some means is required to trigger re-establishing connectivity through the back-up path. Two possible implementations are:

- A multi-service edge router supporting both MPLS and IP networking, which automatically manages the link states
- Implementing Hot Standby Routing Protocol (HSRP) between the access gateways to manage the network re-routing.

Using redundant edge routers provides improved resiliency, particularly for the multi-service router scenario.

As discussed earlier, customer site networking equipment requires a routing exception to route Teleworker devices via the LER, rather than via the IAD. Routing over either device will work; the LER is expected to provide higher QoS than the IAD.

For the data center, there may be an in-line firewall outside of the control of the UC service provider. It is recommended that this external firewall NOT implement NAT. Network design and configuration is simplified if the externally visible IP addresses are configurable by the UC service provider.

MiVoice Border Gateway provides basic firewall capabilities, such as port blocking and forwarding based on Mitel Standard Linux static routing tables. The services implemented effectively provide application layer protection from distributed denial-of-service (DDoS) attacks. Acting as a firewall, MiVoice Border Gateway processing speed is sufficient for most sites as traffic is limited by the WAN pipe. An in-line third-party firewall provides additional protection, including Intrusion Detection System (IDS) and Intrusion Prevention System (IPS) with signature-based updating, UDP flooding protection, advanced port blocking and forwarding capabilities, and advanced reporting. An in-line IDS/IPS system can provide a measure of network layer DDoS protection. Any in-line firewall should disable SIP adaptation as it conflicts with MiVoice Border Gateway SIP handling. And, as mentioned above, any external firewall on the data center edge should not implement NAT.

Data center edge

The principal external gateways for UC systems are the 3300 ICP PSTN gateway and the MiVoice Border Gateway. The 3300 ICP is the service provider gateway for TDM trunks. The MiVoice Border Gateway provides both the service provider gateway for SIP trunks and the access gateway for application and management end-points.

The 3300 ICP provides both the PSTN gateway and call control engine. For smaller sites, both functions may be implemented with a resilient pair whereas for larger sites, these roles may be divided between user controllers and trunking gateways. MiVoice Border Gateway supports two deployment modes, Server-Gateway and Demilitarized Zone (DMZ). For a typical mid-size hosted UC deployment, example network deployments are shown below. Figure 13 shows a service provider network with MiVoice Border Gateway Virtuals in Server-Gateway mode.

Figure 13: Service Provider Medium Large Business External Networking

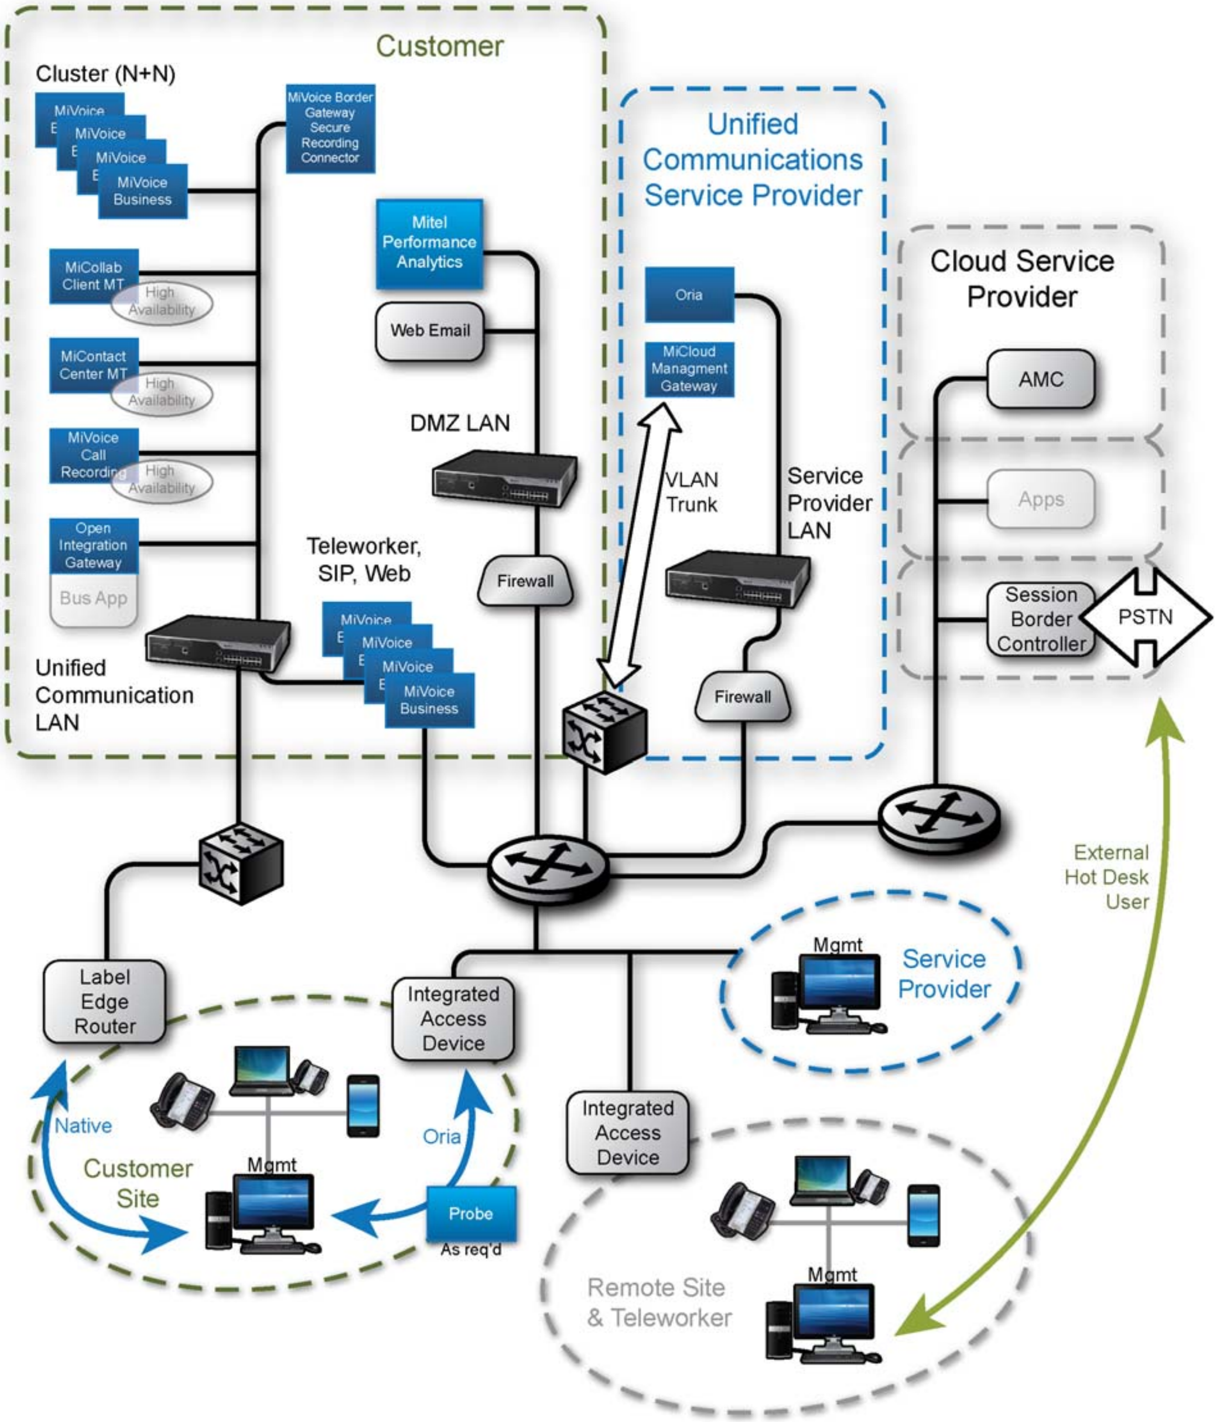

Also shown in the service provider graphic is the Oria routing constraint—Oria requires 1:1 NAT access to the managed UC components. This constraint exists for all single-tenant topologies, including SMB-LD.

For Enterprise deployments, MarWatch server may be deployed in the customer network, or accessed as a cloud service. In the Managed Service Provider Program, MarWatch can be deployed per-customer; MarWatch cloud-based services are not supported. The MarWatch server may be deployed in the DMZ or core LAN with web proxy via a third-party firewall. MarProbe is deployed wherever network health and monitoring is required, typically the customer site and core UC service provider network.

The PSTN is effectively a private network with well-controlled access. Networking for PSTN access to the 3300 ICP is based on dedicated connections—typically PRI or T1. The decision to deploy TDM gateways is largely based on considerations for resiliency, available service provider agreements, and existing equipment. For more information about PSTN networking and 3300 ICP as trunking gateway, refer to the *MiVoice Business Technician's Handbook*.

In both examples shown, MiVoice Border Gateway provides a SIP trunk proxy, Teleworker service, and web proxy.

When deployed as Server-Gateway, it also implements firewall and static routing capabilities. In Server-Gateway mode, the MiVoice Border Gateway is shown in parallel with a third-party firewall. The MiVoice Border Gateway provides access to the UC components, and a third-party firewall provides access to other business applications. This is the recommended deployment for Server-Gateway mode for large sites.

For smaller sites, MiVoice Border Gateway may act as the firewall for all of the core LAN servers. External addresses are required for each MiVoice Border Gateway, either with NAT to the DMZ address of the MiVoice Border Gateway or directly accessible on the MiVoice Border Gateway WAN interface.

The two major factors in choosing the MiVoice Border Gateway deployment mode are:

- the existing IT networking and security policies
- the existing network infrastructure

Other factors to consider include:

- Security zones increase overall security of the network by confining attacks and breaches within a zone. The DMZ provides a security zone, as does separating the UCaaS servers from the other business servers in the Server-Gateway mode.
- The firewall port blocking and forwarding is necessarily similar for DMZ and Server-Gateway modes because firewall rules are required to pass the UC traffic. Whether ports are opened in the firewall managing the DMZ or the same ports are allowed in the MiVoice Border Gateway in Server-Gateway mode has limited impact on overall security;
- Networking equipment costs between the two options depend on the existing infrastructure. DMZ mode requires equipment to implement the DMZ, while Server-Gateway in parallel with the business firewall requires networking equipment from the Internet access drop to the MiVoice Border Gateway. DMZ mode deployed into an existing DMZ is the most economical.
- Configuration and management, particularly trouble shooting, is easier with Server-Gateway mode. MiVoice Border Gateway provides built-in templates for Mitel products. In Server-Gateway mode, MiVoice Border Gateway owns the WAN address, avoiding unex-

pected routing to reach a firewall-owned address. Server-Gateway mode avoids firewall configuration issues such as UDP port mismatch.

- Visibility and control, both monitoring and reporting, is typically more comprehensive with DMZ mode.

The two modes of operation, Server-Gateway and DMZ, affect MiVoice Border Gateway addressing and SIP conversion. Beyond these modes, MiVoice Border Gateway also offers various customization parameters to adapt to non-standard configurations. A custom profile may be used to override the default streaming addresses. IP translation tables may be used to translate IP addresses in call set-up, particularly useful for streaming between MiVoice Border Gateways on the same LAN.

Emergency Location (E911)

It is very important to ensure proper emergency call location is available to first responders.

Locating an end-user on a hosted solution can prove difficult, and is especially crucial when used with emergency location systems, such as E911.

If the business is small enough, the business location may be identifiable simply from the assigned DID/DDI number. This will be registered with the SIP trunk service provider.

However, to locate an individual, especially where the business is physically large, or covers multiple sites, becomes more difficult. To assist in this situation, every user is assigned a DDI/DID number and the location of that individual is maintained in the SIP trunk service provider database and local PSAP authority.

It is important that the customer understand that when end-users change locations, they must inform the hosted service provider of this change so that the SIP service provider can be updated.

Before installation, the service provider must ensure that the location information stored with the SIP Trunk provider can be forwarded to the appropriate PSAP to service the customer. With the advent of IP networks, this is becoming easier, but some physical locations may not subscribe to this service, making E911 location services difficult for the SIP trunk provider. This is especially important where the SIP Trunk provider may host a number of DID/DDIs from many locations, but physically connect to the PSTN at a few major points of presence.

If a survivable gateway has been deployed on the customer premise, the gateway may also use PRI trunks for emergency calling. The MiVoice Business can include location information to be sent over these trunks. See *MiVoice Business Engineering Guidelines* and on-line help for further details.

Chapter 5

NETWORK AND NETWORKING CONSIDERATIONS

Network and networking considerations

The following sections discuss networking, addressing, and QoS for the three main architectures.

- “Small Business architecture” on page 89
- “Small Medium Business architecture” on page 90
- “Medium-Large Business architecture” on page 97

The sections that follow discuss the networking considerations that affect all topologies, and the supported devices:

- “Networking and QoS” on page 105

Small Business architecture

All end customer access to the multi-customer applications is through network routed addresses. This does not preclude the use of MPLS or carrier managed networks with appropriate configuration of the end-customer access gateway.

The MiVoice Business Multi Instance server may be configured for VLAN or non-VLAN operation. The requirement is that the MiVoice Business Multi Instance server and the related multi-customer Mitel applications are reachable within a local area network. As such, the architecture supports the use of VLAN mode, provided that all applications reside in the same VLAN. When the MiVoice Business Multi Instance server is used in VLAN mode, the associated LAN switch must apply Layer 2 tagging at the port level to provide access from the untagged LAN to the tagged LAN.

SB QoS considerations

Quality of Service (QoS) must be provided end-to-end, from the user phone through to connection to the service provider. Consider the QoS settings on the Public Network connections, the customer access connection, and the customer LAN.

Where the end-customer and the service provider share a common Public Network, the service provider may provide traffic shaping and priority queuing through Service Level Agreements to ensure that priority marked packets are handled appropriately.

Peering arrangements between peer-to-peer Border Gateway Protocol (BGP) connected Public networks may also allow QoS to be honored between these networks, rather than over the Public Internet.

Although connections over the Public Internet are possible, SLAs and QoS are not typically honored due to uncertainty of the connection route over different networks.

It is recommended that customers provide a means of enabling QoS and queuing mechanisms within their own network. It may also be necessary to provide a dedicated voice connection; for example, to provide the necessary bandwidth and avoid access congestion with basic Internet routers. Traffic shaping units can also be included in the customer connection, either by the Public Network Provider, or installed on the customer site. These effectively throttle unmarked download traffic to provide sufficient bandwidth for voice traffic.

Connections that use the Public Network infrastructure are called Over-the-Top (OTT) deployments, since they use the Public IP addressing scheme, rather than dedicated connections that are customer and network-specific.

SB addressing considerations

The service provider address space is a self-contained flat network and is totally isolated from the customer networks.

Customers have their own local networks and private addressing schemes, many with overlapping IP addresses.

The service provider provides a public IP address on the user gateways. Users connect to the user gateways using their own public IP addresses provided by the carrier, or statically assigned IP addresses at the customer router. The customer router provides the network address translation from internal LAN address to external public address. The MiVoice Border Gateway translates customer public IP addresses to appropriate internal service provider addresses.

Management access for the end-customer is available via a Public IP address associated with the service provider. The management portal authenticates access and translates connections to appropriate internal service provider addresses. If the service provider deploys a Mitel MarWatch Server within the hosted address space, an additional public IP is required for this service. This allows the Mitel MarWatch probes to connect to the central hosted server.

Small Medium Business architecture

The following sections highlight some considerations that apply to the Small Medium Business architecture.

SMB QoS considerations

Quality of Service (QoS) must be provided throughout the customer premise LAN and the service provider Hosting LAN, across the MPLS network connecting the customer premise to the hosting network, and across the connection to the SIP trunk service provider.

Service Level Agreements should be established with the SIP trunk service provider to ensure that priority marked packets are handled appropriately.

Teleworker connections are made over the public Internet, and SLAs and QoS are not typically honored due to uncertainty of the connection route over different networks. In the absence of SLAs, it is recommended that the customer provide a means of enabling QoS and queuing mechanisms within their own network and within the Teleworker's local network. Traffic shaping units can also be included in the customer connection, either by the Public Network Provider, or installed on the customer site. These effectively throttle unmarked download traffic to provide sufficient bandwidth for voice traffic.

SMB addressing considerations

The customer's local network and the service provider's network that hosts applications for the customer use a private IP addressing scheme. The SIP trunk provider and the Internet service provider use the same public IP addresses to address the MiVoice Border Gateways. The MiVoice Border Gateways must use statically assigned IP addresses. Teleworker users connect to the Teleworker gateway using their own public IP addresses provided by the carrier, or statically assigned IP addresses at the customer router.

The service provider's router also provides the network address translation from internal LAN address to external public address. The MiVoice Border Gateway translates customer public IP addresses to appropriate internal IP addresses.

SMB end-customer, Service Provider, and network considerations

The following sections are not exhaustive, and are intended to provide guidance for reviewing this architecture and ensuring that significant details have not been overlooked.

End-customer considerations

- Customer Edge router for connection to the MPLS network
- Firewall/NAT proxy for connection to external networks
- QoS enabled LAN networking to provide high quality voice service

Note: E911 services are expected to be provided by the SIP Trunk provider based on end-user DID numbers.

Typically, the end-customer is responsible for maintaining and updating the E911 location information. This is particularly important for Teleworker users who are not located in the customer’s office locations.

Network considerations

MiCloud Business is designed with the customer’s address space extended into the hosted environment. Customer network isolation is achieved using VLANs and typical network isolation techniques, for example: firewalls and routing rules. Overlapped customer address spaces are allowed within these isolated networks.

Networks carrying voice must support QoS mechanisms to provide high quality voice services. Use of QoS mechanisms allow networking equipment to:

- Traffic shaping based on layer 2 QoS tagging and/or layer 3 priority
- Segregating voice and data network traffic into separate queues

To ensure good voice quality, it is recommended that you over-provision bandwidth capacity on LAN and WAN connections.

Mitel MarWatch may be used to troubleshoot voice quality issues.

Hosted components require network connectivity to:

- Third-party SIP trunk providers

  Connection to the SIP trunk providers is accomplished via multiple MiVoice Border Gateways, which should be configured with 1:1 redundancy. The MiVoice Border Gateway provides Session Border Controller (SBC) functionality to manage interoperability with third-party SBCs. The MiVoice Border Gateways should be deployed as a cluster to allow sharing of SIP trunk licenses. Load balancing of outgoing connections is based on MiVoice Business Automatic Route Selection (ARS) tables. Load balancing of incoming connections will depend on the third-party SBC capabilities. Overall load balancing should be arranged to share the load across the MiVoice Border Gateways in the cluster. Connection to multiple third-party SIP trunk providers is supported. Multiple providers may be used to improve reliability or select among alternate long distance providers.

- Customer on-premise networks

  The SMB architecture uses a dedicated MPLS connection to the end-customer sites. This MPLS connection allows virtual routing and forwarding (VRF) such that customer end-points securely access the LAN deployed MiVoice Business and MiCollab. End-points are connected to the LAN side of the server gateway. Deploying MiCloud Business Solution topologies with public network connectivity to customer sites is possible; the main concern is the ability to satisfy QoS considerations. For further guidance on possible deployment with public networks, contact Mitel Professional Services.

- Remote client access

  Remote clients access call control and applications over the Internet with connectivity through the MiVoice Border Gateway providing Teleworker service. The MiVoice Border Gateways should be clustered to provide license sharing, MiNet end-point load balancing, and resiliency. As MiCollab Clients do not support redirection by the MiVoice Border Gateway, these clients cannot be load balanced; MiCollab Client resiliency is achieved with VMware high availability.

- External cloud services

  Integration to cloud services is done through the Mitel Open Integration Gateway. The service provider should ensure that sufficient bandwidth is available to support the expected end-customer transaction traffic.

Service Provider considerations

- Application Management Center – Mitel components require web service access to the license server for operational verification
- Mitel Performance Analytics Probe– web-based service for fault and performance monitoring
- E-mail forwarder for alarms and notifications
- Public DNS registrations for MiVoice Border Gateways

SMB-LD considerations

Considerations for the SMB-LD architecture include management access, application programming, and licensing interactions. There are also scaling considerations based on the number of users and the number of Public IP addresses that are required to handle the number of customers.

The SMB-LD deployment also allows full UCC connectivity for those customers that require it, without the service provider having to provide this to all customers.

The MPLS connection also allows access to other hosted applications, and provides the capability for a local trunk breakout MiVoice Business unit.

Use of an MPLS connection allows SLA and QoS settings to be honored across the carrier network.

SMB-LD component interactions and dependencies

The SMB-LD architecture is intended to cover the Small to Medium business market, although it can scale to a higher number of users and devices. The main factor that limits the scaling is the number of users that the MiCollab Client Multi-Tenant Service component is capable of handling. The number of MiVoice Border Gateways and MiVoice Business controllers are also driven by the number of users. Typically, the SMB-LD is deployed with up to 250 users, and associated devices, although this can scale up to thousands of users. Scaling details are discussed in “Small Medium Business - Low Density scaling” on page 52.

The core components of the architecture are the MiVoice Business and the MiVoice Border Gateway Virtual. MiVoice Business and MiVoice Border Gateway Virtuals are needed to provide voice connectivity, even if no other applications are required. Use of MiVoice Business Multi Instance allows high density of deployment and a cost-optimized solution for those customers that require basic voice telephony.

The SMB-LD architecture allows the deployment of a 3300 ICP remote trunking gateway at the customer site because the end-customer network is directly connected to the hosted site, rather than via NAT and a public gateway. Programming the remote trunking gateway requires manual interaction, as it cannot be provisioned via Oria. Manual programming creates a risk of database synchronization conflicts. Consult Mitel Professional Services before deploying a remote gateway.

The SMB-LD architecture can be combined with the SB architecture. In this arrangement, the MiVoice Business Multi Instance may be shared across both architectures, with some instances in the shared network with MiCollab Client Multi Tenant providing light UC services, and some instances in isolated networks with MiCollab instances offering rich UC services. The MiVoice Business Multi Instance must be operating in VLAN mode, with the management address on the untagged VLAN, the shared instances' addresses on a common VLAN, and the isolated instances on unique VLANS. For the other platform components, their deployment is as per the relevant architecture.

The MiVoice Border Gateways provide the following functions:

- SIP Trunk connections to SIP service provider

- Teleworker phones with Mitel proprietary phones
- UC connections for mobile UC clients and SIP soft phones

Scaling of the MiVoice Border Gateways is discussed in "MiVoice Border Gateway scaling" on page 60.

The service provider must consider the number of IP addresses to be provisioned on the MiVoice Border Gateway Virtuals across a number of customers. The service provider requires multiple public IPv4 addresses for the end-customers.

For a core voice-only deployment without Teleworker phones, the MiVoice Border Gateway Virtuals are the only virtualized product within the deployment used only to connect to a SIP trunk service provider. In such a case, an alternative to multiple MiVoice Border Gateway Virtuals is a third-party VLAN-aware SBC. See the Mitel SIP Center of Excellence (SIP CoE) for the SBCs that are suitable for this function.

Where Teleworker users or UC functionality is needed, MiVoice Border Gateway Virtuals must be used to provide the necessary Application Level Gateway (ALG) functions between the public external and customer internal networks. Third-party SBCs do not work for these connections.

Inclusion of the MiCollab Virtual server provides access to the full UCC capability. The MiCollab server needs to be linked to the MiVoice Business for the customer deployment. The number of users and type of users will determine the scaling of the MiVoice Business and MiVoice Border Gateway infrastructure.

A MiCollab server must be deployed per customer for each customer requiring UCC capabilities. The tenanted MiCollab solution is not VLAN-aware and will not work across multiple customers in this architecture.

Different voice mail deployments exist depending on whether the customer is voice-only, or if MiCollab Virtual is deployed for UCC functions. If the customer is voice-only, then the internal MiVoice Business embedded voice mail is used. If MiCollab Virtual is deployed, then the MiCollab Unified Messaging is used. MiCollab Unified Messaging provides additional UCC functions, including visual voice mail and e-mail integration. Different provisioning templates in Oria apply, depending on whether the customer is voice-only or uses MiCollab capabilities.

The MiCollab Client can be associated with mobile devices such as smart phones and mobile PC tablets. These mobile devices must register through the customer MiVoice Border Gateway Virtual if full mobility outside of the business premise is required. The MiVoice Border Gateway must be scaled for these external connections.

Where there are MiCollab Clients that will be external to the customer network, the MiVoice Border Gateway Virtual provides additional Application Level Gateway (ALG) functionality that is not provided by a third-party SBC, so MiVoice Border Gateway Virtuals are required with a UCC deployment.

Access to hosted cloud services, such as Salesforce and Google, is provided by a per customer Mitel Open Integration Gateway Virtual application. An Internet connection per customer is needed for this integration to work.

SMB-LD QoS considerations

The availability and use of an MPLS circuit means that the connection between the customer and the hosted site is managed. This connection can have SLAs applied to it to ensure timely delivery of voice and signaling packets.

DSCP values must be consistent between the customer end-devices and the service provider's hosted equipment. Performance is improved if the customer also employs consistent DSCP markings in their network equipment.

Typically the data bottleneck is at the incoming connection at the customer premise, especially if this connection is shared with an Internet connection. Enabling QoS scheduling at the penultimate (next to last) router, at the provider edge, helps to ensure that voice and associated signaling are delivered to the required SLA.

SMB-LD addressing considerations

The service provider IP address space is a self-contained flat network. The service provider's network is completely isolated from the customer networks.

Hosted customers have their own local networks and private addressing schemes. The customer-hosted IP addresses may overlap, because the service provider VLANs are used to differentiate the different customer networks within the hosted infrastructure.

There are some restrictions of IP addresses that the customers can present to the hosted space on a particular MiVoice Business Multi Instance server. The same customer IP addresses and service provider addresses cannot be used on the same MiVoice Business Multi Instance. In practice, the service provider will be using 10.0.0.0/8 addresses and the customers will be using 192.168.0.0/16 addresses. Details of these restrictions are described in the *MiVoice Business Multi-Instance Engineering Guidelines*.

The service provider configures a public IP address on each customer external gateway. Due to the large number of customers, a large number of public IP addresses must be available. The IP addresses must be statically assigned to the gateways.

For management access into the customer hosted space, a 1:1 NAT router is needed for each customer. The router translates between a private service provider address and the private customer addresses. The router may be a physical device per customer, or it may be a virtual router—part of the virtual infrastructure.

Customer access to the Oria management portal is available though a public IP address at the edge of the service provider network. The Oria management portal authenticates access and translates connections to internal service provider addresses on the NAT routers, and eventually, to hosted units in the hosted customer space.

SMB-LD end-customer, Service Provider, and network considerations

The following sections are not exhaustive, and are intended to provide guidance for reviewing this architecture and ensuring that significant details have not been overlooked.

End-customer considerations

The following items should be considered for end-customers:

- Dedicated MPLS connection to hosted service provider
- Common carrier network with the service provider or carrier peering connections
- Local DHCP on customer network
- Local PSTN access can be provided with a trunking gateway.

Note: Manual management and configuration is required with a trunking gateway.

- Level of UC functionality required and connection to other applications. A different architecture may be more suitable if a customer requires more UC functionality.
- Direct connection to non-standard interfaces (example: SIP ATA for door openers)
- Internet access is needed from the customer network for management and license verification

Network and carrier considerations

The following items should be considered for the network and carrier portions of the solution:

- MPLS is needed for connection between end-customer and hosted site
- Ability of the MPLS network provider to provide SLA and honor QoS
- Customer alarm notification to dealer or service provider. Internet service and e-mail addresses are needed for sending alarm notifications.
- Ongoing monitoring and alarm reporting requiring reseller and service provider access.
- SIP device interoperability and SIP trunk provider interoperability must be verified with the SIP-COE
- Access to applications for scheduled backups and storage, including process ownership and data storage location defined
- MiVoice Border Gateway clustering and/or HA based on scaling requirements
- Service provider VLAN and use of VRF router is a requirement for customer isolation and use of overlapped IP addresses

Service Provider considerations

The following items should be considered for the service provider deploying the solution:

- Provision for remote management access (example: web proxy, VPN)

- VPN connection enabled for remote management access
- NAT router and management plane access to each hosted customer network for customer management
- Template deployment for voice-centric deployment (without MiCollab) may differ from the UCC deployment (with MiCollab)
- Performance, instances per server, and traffic considerations are needed for server scaling
- Licensing model (UCC, a-la-carte, service provider, Enterprise)
- Process defined for storage and access to user information to comply with regulations for security and privacy
- Router ACLs for limited access into customer network (example: VRF)
- A defined process for collection and provisioning of billing information from SIP service provider and communication to end-customer
- E911 and emergency location information must be consistent with that at the PSAP, regardless of whether calls are made over SIP trunks or local trunk breakout
- Turnover of service provider personnel and password access updates (example: password expiry, password updates)
- Backup and recovery policies
- Pre-deployment “sandbox” for pending updates
- Inclusion of Windows and VMware components and licensing
- Data center considerations including power, geo-location, cooling, security, and access
- Maintenance and provisioning staff for local and remote date centers
- Management access to local and remote data centers
- SIP trunk service provider and scaling including consolidation of trunks over many customers
- Pre-provisioning of IP addresses prior to delivery of customer end devices or use of SRC for end-devices
- Overall service resiliency and redundancy risk plan considering single points of failure (example: physical environment issues, power failures, etc.)
- Embedded per instance voice mail for voice-centric deployments or MiCollab for UCC deployments

Service provider needs to supply IP addresses for MiVoice Border Gateways.

Medium-Large Business architecture

Flow Through provisioning changes the way MiCollab is networked with the MiVoice Business cluster. With MiCollab Release 7.2+, MiCollab participates in SDS communications for synchronizing changes with the clustered MiVoice Business controllers.

Flow through provisioning changes the way Oria is used to provision users. With MiCloud 3.x using Flow Through Provisioning, Oria must be configured with the IP addresses of the clustered MiVoice Business controllers, both the address in the service provider address space and the address in the customer address space.

DNS servers in both the service provider address space and the customer address space are still recommended to support browser access via FQDN rather than IP address. DNS is still required to support reach through capabilities to the embedded management interfaces.

The following sections describe considerations for deployment of the MLB architecture.

MLB QoS considerations

Quality of Service (QoS) must be provided throughout the customer premise LAN and the service provider Hosting LAN, across the MPLS network connecting the customer premise to the hosting network, and across the connection to the SIP trunk service provider.

Service Level Agreements (SLA) should be established with the SIP trunk service provider to ensure that priority marked packets are handled appropriately.

Teleworker connections are made over the public Internet. SLAs and QoS are not typically honored due to uncertainty of the connection route over different networks. In the absence of SLA, it is recommended that the customer provide a means of enabling QoS and queuing mechanisms in their own network and in the Teleworker local network. Traffic shaping units can also be included in the customer connection, either by the Public Network Provider, or installed on the customer site. These effectively throttle unmarked download traffic to provide sufficient bandwidth for voice traffic.

MLB addressing considerations

The customer's local network and the service provider's network that hosts applications for the customer use a private IP addressing scheme. The SIP trunk provider and the Internet service provider use public IP addresses to address the MiVoice Border Gateway Virtuals. The MiVoice Border Gateway Virtuals must use statically assigned IP addresses. Teleworker users connect to the Teleworker gateway using their own public IP addresses provided by the carrier, or statically assigned IP addresses at the customer router.

The service provider's router also provides the network address translation from internal LAN address to external public address. The MiVoice Border Gateway Virtual translate customer public IP addresses to appropriate internal IP addresses.

MLB end-customer, Service Provider, and network considerations

The following sections are not exhaustive, and are intended to provide guidance for reviewing this architecture and ensuring that significant details have not been overlooked.

End-customer considerations

- Customer Edge router for connection to the MPLS network

- Firewall/NAT proxy for connection to external networks
- QoS enabled LAN networking to provide high quality voice service

Note: E911 services are expected to be provided by the SIP Trunk provider based on end-user DID numbers.

Typically, the end-customer is responsible for maintaining and updating the E911 location information. This is particularly important for Teleworker users who are not located in the customer’s office locations.

Network considerations

MiCloud Business Solution topologies are designed with the customer’s address space extended into the hosted environment. Customer network isolation is achieved using VLANs and typical network isolation techniques, for example: firewalls and routing rules. Overlapped customer address spaces are allowed in these isolated networks.

MiCloud Management Gateway provides the 1:1 NAT capabilities to allow managed network connectivity between the service provider address space and the customer's address space.

Networks carrying voice must support QoS mechanisms to provide high quality voice services. Use of QoS mechanisms allow networking equipment to:

- Traffic shaping based on layer 2 QoS tagging and/or layer 3 priority
- Segregating voice and data network traffic into separate queues

To ensure good voice quality it is also recommended that you over-provision bandwidth capacity on LAN and WAN connections.

A Mitel Performance Analytics Probe may be used to troubleshoot voice quality issues.

Hosted components require network connectivity to:

- Third-party SIP trunk providers

  Connection to the SIP trunk providers is made via multiple MiVoice Border Gateways, which should be configured with 1:1 redundancy; MiVoice Border Gateway provides Session Border Controller (SBC) functionality to manage interoperability with third-party SBCs. The MiVoice Border Gateways should be deployed as a cluster to allow sharing of SIP trunk licenses. Load balancing of outgoing connections is based on MiVoice Business Automatic Route Selection (ARS) tables. Load balancing of incoming connections depends on the third-party SBC capabilities. Overall load balancing should be arranged to share the load across the MiVoice Border Gateways in the cluster. Connection to multiple third-party SIP trunk providers is supported. Multiple providers may be used to improve reliability or select amongst alternate long distance providers.

- Customer on-premise networks

  The MLB architecture uses a dedicated MPLS connection to end-customer sites. This MPLS connection allows virtual routing and forwarding (VRF) such that customer end-points have secure access to the LAN-deployed MiVoice Business and MiCollab. End-points are connected to the LAN side of the server gateway. Deploying MiCloud Business Solution architecture with public network connectivity to customer sites is possible; the main concern is the ability to satisfy QoS considerations. For further guidance on possible deployment with public networks, contact Mitel Professional Services.

- Remote client access

  Remote clients access call control and applications over the Internet with connectivity through the MiVoice Border Gateway providing Teleworker service. The MiVoice Border Gateways should be clustered to provide license sharing, MiNet end-point load balancing, and resiliency. MiCollab Clients do not support redirection by the MiVoice Border Gateway, so these clients cannot be load balanced. MiCollab Client resiliency is achieved with VMware high availability.

- External cloud services

  Integration with cloud services is accomplished using the Mitel Open Integration Gateway. The service provider should ensure sufficient bandwidth is available to support the expected end-customer transaction traffic.

Service Provider considerations

- Application Management Center – Mitel components require web service access to the license server for operational verification
- Mitel Performance Analytics Probe – web-based service for fault and performance monitoring
- E-mail forwarder for alarms and notifications
- Public DNS registrations for MiVoice Border Gateway gateways

Networking the applications

This section discusses the networking considerations specific to each application.

MiCloud Management Gateway

The MiCloud Management Gateway is optimized for deployments spanning multiple VLANs. Specific networking requirements include:

- Customer networks with overlapping addresses are supported.
- Customer networks require unique VLANs accessible via a single VLAN trunk. A single customer may have multiple VLANs.
- The service provider network must have non-overlapping addresses in a single tagged or untagged VLAN. If the service provider uses a VLAN, it must be unique from the customer VLANs.
- All service provider applications connecting through the MiCloud Management Gateway must operate from trusted networks of the management plane gateway. (The MSL management interface is used to specify the trusted networks.)
- The MiCloud Management Gateway requires an address in the service provider network; that is, the MSL system address for the application.
- For connections initiated to customer end-points, referred to as "southbound," addresses must be allocated in the service provider network space for each customer end-point, and a single primary service address must be allocated in the customer network. Multiple management applications can use the same primary service address of the MiCloud Management Gateway.

  For connections initiated to service provider end-points, referred to as "northbound," service addresses must be allocated in the customer network space for each service provider end-point. Customer end-points must be configured as destination addresses with associated addresses in the service provider network.

  Port translation is not performed.
- MiCloud Management Gateway supports IPv4 only.
- MiCloud Management Gateway requires two network interfaces:
  - One using a tagged or untagged VLAN addressed in the service provider network
  - One using a VLAN trunk to reach customer networks

To ensure network security, MiCloud Management Gateway provides filtering based on IP address, protocol, and port. More specifically, all traffic from the management network is blocked unless the source IP is in the local networks list, and all traffic from the customer network is blocked unless the source IP and TCP/UDP port or ICMP is configured (in the Services tab of the MMG configuration interface in MSL). All traffic from one customer network to another customer network is blocked.

For information and configuration instructions, see the MiCloud Management Gateway Help.

MiContact Center Business

MiContact Center Business networking requirements include connections to:

- MiVoice Business controllers using IP trunks to trusted 5020 ports for IVR media streaming capabilities (using MiTAI for CTI events and MiContact Center Business routing capabilities).
  - For larger contact centers with two tier MiVoice Business architecture, IVR servers connect to 5020 ports on the Queuing gateways and MiContact Center Business server connects (using MiTAI) to the Agent controllers.
- Remote servers require connectivity to the Contact Center Business server for configuration, statistics, media files, and software updates.
- Contact Center desktop clients using standard sockets
- Contact Center web clients using HTTP or HTTPS
- External e-mail servers using IMAP for report distribution and alarm notifications.
- Multi-media Contact Center requires connectivity to the media sources:
  - IMAP and SMTP (or Secure IMAP and Secure SMTP) to the e-mail provider. Multi-media Contact Center supports both premise-based servers, such as Microsoft Exchange, and cloud services such as Google Gmail and Microsoft Exchange Online.
  - HTTP or HTTPS to the web server, Microsoft IIS and Apache. for hosting the customer facing chat window for real-time web chat
  - HTTP or HTTPS to the SMS provider web services API for SMS messages

For remote desktop and web clients, MiVoice Border Gateway provides pre-configured connectors to support Contact Center clients.

MiVoice Call Recording

For MiVoice Call Recording, networking requirements include MiVoice Call Recording Server connections to:

- Off-board SAN or NAS storage and off-board SQL server, when used
- MiVoice Border Gateway-SRCs using a control interface over TCP, and audio streams (G.711 or G.729) over RTP/UDP
- MiVoice Business controllers using MiTAI for CTI events
- MiContact Center Business control interface over TCP for Contact Center Client integrated recording. This is supported only for single tenant (MLB) deployments

- Desktop clients using standard sockets. Audio playback is streamed in GSM mono or G.729 stereo.
- Web clients using HTTP or HTTPS with audio playback using MP3 format.

For remote desktop and web clients, access to MiVoice Call Recording must be configured through a third-party router and firewall. MiVoice Border Gateway does not support MiVoice Call Recording remote client connections.

MiCloud Business Analytics

For Mitel Business Analytics, connections are required to:

- Delivery Controller to each MiVoice Business for SMDR, using TCP/IP output streaming with the Delivery Controller acting as a TCP/IP client
- Delivery Controller to the cloud server using TCP, with an encrypted connection through a third-party firewall

For multi-tenant architectures, the collector node may be networked to multiple MiVoice Business controllers, all residing in the shared customer address space. For single tenant architectures, the collector node is deployed in the service provider address space, and connects to MiVoice Business controllers in the customer address space using the MiCloud Management Gateway.

Delivery Controller

The Delivery Controller runs on the Windows operating system and is supported on:

- Microsoft Windows 7, 8
- Microsoft Windows Server2008 R2 SP1, 2012
- Industry Standard Server (ISS)
- Virtual appliance on VMware platform

The Delivery Controller provides pre-processing of the SMDR records. For nominal traffic rates, server specifications are listed in the table below.

Table 15: Mitel Delivery Controller server specifications

END POINTS	VCPU	CPU (GHZ)	RAM (GB)	DISK SPEED (RPM, IOPS)	DISK (GB)
Up to 15,000	2	5	8	10,000; 125	10
Up to 60,000	4	10	16	10,000; 140	100
Up to 120,000	8	20	16	15,000; 175-210	200

Specifications for use on Industry Standard Servers are similar, translating vCPU to cores with minimum speed 2.6 GHz each.

End-customer access is made using web-based browsers from the desktop or mobile device using iOS, Android™, or Windows® operating systems.

Vidyo

For Vidyo services, the following connections are required:

- MiCollab Server to VidyoPortal using the web services interface
- MiCollab Server to MiCollab Clients (using the MiCollab Client user interface for creating connections)
- MiCollab Client desktop and mobile clients to VidyoPortal using the web services interface
- Vidyo clients to VidyoRouter using SIP-based connections

For authentication services, MiCollab server and VidyoPortal must have a valid certificate to support the HTTPS connection. Remote clients can connect to the MiCollab Server using the MiVoice Border Gateway and to the Vidyo cloud servers using third-party firewalls and routers.

Redirect Server and MiCollab Client Deployment blade

The Redirect Server is provided as a Mitel cloud service and is publicly reachable over the Internet. The Redirect Server supports secure web services connections from MiCollab for Mobile clients and MiCollab Client Deployment components.

Networking requirements for the MiCollab Client Deployment component include connections to:

- Redirect Server to push the deployment URLs and related deployment credentials, using the web services
- MiCollab mobile clients to provide configuration data using web services
- MiCollab for management integration using either SAS Thrift™ interface for integrated mode, or MiCollab Client SOAP interface for co-located mode
- MiVoice Border Gateway to retrieve Teleworker parameters using web services (REST interface)
- Oria for management integration using Apache Thrift™

MiCollab Client Deployment must be publicly reachable from the Internet. A MiVoice Border Gateway may be used as a web proxy to the MiCollab server.

Networking and QoS

The choice of networking is largely determined by the service provider's available networking infrastructure and their commercial agreements, the service provider's MiCloud-based service offers, and the competitive pricing landscape for both services and networking.

This chapter addresses networking requirements to consider prior to deploying a UC solution. The networking requirements discussed here are relevant to service providers providing hosted solutions, data center operators, and installers and operators of any equipment located on customer premises.

- “Quality of Service, network assessment and end-point configuration” on page 105
- “Wi-Fi networks” on page 108
- “Wide Area Networks - QoS and Service Level Agreements” on page 109
- “Network infrastructure for IP phones” on page 109
- “Public network vs. MPLS” on page 112
- “Oria access to customer sites” on page 114
- “Bandwidth considerations” on page 115

Note: Much of the material discussed in this chapter is covered in greater detail in the *MiVoice Business Engineering Guidelines*.

Quality of Service, network assessment and end-point configuration

Network congestion due to high traffic levels can reduce the speed at which packets are transmitted, and cause delay in receipt of the packets. As network congestion increases, network switches and routers will often be forced to discard packets rather than forwarding the packets to the recipient. Packets are discarded when the network equipment cannot receive, process, and forward the incoming packets quickly enough. In some cases, packets may be discarded if they have been marked as low priority packets, and processing precedence is being given to higher priority packets.

When packets belonging to a real time voice or video stream are discarded by a switch or router, the audio quality or video quality experienced by the receiving end is reduced. When packets are discarded, a telephone call may break up to such a degree that the conversation becomes unintelligible, and participants in a video conference may receive an image stream that is delayed, and possibly heavily pixelated.

L2 and L3 QoS mechanisms are used to mark transmitted packets with a priority level. The priority level, or QoS setting, is used by network switches and routers to give packets with a higher QoS setting precedence over packets with a lower QoS setting or no QoS setting.

Network priority—QoS settings—are required to ensure the timely delivery of packets carrying data for real-time voice and real-time video over LANs and WANs. Use of the correct QoS settings on end-points, switches, and routers helps to ensure that users receive high quality voice and video services.

Network assessment

An assessment of the LAN should be conducted prior to installation of VoIP or IP video equipment. It is essential to assess the network, and if necessary, configure it to maintain good voice quality, video quality, and product functionality for users.

Depending on the results of the network assessment, the existing network may need to be modified, or equipment with QoS capabilities may need to be installed.

The network should be re-assessed any time there have been any major changes to the network design, or if there has been a significant increase in the number of users.

The main network issues affecting voice and video quality are delay, jitter, and packet loss.

Use the network limits shown in the following table to evaluate the results from a network assessment. For ideal voice and video packet transmission the LAN or WAN should comply with the values shown in the GO row.

Table 16: Network Limits

STATUS	PACKET LOSS	JITTER	END-TO-END DELAY	PING DELAY
GO	< 0.5%	< 20 ms	< 50 ms	< 100 ms
CAUTION	< 2%	< 60 ms	< 80 ms	< 160 ms
STOP	> 2%	> 60 ms	> 80 ms	> 160 ms

See the following guides for additional information about network assessment:

- *MiVoice Business Conference Phone/Video Phone Engineering Guidelines*
- *MiVoice Business Troubleshooting Guide*
- *MiVoice Business Voice Quality Troubleshooting Guide*
- *Voice Quality Solutions Guide*

L2 and L3 priority mechanisms

There are two areas where priority mechanisms can operate in the network to ensure that specific types of traffic is given higher priority than other types of traffic:

- Layer 2 in the LAN, through use of VLANs and packet tagging
- Layer 3 at network routers, and for WAN connections using Differentiated Services Code Point (DSCP) values

The following table shows the recommended L2 and L3 QoS settings for various traffic or service classes.

Table 17: Mitel-recommended Network QoS settings

SERVICE CLASS	RECOMMENDED L2 VALUE	RECOMMENDED L3 VALUE
Telephony (voice)	6	46 (EF)
Signaling	3	24 (CS3)
Multimedia conferencing	4	34 (AF41)
Real-time interactive	4	32 (CS4)
Standard	0	0 (DF) (BE)
Telephony (voice)	6	46 (EF)

For additional information related to:

• QoS settings for a particular application, end point or product	Refer to the appropriate product documentation found on Mitel OnLine.
• Troubleshooting voice quality, video quality and network QoS issues	Refer to the *Voice Quality Solutions Guide*, and the *MiVoice Business Voice Quality Troubleshooting Guide*, found on Mitel OnLine.
• Network QoS settings for third-party networking gear	Refer to the manufacturer's documentation.

Obtaining network parameters

Installing IP phones requires programming each phone's networking parameters, including QoS settings.

Mitel IP phones, consoles, and conference units have a number of different methods that they can use to obtain networking parameters such as VLAN and QoS information. Each network parameter source is assigned a priority level. An IP phone seeking network parameters starts with the priority level five method, which has the highest priority, If all of the necessary parameters are not available from this source, the phone uses each decreasing priority level until all the required parameters are found.

Table 18 lists the various sources of networking parameters and the priority level.

Table 18: Network parameters and priority level

SOURCE OF NETWORK PARAMETERS	PRIORITY LEVEL	NOTES
Manual entry (static)	5	Network parameters may be manually programmed by an installer through the phone set UI.
LLDP-MED	4	The IP phone's network parameters are obtained from an LLDP-MED-compliant L2 switch.

Table 18: Network parameters and priority level

SOURCE OF NETWORK PARAMETERS	PRIORITY LEVEL	NOTES
CDP Cisco Discovery Protocol	3	CDP can provide VLAN information to the IP phone and QoS values that are compliant with Cisco. Compatibility of the equipment can be inferred by the IP phone based on the fact that Cisco gear is present on the LAN.
DHCP	2	A DHCP server can provide the IP phone with network parameters
Factory default values	1	The IP phone contains factory default networking parameters

Teleworker Phones - obtaining a call server IP address

When an IP phone is first powered on in Teleworker mode, it attempts to find the IP address of the call server. In the case of a Teleworker phone, the call server is a MiVoice Border Gateway.

The Teleworker phone has three different sources that it can use to obtain the call server IP address. These sources, in descending order of precedence, are:

3	Manual (static) programming (via the IP phone's UI) Manual (static) programming (via the IP phone UI)
2	DHCP server
1	Mitel Redirection and Configuration Service server

Wi-Fi networks

When Wi-Fi networks are used to provide connectivity to IP phones and/or IP video devices, the wireless network must be evaluated carefully to ensure that it can support good audio and video quality.

The QoS standard for Wi-Fi networks is called WMM (Wi-Fi Multimedia).

Along with geographical coverage and bandwidth considerations, the Wi-Fi network evaluation must ensure that WMM is employed and that the WMM settings are correctly mapped to L2 or L3 QoS settings. The following table shows the recommended mapping of values.

Table 19: L2, L3, and WMM QoS mappings

SERVICE CLASS	L2 PRIORITY	L3 PRIORITY	WMM ACCESS CATEGORY	WMM CATEGORY
Telephony (voice)	6	46	AC_VO	Voice
Signaling	3	24	AC_BE	Best effort
Multimedia conferencing	4	34	AC_VI	Video

Table 19: L2, L3, and WMM QoS mappings

SERVICE CLASS	L2 PRIORITY	L3 PRIORITY	WMM ACCESS CATEGORY	WMM CATEGORY
Standard	0	0	AC_BK	Background

For additional information, refer to the Wi-Fi access point product documentation.

Wide Area Networks - QoS and Service Level Agreements

The UC solution uses Wide Area Network connections to:

- Connect remote office networks to the headquarters network
- Connect user IP devices located on the customer's premises to a private cloud or data center
- Connect user IP devices located on the customer's premises to a Hosted service provider

To ensure good voice and video quality, L3 QoS (DSCP) must be employed, and a Service Level Agreement (SLA) should be in place to ensure that the QoS requirements are honored by the WAN provider.

In some situations, there might be more than one network provider involved in establishing an end-to-end WAN connection between two particular locations. To ensure that QoS markings are honored end-to end across a WAN connection, it is imperative that the SLAs with all of the network providers involved be correctly set up with the same definitions, and that all BGP routers are configured according to the SLAs.

Sites requiring a high level of availability require both a primary and a secondary WAN link to provide a level of connection resiliency.

- For some customers, the same level of service is required on both of the WAN links. If this is the case, then both the primary and secondary links need to have similar characteristics and the same SLA must be in place for both connections.
- Customers not requiring the same level of service on both of the WAN links may be willing to accept a lower level of service in return for a cost savings by using a lower grade link as the secondary connection. In this situation, the customer may use an MPLS connection with an appropriate SLA as its primary WAN link, and the secondary link could be made over the public Internet where an SLA may or may not be available.

For VoIP engineering guidelines information, refer to the *MiVoice Business Engineering Guidelines*.

For information related to video conferencing over a WAN, refer to the *MiVoice Business Conference Phone/Video Phone Engineering Guidelines*.

Network infrastructure for IP phones

IP phones require basic networking infrastructure so that they may obtain firmware loads and networking parameters. Additional networking infrastructure is required to provide connectivity between the phones and the MiVoice Business controller. In some cases, network infrastructure may also be used to provide power to the phones over the LAN cabling.

For a list of the supported devices, see the *MiCloud Business Solution Blueprint*.

The following sections discuss the components that form the network infrastructure.

It is necessary to consider where the network infrastructure is physically located. For instance, some components, such as TFTP and DHCP servers, may be physically located on the customer premise or they may be located off-site in a cloud or data center.

When deciding where to locate the TFTP and DHCP servers, consider if there are requirements for local survivability. For example, if the WAN link is down and the servers are located off-site, the phones will not be able to communicate with the TFTP and DHCP servers.

Additional notes about the Small Medium Business - Low density architecture.

- There are two different connection paths to consider for the SMB-LD architecture, the internal path within the customer network, and the external path via the Teleworker gateway.
- In the SMB-LD architecture, the customer network is extended into the remote hosting site over a dedicated MPLS circuit. There are no gateways or breaks in the network. The MPLS connection is effectively transparent to the end-customer. As a result of this networking transparency, the customer can deploy any of the devices that are listed in the *MiCloud Business Blueprint*. There are some device restrictions for Teleworker connected devices.
- Some additional configuration is needed at the customer MiVoice Border Gateway Virtual to proxy through the MiCollab Client Multi-Tenant Service connections when using UC-specific soft phones, For Teleworker-connected devices, there are some restrictions on the devices that can be supported. See the *MiCloud Business Blueprint* for details.

TFTP server

MiVoice Business has an integral TFTP server which may be used to provide IP phones that are located on the same LAN with their application software. External TFTP servers may also be used, but to ensure that the phone software is at the correct revision for a given version of MiVoice Business software, it is recommended that the MiVoice Business integral TFTP server be used.

When remote or Teleworker phones are part of the UC solution, the MiVoice Border Gateway, which will have a copy of the current phone application software, will provide the Teleworker phones with their application software. If the phone application software on the MiVoice Border Gateway is out of date, the MiVoice Border Gateway obtains the latest phone software load from MiVoice Business.

For more information, refer to the MiVoice Business product documentation, the *MiVoice Business Engineering Guidelines* and the *MiVoice Border Gateway Engineering Guidelines*.

DHCP server

IP phones can use DHCP to obtain their networking parameters such as IP addresses, L2 priority settings, L3 priority settings, and VLAN information.

The following MiVoice Business products support integral DHCP servers:

- MiVoice Business on a 3300 ICP platform
- MiVoice Business Multi Instance

> **Note:** When MiVoice Business Multi Instance is used in hosted site deployments that support multiple customers, the integral DHCP server should not be used because a single DHCP server cannot be shared by multiple customers.In the case of servicing multiple customers, multiple third-party DHCP servers must be used.

- MiVoice Business for Industry Standard Servers (MSL-based DHCP server)
- MiVoice Business Virtual (MSL-based DHCP server)

For more information, refer to the MiVoice Business product documentation and the MiVoice Business System Administration Tool Help.

Layer 2 and Layer 3 networking equipment

To ensure good voice and video quality, L2 and L3 networking equipment must support QoS mechanisms, and the switches and routers should be configured as per the recommendations shown in the Network Assessment section.

To increase system availability, L2 and L3 redundancy and resiliency mechanisms, and networking protocols should be employed.

Some specific examples of redundancy mechanisms that can be used in network design are:

- Use of fully redundant L2 and L3 networking equipment
- Duplication of L2 switches and L3 networking devices
- Duplication of stored data, i.e. duplication of DHCP servers
- Duplication of storage devices, e.g. SAN, NAS, and RAID
- Duplication of transmission paths via partial mesh networking to support redundant communication paths
- Resilient topologies, i.e. hierarchical network design
- Networking protocols, for example Spanning Tree, Open Shortest Path First, VRRP, or Cisco HSRP

The *MiVoice Business Resiliency Engineering Guidelines* discuss network design for resiliency and L2 and L3 resiliency/redundancy protocols.

The *MiVoice Business Engineering Guidelines* discusses network design practices for network maintainability and scalability.

The Mitel white paper called Network Design for Availability provides a detailed discussion on network design practices for achieving higher availability.

Configuration information for specific models of switches and routers is covered in the product vendor's documentation.

Mitel's portfolio of IP phones are LLDP-MED compliant. If LLDP-MED compliant L2 switches are deployed the IP phones will be able to use the LLDP-MED protocol for obtaining networking parameters. LLDP-MED is also useful or providing phone physical location information for E911 purposes.

Power considerations

Consider the entire voice path from one device to another when distributing power. Consider especially which devices need to maintain power during a general power outage. Devices (such as phones) and the underlying network infrastructure continue to need power for phone service to be maintained.

The networking infrastructure requires UPS systems so that power can be maintained. Refer to the manufacturer's data sheets for each product's power requirements.

L2 switches that support PoE and are compliant with IEEE 802.2af and/or IEEE 802.2at should be considered for use as access layer switches to provide connectivity and power over Ethernet cabling to the IP phones.

PoE L2 switches allow the IP phones to be powered and managed from a common location. Additionally, a UPS system can be co-located with the L2 PoE switches so that L2 switch power and IP phone power can be maintained from a centralized location during a mains power outage.

Additional information can be found in the *MiVoice Business Engineering Guidelines* and the L2 switch and router product documentation.

Cabling infrastructure

LAN cabling should, at a minimum, be Category 5-compliant end-to-end. For information about LAN cabling refer to the *MiVoice Business Engineering Guidelines*.

For installations where the LAN cabling is Category 3 and there are circumstances that prevent the customer from upgrading the wiring plant, the StreamLine L2 switch can provide both power and connectivity to the IP phones. For additional information, refer to the StreamLine product documentation on Mitel OnLine.

Public network vs. MPLS

How can a carrier use MPLS to provide a network infrastructure that is both publicly-accessible while providing private connections for customers when required?

There are a number of ways that an MPLS network can be deployed, whether for private use only, for public use, or a mix of these requirements. Figure 14 highlights a typical configuration that deploys both private and public connections:

Figure 14: Network with Public and MPLS Connections

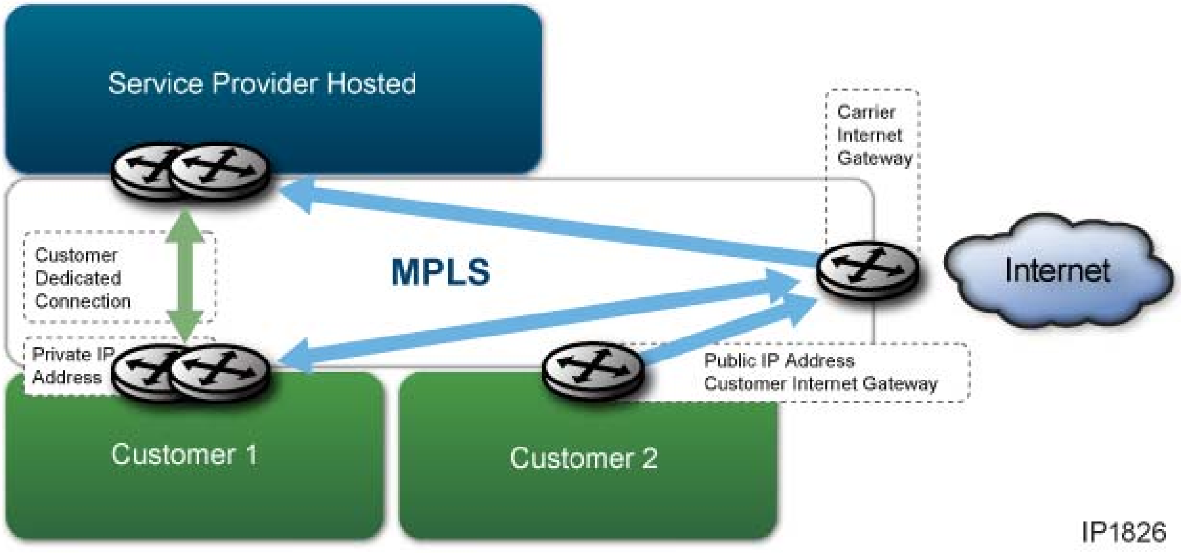

For a public MPLS connection, the end-customers and the service providers each have a router that has access to a unique public IP address. The IP addresses are provided by the carrier. Typically the carrier then directs any public IP traffic to the carrier's Internet gateway. Connections to public IP addresses in the carrier network are internally routed. Other public IP addresses are routed externally.

The advantage for the hosted service provider is that the public IP address can be reached from any network globally. Service Level Agreements can be applied to end-customers and service providers that share a common carrier/MPLS provider. One disadvantage is that all traffic must go through the common carrier Internet gateway. The carrier may be able to overcome this disadvantage by providing public routing within the network, but this also incurs management overhead of the devices and programming of routes, which the carrier may be less willing to do on a public connection.

For a private MPLS connection, the end-customer and hosted service provider share the same carrier network. The carrier network may use public IP addresses that are restricted from going out of the carrier Internet gateway by the router's access control list, or may opt to use private IP addresses that are not Internet routable. The routers at the customer and service provider are provisioned with dedicated routing rules and connections. In effect, the two end-connections and networks are provided with a transparent tunnel across the MPLS network. Advantages include security, plus the ability to provision Service Level Agreements with dedicated bandwidths, which may not be available via the carrier's Internet gateway.

MPLS carriers that wish to remain completely private do not provision an Internet Gateway. All connections are then private within the carrier space, and each connection has to be uniquely provisioned.

Oria access to customer sites

Figure 15, “Oria and NAT routing,” on page 115 shows the service provider network, the end-customer network and the NAT router that allows the two separate networks to communicate with each other.

The service provider's network uses private IP addresses in the 10.0.100.0/24 range, and the end-customer's LAN uses private IP addresses in the 192.0.101.0/24 range.

The Oria Management Application resides in the service provider's network and needs to be able to communicate with the applications that reside in the end-customer's LAN.

Since the service provider's LAN and the end-customer's LAN are using different IP address ranges, a 1:1 NAT router is needed to connect from the SP network into the customer network. Oria supplies the address required to allow the customer hosts to communicate with each other in the customer network.

During the initial network configuration, IP addressing relationships need to be established. Using Figure 15 as an example, the service provider and end-customer DNS machines would be programmed with the values in Table 20, “Service Provider and end-customer DNS entries,” on page 114 and the NAT Router would be programmed with the IP address mappings in the Table 21, “NAT Router configuration,” on page 114.

Table 20: Service Provider and end-customer DNS entries

DNS CONFIGURATION	FQDN	IP ADDRESS
Service provider	Customer_1_MiCollab	10.0.100.2
	Customer_1_MiVoice_Business_A	10.0.100.3
	Customer_1_MiVoice_Business_B	10.0.100.4
	Customer_1_MiVoice_Border_Gateway	10.0.100.5
End-customer	Customer_1_MiCollab	192.0.101.20
	Customer_1_MiVoice_Business_A	192.0.101.30
	Customer_1_MiVoice_Business_B	192.0.101.40
	Customer_1_MiVoice_Border_Gateway	192.0.101.50

Table 21: NAT Router configuration

END-CUSTOMER IP ADDRESS	SERVICE PROVIDER IP ADDRESS
192.0.101.20	10.0.100.2
192.0.101.30	10.0.100.3
192.0.101.40	10.0.100.4
192.0.101.50	10.0.100.5

Figure 15: Oria and NAT routing

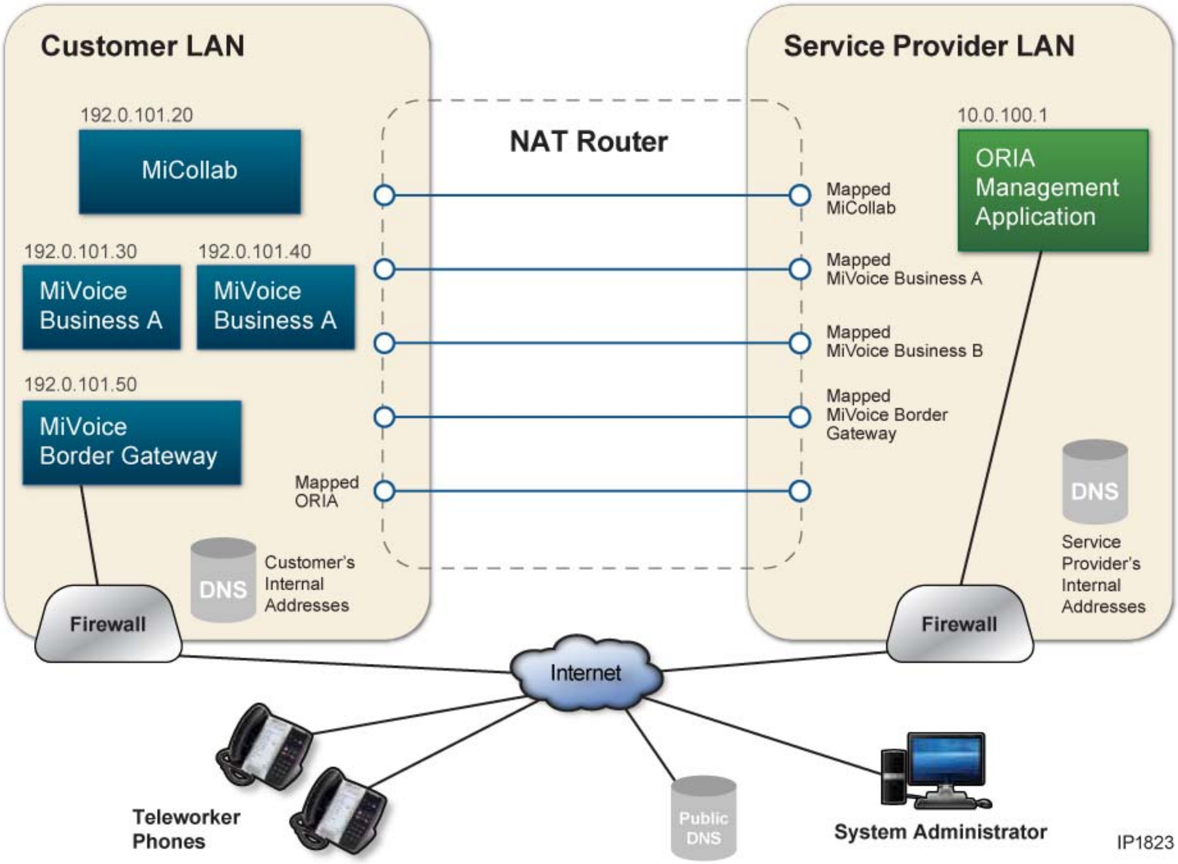

Figure 15 shows how Oria and the applications in the Customer's network communicate with each other. NAT routing can be done using MiCloud Management Gateway or VMware vCloud Networking and Security (vCNS).

Bandwidth considerations

Even when QoS mechanisms are employed in the network so that video and voice packets are handled with the requested priority by L2 switches and routers, it does not alleviate the requirement to verify that there is sufficient network bandwidth to carry all of the expected traffic.

QoS mechanisms are designed to ensure specific classes of traffic receive consistent treatment from networking equipment, but with insufficient bandwidth, QoS cannot guarantee performance for the high priority traffic such as video, nor does QoS ensure that low priority traffic will not be completely blocked by higher priority traffic such as video.

If a network interface does not provide enough bandwidth, or a network router is unable to process the volume of packets it is receiving quickly enough, the packets will be at risk of being corrupted, delayed, or completely lost, regardless of the QoS setting applied to the packets.

Bandwidth consumption

Before determining the bandwidth requirements for a particular communication link, it is important to consider the traffic flow and where devices are located relative to their controllers. The use of compression zones and IP networking may also have a bearing on traffic flow in parts of the network.

To determine what the total bandwidth consumption will be for a particular link, consider:

- Call traffic patterns
- Number of voice calls and whether they will be compressed
- Number of video calls and which CODEC will be used
- Number of video conferences and which CODEC will be used
- Data traffic, regular business traffic, and traffic patterns
- Maintenance traffic, file backup processes, and when they run

Information about how to calculate voice media bandwidth, the effect on bandwidth when using different CODECs and IP trunking are discussed in the *MiVoice Business Engineering Guidelines*.

Bandwidth consumption for video conferences using the MiVoice Video Phone is covered in the *MiVoice Conference and Video Phone Engineering Guidelines*. For information about the bandwidth required for MiCollab, see the *MiCollab Engineering Guidelines*.

The bandwidth required for third-party applications should be available in the vendor's documentation. If bandwidth utilization information is not available, network monitoring tools can be used to determine total bandwidth and peak bandwidth requirements. Network monitors can be run over a period of time to determine patterns.

Many routers provide embedded tools that can be used for measuring bandwidth consumption, and this is another option for determining bandwidth utilization.

Compression zones and bandwidth management

CODECs are devices or programs that encode or decode a signal into a digital format. The payload might be voice, video, or FAX data. Different CODECs can provide different sized payloads given the same input information. A reduction in payload is often referred to as compression.

IP phone calls that typically require compression are those for which the call traverses an IP trunk or a WAN connection with limited bandwidth. Compression zones are used to define where compression will be used. The decision about whether to use compression for a call is based on how compression zones have been configured on MiVoice Business, and the zone with which a particular phone is associated.

Establishing compression zones defines where compression will be used for a voice call, but there is nothing limiting the number of calls that can be placed across a connection. In fact, it is possible to oversubscribe a link by placing too many calls across the link. There is a potential for all calls on the link to suffer transmission impairments when a link is oversubscribed.

Use bandwidth management and Call Admission Control to monitor the usage of the communication links, and determine whether a call should proceed (or not proceed) across the communication link based on whether or not the link has reached its maximum capacity.

The terms “Bandwidth management” and “Call admission control” are often used interchangeably to describe the management, and potential re-routing, of calls across an IP network between end devices. These are actually two separate concepts:

- Bandwidth management gathers information about the availability and use of bandwidth on particular connections and links.
- Call Admission Control uses this information to decide whether a call should be completed or not.

Although the IP network is often considered as a “cloud,” it is actually made up of many parts, including LANs, MANs, and WANs. There are constraints on the amounts of data that can be handled at the transitions between the different networks, and often within the networks themselves.

If a link is bandwidth limited, it is desirable to be able to determine the available bandwidth across the link and manage it to ensure that it is available for voice use. After the bandwidth is known, it is possible to determine how many virtual channels can be established and to admit, redirect, or reject calls based on current available resources, that is, bandwidth. The latter is the task of Call Admission Control between end-nodes.

The *MiVoice Business Engineering Guidelines* and the MiVoice Business System Administration Tool Help provide detailed information about establishing compression zones, and how to use Bandwidth Management to prevent over-subscription on a connection.

Appendix A

GLOSSARY

Glossary

TERM	DESCRIPTION
ACD	Automatic Call Distribution. A package of advanced call processing features, relating to groups of agents who handle calls and agent supervisors.
AMC	Applications Management Center. Used to activate new hardware and software licenses for Mitel products.
ARID	Application Record ID - License ID created by and for the Mitel AMC.
ARP	ARP – Address Resolution Protocol. Used to identify a MAC address against an IP address.
ARS	ARS – Automatic Route Selection. This is a method whereby call control can best determine the path from one controller to another and provide a seamless connection to the user.
ASU	ASU – Analog Services Unit. This unit provides a combination of analog ONS interfaces for phones and/or LS trunks.
BRI	BRI – Basic Rate Interface. Digital ISDN connection to PSTN or local digital phone. This is the smallest quantity of digital channels that can be delivered, and consists of two digital channels for voice and data. Variants include the U interface for North America and S0 in Europe.
Call Control	Software to create connections and paths between end-user devices.
CAT 3	Category 3 Cable. A type of UTP cable for use in a LAN, capable of 16 Mbps. Typically used for voice and data on 10BASE-T Ethernet.
CAT 5	Category 5 Cable. A type of UTP cable for use in a LAN, capable of 100 Mbps.
CCS	Centum Call Second. A measure of call traffic. One call lasting 100 seconds is referred to as 1 CCS.
CDE	Customer Data Entry. A command line interface used to configure the Mitel 3300 ICP.
CDP	Cisco Discovery Protocol. A Cisco proprietary protocol that allows IP devices and L2 switches to communicate with each other for configuration purposes.
CEID	Cluster Element ID. A means of identifying different system units to maintain a consistent number plan.
CESID	Customer Emergency Services Identifier. A means of correlating a user and a directory number to information stored in a physical location data base.
CIM	Copper Interface Module. A TDM interface module used to connect the ICP to various peripherals via CAT 5 UTP.
CIR	Committed Information Rate. A means to identify how much information MUST be carried in a connection, e.g. CIR = 64 Kbps for voice.
CODEC	COder and DECoder. Coder and decoder commonly used as a single function. A means to convert analog speech into digital PCM and vice versa.
Controller	Control element of ICP (see also RTC).
COS	Class of Service. This refers to the priority value in the Layer 2 part of an IP packet when IEEE 802.1p is used.
CPH	Calls Per Hour. For example, six CPH means six calls per hour.

TERM	DESCRIPTION
CSM	Customer Service Manager. Former name for MiContact Center Office, an entry level contact center solution hosted on MiCollab for basic contact centers or workgroups with up to 100 agents.
CSMA/CD	Carrier Sense Multiple Access Collision Detect. The mechanism used on shared Ethernet connections to ensure that devices are not sending at the same time, and if they are, to initiate a back-off and retry algorithm.
CTI	Computer Telephony Integration. Means of combining computer functions to control operation of telephony equipment.
Datagram	A logical grouping of information sent as a network layer unit over a transmission medium without prior establishment of a virtual circuit. IP datagrams are the primary information units in the Internet. The terms “frame”, “message” and “packet” are also used to describe a datagram.
DECT	Digital Enhanced Cordless Telephony. Originally this was a European standard for digital cordless phones. This is now a worldwide standard, hence, the name change to Enhanced. Standard DECT phones are not available in North America.
DHCP	Dynamic Host Configuration Protocol. A means of passing out IP addresses in a controlled manner from a central point/server.
DiffServ	Differentiated Services. DiffServ is a protocol for specifying and controlling network traffic by class so that certain types of traffic get precedence. For example, voice traffic, which requires a relatively uninterrupted flow of data, might get precedence over other kinds of traffic. Differentiated Services is the most advanced method for managing traffic in WAN connections. This uses the Type of Service field at Layer 3 in an IP packet. See also DSCP.
DLM	Distributed License Manager
DN	Directory Number. A telephone or extension number.
DNIC	Digital Network Interface Circuit. A chip used as the basis for several sets which handle both voice and data.
DNS	Domain Name Server. A means of translating between typed names and actual IP addresses, e.g. microsoft.com = 207.46.134.222
DPNSS	Digital Private Network signaling System. A British common channel signaling protocol for requesting or providing services from/to another PBX.
DSCP	Differentiated Services Code Point. This is a value that is assigned to the Type of Service byte in each outgoing packet. The value can be in the range of 0 to 63 and allows routers at Layer 3 to direct the data to an appropriate queue. Value 46 is recommended for voice and will use the Expedited Forwarding queue or Class of Service.
DSP	Digital Signal Processor. This is a programmable device that can manipulate signals, such as audio, to generate and detect a range of signals, e.g. DTMF signaling.
DSU	Digital Service Unit. A peripheral which provides digital ports for the ICP.
DTMF	Dual Tone Multi-Frequency. In-voice-band tones used by telephones to signal a particular dialed digit. Also known as touch tone.
E	Erlang. A measure of usage of a resource, e.g. 0.75 e = 75%. 1 e = 36 CCS.
E1	Primary Rate running at 2.048 Mbps providing 30 channels of voice of PCM.
E2T	Ethernet to TDM. This is the conversion of voice streaming between TDM and IP.

TERM	DESCRIPTION
E911	Enhanced 911 (Emergency Services). Also 999 (UK) and 112 (International).
eMOH	Embedded Music On Hold.
ESM	Embedded System Management. Means to program a system from the System Administration Tool, Group Administration Tool, or Desktop Tool.
FAX	A means of transmitting printed text or picture information with acoustic tones.
FIM	Fiber Interface Module. A fibre optic TDM interface module used to connect the ICP to various peripherals.
FQDN	Fully Qualified Domain Name. The complete domain name for a specific computer, host, or IP end-point. The FQDN consists of two parts: the host name and the domain name. For example, an FQDN for a hypothetical server might be MyServer.Business.com. The host name is MyServer, and this host is located within the domain called Business.com. A Domain Name Server (DNS) is used to resolve the FQDN to an actual IP address.
FTP	File Transfer Protocol. An electronic method to transfer file information.
G.711	G.711 – PCM Voice Streaming. ITU standard for conversion of voice-streaming to digital PCM (64 kbps).
G.729	Voice Streaming CODEC. Reduced bit rate from G.711 (8 kbit/s)
GARID	Group Application Record ID (Group ARID)
Group Controller	The call control of the ICP is in control of a number of units, where the functions are more dedicated, e.g. to a separate gateway
GRP	Gateway Routing Protocol. A generic term which refers to routing protocols.
HSRP	Hot Standby Routing Protocol. A Cisco proprietary protocol used to increase availability of default gateways used by end hosts.
ICMP	Internet Control Message Protocol. Messages to help identify when devices are present and create warnings when they fail.
ICP	IP Communications Platform. Includes gateway function, call control, plus a number of other features, such as voice mail.
IP Address	Internet protocol address. A 32-bit address assigned to hosts using TCP/IP. An IP address belongs to one of five classes (A, B, C, D, or E) and is written as four octets separated by periods (dotted decimal format). Each address consists of a network number, an optional subnetwork number, and a host number. The network and subnetwork numbers together are used for routing, while the host number is used to address an individual host within the network or subnetwork.
IP	Internet Protocol. An encapsulation protocol that allows data to be passed from one end-user to another. Typically this was over the Internet, but the same protocol is now used within businesses.
IRDP	ICMP Router Discovery Protocol. An extension to the ICMP protocol that provides a method for hosts to discover routers and a method for routers to advertise their existence to hosts.
ISDN	Integrated Services Digital Network. The digital PSTN network. Integrated because this network carries both voice and data and provides direct digital connectivity to the user via BRI or PRI connections.
ISL	Inter-Switch Link. Cisco-proprietary protocol that maintains VLAN information as traffic flows between switches and routers.

TERM	DESCRIPTION
L2	Layer 2. The second layer of encapsulation of data to be transferred. Typically with TCP/IP this includes the MAC layer.
L3	Layer 3. The third layer of encapsulation of data to be transferred. Typically with TCP/IP this includes the IP address.
LAN	Local Area Network. This is a network within a local area, typically within a radius of 100 m. The transmission protocol is typically Ethernet II.
Leased IP	An IP address that has been assigned through DHCP and is valid only for the duration of the agreed lease time.
LLDP	Link Layer Discovery Protocol. A low level protocol used to pass information about the connection configuration between two end devices, for example VLAN. Typically this would be between an end device such as a PC or IP phone and the network access port on the Layer 2 switch.
LLDP-MED	Link Layer Discovery Protocol - Media End-point Discovery. LLDP-MED is an extension of LLDP that provides auto-configuration and exchange of media-related information such as voice VLAN and QoS. It is designed to provide enhanced VoIP deployment and management.
LS	Loop Start – This is a particular analog trunk protocol for signaling incoming and outgoing calls.
MAC	Media Access Controller. This is the hardware interface that data (media) travels through. Typically this will be assigned a world-wide unique address.
MAN	Metropolitan Area Network. This is a larger network that may connect a number of LANs within a business, as well as a number of businesses. Typically, this would cover a city area, and use fibre optics to get maximum bandwidth.
Mbps	MegaBits Per Second. Million bits per second is a measure of bandwidth on a telecommunications medium. May also be written as Mbits/s or Mb/s. Mbps is not to be confused with MBps (megabytes per second).
MDUG	Multi-Device User Group. A software construct used for grouping devices.
MDUL	Multi-Device User License. A licensing construct used for license sharing among devices.
MFRD	Mitel Feature Resources Dimensions. This is a definition of the number of features that can be used on a particular unit.
MHz	Megahertz. Frequency measurement.
MiNet	Mitel Network Protocol. This is Mitel’s proprietary stimulus-based protocol that is used to signal between phones and controllers, for example key and display information.
MiTAI	Mitel Telephony Application Interface. This Mitel implementation of TAPI is used to connect to external applications, e.g. ACD controllers.
Mitel OIG	Mitel Open Integration Gateway
MLB	MiCloud Business Solution Medium Large Business architecture
Modem	MOdulator-DEModulator. Device that converts digital and analog signals. At the source, a modem converts digital signals to a form suitable for transmission over analog communication facilities. At the destination, the analog signals are returned to their digital form. Modems allow data to be transmitted over voice-grade telephone lines.
MOH	Music on Hold

TERM	DESCRIPTION
MTBF	Mean Time Between Failures. The statistical time between expected component failures.
MTU	Maximum Transmission Unit. An MTU is the largest size packet or frame, specified in octets (eight-bit bytes), that can be sent in a packet- or frame-based network, such as the Internet.
MWI	Message Waiting Indicator. A visual indicator in a telephone that indicates to the user that a message is waiting.
NAT	Network Address Translation. A means of translating internal IP addresses to a defined limited range of Internet IP addresses. The benefit is the ability to use a limited range of Internet addresses and map these to a much larger internal range.
NIC	Network Interface Card. Physical connection to the network. In a PC, this is often a plug-in card.
NSU	Network Services Unit. This interface connects between the PSTN Primary Rate trunks and the ICP.
On-Premise	UC Enterprise Solution Solution On-Premise architecture
ONS	On-Premise Line. This is a two-wire analog telephony interface, within an office environment, and not passed outside.
OPS	Off-Premise Line. This is a two-wire analog telephony interface, typically installed external to a building, e.g. external shed or guard house.
OSPF	Open Shortest Path First. A link-state routing protocol used for routing IP traffic over the most cost-efficient route.
PC	Personal Computer
PCM	Pulse Code Modulation. The digital representation of analog signals.
PDA	Personal Digital Assistant. A handheld personal organizer that can interface to a PC or a Mitel PDA Phone.
Permanent IP	An IP address that has been leased (from DHCP) on a permanent basis.
PI	Performance Index. A calculation of the performance limits of a system. Different weighting values are assigned to various types of calls. Based on the expected calls per hour (CPH) of all of the user ports on the system, a system performance index (PI) can be calculated. The system PI is used as an indication of how much traffic the 3300 ICP can handle at any one time.
Ping	This is a means of sending a test message and waiting for a reply to determine if a network device is reachable. On a PC, this is invoked with the command ping.
PPM	Parts Per Million. This is a measurement of accuracy, or the expected error in one million events. Therefore one PPM means that 999,999 to 1,000,0001 events occurred when 1,000,000 were expected. This is 0.0001% error. For example, a household clock that is one second accurate per day is 11.5 ppm, or would need to be 0.086 seconds incorrect per day to be one ppm.
PRI	Primary Rate Interface. This is a connection to the PSTN where a number of trunk channels are multiplexed onto a common connection. Both T1 and E1 variants are available.
Private Cloud	UC Enterprise Solution Private Cloud architecture
Private Cloud, Shared Services	UC Enterprise Solution Solution Private Cloud, Shared Services architecture

TERM	DESCRIPTION
PSTN	Public Switched Telephone Network. The telephone network that provides local and long distance connections, e.g. Bell, AT&T, British Telecom.
PTT	Poste Telefonie Telegrafie. PSTN services. Some countries combine postal services and telephony under a common service provider, e.g. the government.
RAD	Recorded Announcement Device.
RAID	Redundant Array of Independent Disks. Array of hard drives on which the information is duplicated. A controller manages the disks, switching automatically from the primary to the secondary in the event of the failure of the primary hard drive.
RDN	Remote Directory Number. The Remote DN Table is used to identify alternate ICPs to check for availability of devices, and to determine if a device is located on the Primary or Secondary ICP.
RFC	Request For Comments. A document that is created, maintained and distributed by the Internet Engineering Task Force. An RFC is the vehicle that is used to discuss and evolve a networking related protocol. RFCs usually get approved and issued as standards.
RFP	Radio Fixed Parts. The Radio Fixed Parts (RFPs) connect to the 3300 ICP through the LAN. The wireless phones communicate with the RFPs using standard Digital Enhanced Cordless Telecommunications (DECT) protocol.
RGP	Router Gateway Protocol. A means whereby routers on a common subnet can communicate with and identify each other. Useful when ICMP Re-direct is needed to identify an alternative path.
RIP	Routing Information Protocol. A networking protocol that maintains a database of network hosts and routers and exchanges information about the architecture of the network.
RSTP	Rapid Spanning Tree Protocol. A version of STP that will converge networks more rapidly than STP (see STP).
RTC	Real Time Complex. This is the control block within an ICP. This includes Call Control and internal controls for the unit.
RTP	Real Time Protocol. Protocol used to identify sequence of voice packets with timing information before being sent to a user via UDP.
SAC	Switch Application Communications
SB architecture	MiCloud Business Solution Small Business architecture
Scalable	MiCloud Business Solution Scalable architecture
SET	System Engineering Tool. Used for calculating system parameters, limits and allowable additions.
SIP	Session Initiation Protocol. An IETF standard for signaling over IP.
SMB	MiCloud Business Solution Small Medium Business architecture
SMB-LD	MiCloud Business Solution Small Medium Business, Low Density UCC architecture
Static IP	An IP address that has been manually assigned and fixed. Typically, static addresses are exceptions within DHCP.
STP	Spanning Tree Protocol. A means whereby the network can determine multiple paths between two points and disconnect them to leave a single path, removing broadcast issues.

TERM	DESCRIPTION
Subnet	A subnet (short for “subnetwork”) is an identifiably separate part of an organization's network. Typically, a subnet may represent all the machines at one geographic location, in one building, or on the same local area network (LAN).
SWB	Mitel Sales WorkBench
T.37	Internet Protocol for FAX (Store and Forward). A means of taking a TDM FAX, converting it to data, passing it via IP and reconverting it back to TDM.
T.38	Internet Protocol for FAX (Real Time). Similar to T.37 in function, but carried out in real time, i.e. with minimum delay.
T1	Primary Rate. Provides 23 or 24 channels of trunks per connection.
TAPI	Telephony Applications Programming Interface. TAPI is a standard programming interface that lets you and your computer communicate over telephones or video phones to end-users or phone-connected resources.
TAR	Tape Archive and Retrieval. A file transfer utility.
TCP	Transmission Control Protocol. The methods of transmitting data between two end-points using IP with acknowledgement.
TDM	Time Division Multiplex. A means of combining a number of digitally encoded data or voice channels onto a common digital stream, e.g. T1.
TFTP	Trivial File Transfer Protocol. A simplified version of FTP used to transfer data with minimal overhead.
TOS	Type of Service. A field within the Layer 3 (IP) encapsulation layer to identify some properties relating to service parameters; in this case, delay and priority of handling.
UDP	User Datagram Protocol. A layer 4 protocol with minimal handshaking and overhead. Used to stream voice. Considered connection-less.
ULM	Unified License Manager
Unicast	A process of transmitting messages from one source to one destination, as opposed to a broadcast or multicast.
UPS	Uninterruptible Power Supply. A unit capable of providing output power for a period of time when the local mains supply fails. Usually relies on storage devices such as batteries.
UTP	Unshielded Twisted Pair. Cable that reduces emissions and maintains an impedance match through the twists per metre in the cable without resorting to shielding.
VLAN	Virtual LAN. A means of providing virtual LANs on a network using common physical components. Such VLANs are logically unconnected except through some Layer 3 device.
WAN	Wide Area Network. A network connection to a network that could be global, e.g. via Frame Relay.
Wi-Fi	Wi-Fi Alliance technology for Wireless LAN based on IEEE 802.11.
WLAN	Wireless LAN
WAV	WAVe file. WAV is an audio file format, created by Microsoft, that has become a standard PC audio file format for everything from system and game sounds to CD-quality audio. A Wave file is identified by a file name extension of .wav.

Mitel
Powering connections

mitel.com