MiCloud Flex

Solution and Engineering Guidelines

January 2023

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.

No part of this document can be reproduced or transmitted in any form or by any means - electronic or mechanical

- for any purpose without written permission from Mitel Networks Corporation.

TRADEMARKS

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

Contents

Revision History

Document Version	Document Release Date	Updates
7.0	January 2023	• Updates for MIR Release 7.1
6.0	December 2022	• Further WebRTC Pro clarifications • Removal of EML Speech Analysis • Inclusion of High Traffic Resource settings for MiCollab • Inclusion of SMBC for resilient survivable gateway
5.0	August 2022	• MBG WebRTC Pro section mentions limitation to support WebRTC Anonymous Mode.
4.0	April 2022	• MIR on-premises storage • Inclusion of WFM into MiCC-B Large • Web-Proxy updates for MICAM, MiCC-Outbound and MIR • Support for up to 7500 users on MiCollab • Inclusion of Azure Cognitive Services for Speech Analysis
3.0	December 2021	• MiCollab Advanced Messaging • MiCC music in queue requirements • Solution Diagram • MIR deployment and disk partition diagram update
2.1	October, 2021	• MiContact Center Speech – Deployment architecture • WFM as a standalone virtual machine
2.0	August, 2021	• MBG Cluster and its limits. • Additional information on deployment disk partitions (MIR and Outbound). • More details about MiCC Outbound. • Web-proxy for MiCC Outbound. • Nuance/MiCC Speech definitions and architecture • SD-WAN and MiCloud Edge
1.0	May, 2021	• Initial release

Solution Overview

MiCloud Flex is a secure, hosted, scalable Unified Communications and Collaboration (UCC) solution with a focus on high availability and dependability, allowing customers to run a wide range of Unified Communications and Contact Center applications.

The MiCloud Flex solution is comprised of different Mitel Unified Communications reference designs and topologies to meet service provider and customer requirements, network connectivity requirements, and system Unified Communications scaling. The solution covers scaling from a few users up to and beyond 10,000 users, with each. user being associated with multiple devices (for example, office desk phone, softphone, mobile phone, etc.).

Unlike Unified Communications as a Service (UCaaS) offers that are typically multi-tenanted solutions, MiCloud Flex is a hybrid solution utilizing both multi-instance and multi-tenant components. It provides the capability to scale and customize the deployment to meet customer requirements, rather than a one size fits all methodology. MiCloud Flex enables both Over the Top (Internet) and private networking connectivity.

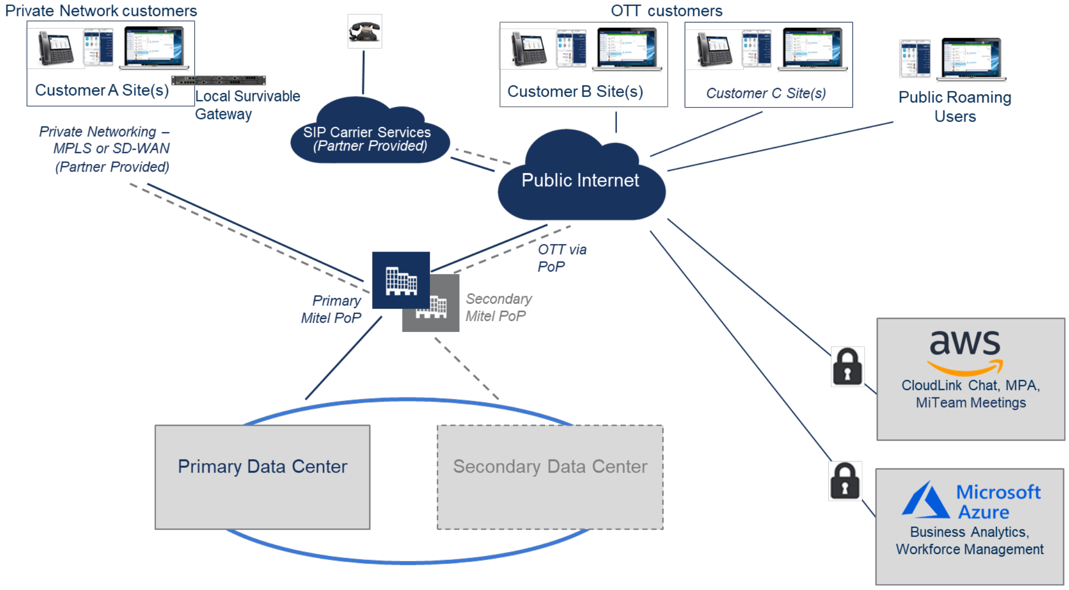

The core telephony, collaboration and some optional services operate in VMware hosted (in Tier3+) data centers. Each customer instance of these core services is deployed in customer dedicated virtual machines that are installed on their own isolated Virtual LANs (VLANs). This means that each customer’s instance is isolated from all the other customer instances within the data center infrastructure. A customer cannot contact another customer without routing through the Public Switched Telephone Network (PSTN), as customers are isolated from each other.

Shared Responsibility Model

In the MiCloud Flex solution, there are a number of key functional areas that are linked together to provide a service to the end-customer which creates the need for some demarcation but also a shared responsibility model. Each of these areas have different levels of responsibility and access, which may be best described in the following diagram:

End User / Customer	**Customer**
Customer Configurationand Support	**Partner**
Resource Definition, Licenses	**Mitel**
Infrastructure	**Hosting Provider/Mitel**

The infrastructure is provided by the hosting provider, in this Mitel. Mitel will provide the Public IP addresses and provide access to virtual or physical Points of Presence at defined Cloud Exchange, or colocation facilities. It is the Partner’s responsibility to provide any connections to that infrastructure from a private hosted network, for example, MPLS, or other connections via a separate Connection Provider.

Mitel will provide the definitions of the resources required on that hosted infrastructure, either in this document, or via Mitel Doc Center, based on the partner’s input for the customer requirements. Mitel will provide any appropriate infrastructure or non-Mitel operating system licenses, specifically any Windows OS licenses and Windows SQL licenses. Mitel will provide an Anti-Virus solution.

First and second line customer support, installation and configuration are the responsibility of the partner. Mitel Professional Services may be employed to provide any necessary assistance

End user configuration of features, and any Moves/Adds/Changes is the responsibility of the customer, or system administrator, who may be customer of partner based.

Backups and upgrades are the responsibility of the partner.

For further details see the *MiCloud Flex Partner Guide* on Mitel PowerUP.

Deployment Overview

This section provides an overview of the MiCloud Flex deployment in a hosted infrastructure. Each functional box may consist of one or multiple virtual machines, based on requirements and scaling. These are running in customer specific and isolated networks, even when hosted by a public Cloud provider. Cloud Services, including CloudLink and others such as MPA, MBA and WFM, are considered as Internet connected even though the clients are running in the customer network.

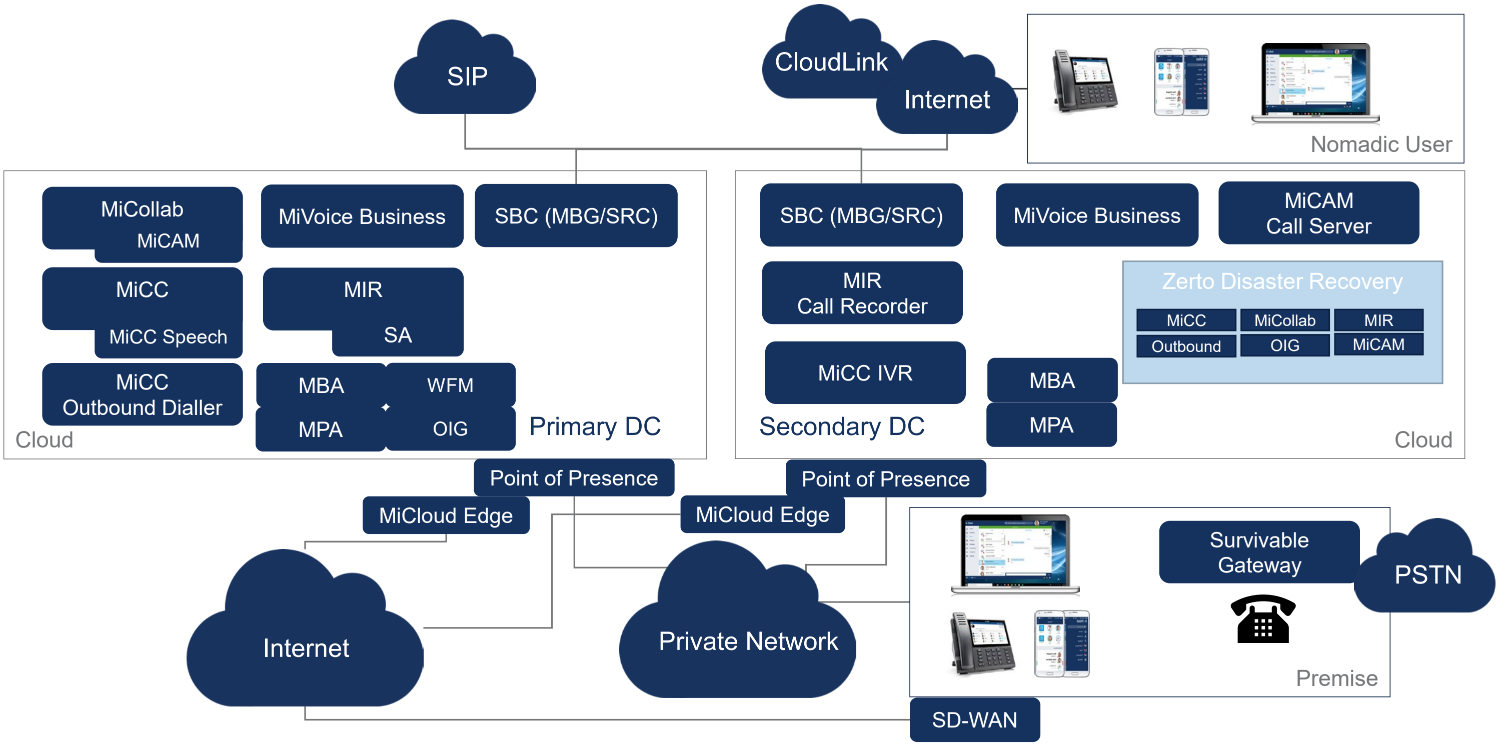

The deployment can be configured and deployed in the following scenarios:

- Single Data Center in a non-resilient mode. Consists of the Cloud services plus the Primary DC.
- Dual Data Center with voice resiliency. Consists of all components in the drawing. Non-voice components are not deployed as resilient and may require custom configuration and deployment work to do so.
- Dual Data Center with voice resiliency and Disaster Recovery as a Service (DRaaS using Zerto)

Non-voice components are not deployed as resilient However, they are protected with a disaster recovery service. Disaster recovery includes key block level information replication between the primary and secondary data centers and will allow the non-resilient components in the primary data center to be instantiated in the secondary data center as a replacement when connected to the replicated data, in the case of a longer-term outage or disaster.

Not all deployment configurations may be available in all regions. Refer to the *General Information Guide* for further details on specific areas and capabilities.

Information on device and application compatibility can also be found in the *General Information Guide* and the Product Compatibility Matrix.

This document covers the engineering guidelines for MiCloud Flex -deployments and includes the following products as part of the solution offering.

Core Applications – Multi-Instance

The core applications, telephony, voice and collaboration tools are deployed as customer dedicated solutions within the VMware infrastructure. The Core components include:

- MiVoice Business (MiVB). Mitel IP-PBX/Call-Control
- MiCollab: Mitel's Unified communication and collaboration solution which includes:
- MiCollab Audio, Web and Video (AWV) Collaboration
- MiCollab Presence engine with calendar integration
- MiCollab corporate instant messaging
- MiCollab (NuPoint) Unified Messaging that also includes a feature rich Auto Attendant
- Optional MiCAM Advanced Messaging
- MiVoice Border Gateway (MiVBG, aka MBG); an Internet facing edge device that acts as a Session Border Controller (SBC) for SIP trunks and phones, as well as for Mitel Proprietary IP Phones and applications.

Optional Applications – Multi-Instance

The following optional applications are deployed as multi-instance solutions in the data centre:

- MiContact Center Business – A feature rich omnichannel contact center solution including an optional speech recognition module
  - Optional Workforce Scheduling
  - Optional Mitel Workforce Management (WFM)
  - Optional MiCC Speech, Mitel Advanced Speech Recognition (ASR) and Text To Speech (TTS) integration with MiContact Center Business Interactive Voice Recognition (IVR) units.
- MiContact Center Outbound – comprehensive, integrated outbound campaign management suite that supports all modes of outbound dialing
- Open Integration Gateway – an Application Programing Interface (API) platform with pre-existingSFDC and Google integrations.
- Mitel Workforce Optimization suite which includes
  - Mitel Interaction Recording, (Call and Screen Recording)
  - Mitel Quality Management, and

  - Mitel Speech Analytics using Azure Cognitive Services
- MiCollab Advanced Messaging (MiCAM) with optional ASR and TTS
- MiCloud Edge SD-WAN network connectivity
- Revolution (available on request)

Optional Applications – Multi-Tenant

The following optional applications are deployed as multi-tenant solutions:

- MiTeam Meetings – A cloud-based team instant messaging tool hosted on Amazon Web Services(AWS) utilized by MiCollab
- MiCollab Chat and MiContact Center Business Web Chat are both delivered using underlying technology known as CloudLink that enables a hosted multi-tenant chat capability and is hosted byAmazon Web Services (AWS).
- Mitel Business Analytics (MBA) – A suite of cloud-based Business Analytics tools hosted on Microsoft’s Azure platform. Connection is established via a secure connection over the Internet from locally deployed collector/client units.
- Mitel Workforce Management (WFM). A suite of cloud-based contact center focused analytic tools hosted on Microsoft Azure platform. Connection is established via a secure connection over the Internet from locally deployed collector/client units.
- Mitel Performance Analytics (MPA)* for system performance analysis and recording, including speech quality and alarming. MPA also provides a secure tunnel for remote management and is the preferred management connection solution. Hosted on Amazon Web Services public cloud platform.

Note*: Although MPA is identified as an optional component, this is the main management access portal to the solution when other alternatives such as dedicated network access, or Virtual Private Network (VPN) connections are not already provided.

The topologies include different virtualization technologies, and in many cases a hybrid of technologies, to take advantage of different features. The applications and call server platforms rangefrom on-premises servers and appliances to Mitel-optimized virtual call **servers** and virtualized application packages running on VMware combined with Public Cloud.

Connectivity Options

MiCloud Flex supports multiple connectivity options, including:

Over the Top (Internet)

Secure, encrypted Over the Top (OTT) services are supported across the Internet without the need for a VPN with MiCloud Flex. For increased service availability Mitel utilizes multiple Internet Service Providers for OTT services into the data centers. Note that not all services are available OTT. This uses public IP addresses.

Private On-Net

This is a private network, or extension of a customer private network and private addresses into the hosted data centre. Connection options include:

MPLS

Mitel supports the use of customer's supplying MPLS for creating a private connection(s) to the data centers with a Service Level Agreement (SLA) from their chosen carrier provider. from the customer site with MiCloud Flex.

SD-WAN

Mitel supports the use of customer's supplying their own SD-WAN (BYOD) for creating a private connection(s)to the data centers from the customer site with MiCloud Flex. The SD-WAN vendor must have been through an interoperability test with Mitel to ensure compatibility.

Mitel's MiCloud Edge SD-WAN solution is also available.

Feature Comparison (Private and Public Networks)

The following information provides a comparison of feature availability between private network and public network connections.

A private network provides the ability to route information within a network domain using Layer2 or Layer3 protocols and does NOT include any Network Address Translations (NAT) to obscure the underlying connection information. Networks in disparate locations may be interconnected to form a private network and may deploy point to point solutions on the WAN connection, using technologies such as SD-WAN, MPLS and VPN. IP addresses used are typically those for private address, as defined in RFC1918. VPN and SD-WAN may use public addresses to connect between gateways in a public network, but the information in the internal payload is not altered and IP addresses remain unchanged.

A public network connection, also referred to as Over the Top (OTT), provides the ability to route information between different network edges and gateways. It will employ IPv4 network address translation (NAT) between private addresses and public addresses and may involve changes to the internal data payload. Data will transition through a gateway, and have address translation, and possibly application-level translation as well. An example would be a Teleworker phone on a home network connecting to a central office gateway. Another example would be connection to a SIP Trunk provider SBC gateway. The payload is transmitted across the public network in native form, and NOT hidden within a tunnel protocol. As such, it should also employ higher levels of security and encryption on the connections to prevent snooping of the contents or masquerading of the connection by a third party. Connections via public networks are often over an undefined path between multiple network providers. A public connection within a network provider is still referred to as OTT, as it's still unclear to the user at either end of the connection what path is taken even within this defined network.

Devices

MiCloud Flex Feature	Public OTT	Private On-Net	Notes
Survivable Gateway	No	Yes	On-premises remote and survivable (secondary controller) gateways are availablewhen connected via a private on-net connection. It is not available OTT
IP Phones (53xx Series)	Yes	Yes	5304, 5312, 5320, 5324, 5320e, 5330e, 5340e, 5360 can support FQDN use. FQDN support is included in Phone Firmware load 6.5.0.128 and higher. MiVB 9.1 SP1 includes this firmware load. Earlier versions of 5300 phone, including 5302, 5330, 5340 (non-'e' version) and those with only 10/100Mbps Ethernet connections cannot be upgraded to support FQDN. It is recommended that new installation use FQDN only along with phones that support this. Older installations may include phones that do not support FQDN
IP Phones (69xx Series)	Yes*	Yes	*Multicasting to 6900 phones does not work via MiVBG, for example, Revolution. XMLPush Notifications are also not available via MiVBG
MiVB Console	Yes	Yes	
MiCollab Softphone	Yes	Yes	Desktop client (SIP), Mobile Client (SIP) (iOS/Android) or Web Client (WebRTC) (Additional solution and product limits may apply to WebRTCconnections)
112 DECT and RFP12	Yes	Yes	Restrictions on number of FQDN characters may limit deployments (limited to63 characters)
IP DECT	No	Yes	Only Onsite IP-DECT Base-station is supported. It is not possible to host the physical base-station in the Data Centre. This needs to be deployed on the customer premises.
SIP DECT 600	Yes	Yes	OMM (physical base-station (RFP) or Linux) must be deployed at the customer premises. It cannot be hosted in the Data Centre. Local OMM management is required. It cannot be hosted in the data-centre. M-OMM is not supported. The OMM emulates SIP phones, one per DECT device. Each emulated SIP phone can register as a teleworker for OTT deployments.
WebRTC and WebRTC Pro Clients	Yes	Yes	WebRTC clients are always treated as Teleworker devices, even those on the customer private network. Internet access is required. Enabled ‘WebRTC Pro’ on the MBG to enable Call Recording, participation on contact centre agent features and resiliency. This capability is only available with the Chrome browser. WebRTC Pro settings on MBG may be incompatible with existing functions. Check under MiCollab and MBG for further details.

ATA TA71xx	Yes	Yes	ATA device sourced through Partner
5485 IP Pager Unit	No	Yes	On-Net connectivity only

Solution Features

MiCloud Flex Feature	Public OTT	Private On-Net	Notes
On-Site CRM Integration	No	Yes	A secure on-net path to the customer CRM is required
External CRM Integration	Yes	Yes	External CRM integration is offered through use of OIG.
Active Directory Integration / Support	Yes	Yes	Office365 AD (Azure) and onsite AD server are supported. Connections to an on-premises AD system via an OTT connection is not available.
Email Integration	Yes	Yes	NuPoint Messenger (NPM) as part of MiCollab supports authenticated SMTP forwarding and may be used with public e-mail services. Embedded Voice Mail (MIVB) supports native SMTP authentication and may be used with public e-mail services (requires MiVB Release 9.2, or higher).
Hospitality features: • PMS integration • SMDR logs • Hotel logs	No	Yes	Only available to On-Net connections.

Interconnection

MiCloud Flex Feature	Public OTT	Private On-Net	Notes
Hybrid Mode – Multiple sites, each with their own system. Ability to connect cloud to premises systems	No	Yes	On-premises MIVB and applicationssupported
Syslog	No	Yes	Available on MSL based products:MiVBG, MiCollab, MiVB
MiCloud Edge SD-WAN	Yes	Yes	Extends private customer network to cloud services over public Internet

Contact Center

MiCloud Flex Feature	Public OTT	PrivateOn-Net	Notes
Email Queue as part of Multimedia	Yes*†	Yes*	* Customer must enable/configure IMAP and SMTP access to their mail server. Mitel recommend that TLS/SSL is enabled for IMAP and SMTP connections, including use of 'STARTTLS' † Connection to an OTT e-mail server may be limited and may require deployment of an in-line authentication E-mail Forwarder
Chat Queue as part of Multimedia	Yes	Yes	* Chat is CloudLink based and uses OTT connections, and requires Internet access
SMS Queue as part of Multimedia	Yes*	Yes*	* The customer will need to sign up to Twilio for the SMS service. The implementation requires Mitel Professional Services
Contact Center Client	Yes	Yes	
Ignite Web Client	Yes	Yes	
IVR Customer Onsite Database Lookup	No	Yes	A secure on-net path to the customer DB is required.This does not affect SFDC.
IVR Customer Web/ Integration Onsite Server	Yes	Yes	If the Web/ Integration server is onsite there must be path for OTTservices to reach the server. Web proxy is available via MiVoice Border Gateway
Workforce Scheduling (WFS)	No	Yes	On-premises connections only. Requires connection to local SQL server.

Call and Screen Recording (WFO MIR)

MiCloud Flex Feature	Public OTT	PrivateOn-Net	Notes
Call Recording IP Stations Internal	N/A	Yes	Requires On-Net SRC. Note that MiCollab softphones (mobile phone or PC) should register as Teleworker (external devices), even if on the internal office network
Call Recording IP Stations External	Yes	N/A	Requires SRC at external MiVBG. Consider scaling for parallel recording (2 x 50%overhead)
Call Recording Trunks	Yes	Yes	Requires ON-Net SRC, typically part of the SIP gateway. Consider scaling for parallel recording (2 x 50% overhead)
Screen Recording	No	Yes*	* This requires a private connection and enough bandwidth to support the number of recorded screens to stream to the Cloud data centre Virtual desktop screen recording is not supported.
Manual Recording on Demand (Stop / Start Recording)	Yes*	Yes	Via Web service. Screen recording only available via On-Net connection.
Automatic Recording on Demand (Stop / Start Recording)	No	Yes	This capability requires use of screen-scan, or screen-mute, licenses and use of the thick client application and therefore only available via On-Net connection.
Playback (audio) web and scoring (PowerPlay Web)	Yes*	Yes	Scoring is provided via a license option: "WFO MIR Quality Monitoring".
Powerplay (audio) thick client (PowerPlay Pro)	No	Yes	Requires dedicated connections and limited to On-Net connections
Playback (Video)	No	Yes	Only available for On-Net connections
MiVoice Business Console	Yes	Yes	Only available with Trunk Side Recording
Reports/Dashboard	Yes	Yes	
XML Phone control	No	No	Contact Mitel Account Team

Speech Analysis (WFO SA) - Transcription and Keyword Spotting

MiCloud Flex Feature	Public OTT	Private On-Net	Notes
Speech Analysis Web	No	Yes	On-net connections only

MiContact Center Outbound

MiCloud Flex Feature	Public OTT	Private On-Net	Notes
Agent Interface	Yes	Yes	Two options: Standalone with call scripting Integrated with MiCC-B Web Ignite OTT connections via MiVoice Border Gateway web proxy
Dashboard (was Live Monitor)	Yes	Yes	Limited to Administrators
Call Monitor (was NVP Monitor)	Yes	Yes	
Route Manager (was Strategy Manager)	Yes	Yes	
NVP Monitor	Yes	Yes	
Inbound Wallboard	Yes	Yes	
Interaction Studio	No	Yes	Downloadable application that runs locally on workstation. Limited to Administrators and Developers.
Inbound Voice Management	Yes	Yes	Limited to Administrators
Campaign Management	Yes	Yes	Limited to Administrators
Predictive Dialing	No	Yes	Options Available include: 1) Integrated with MiVB over MiTai. In this mode a number of IP5020 devices licenses equal to twice the number of concurrent agent licenses are required. These are non-trusted devices. (Suitable for Legacy and On-Premises deployments) 2) Integrated with MiVB over SIP Trunks. In this mode a number of SIP Trunk licenses (zero cost) are required equal to three times the number of concurrent agent licenses. Advanced Dialer licenses are also required. Additional SIP Trunk licenses needed for predictive dialing (Suitable for hosted and new deployments) Option 2) is the preferred and recommended solution for Flex.
Dialer and Telephony SDK	No	Yes	Limited to Developers

MiCAM

MiCloud Flex Feature	PublicOTT	PrivateOn-Net	Notes
MiCAM Web Server	Yes*	Yes	Includes user login for mailbox access, notifications and reports. Some admin functions may not be available OTT
E-Mail Server	* (See note)	* (See note)	Available from On-Net clients. Only available externally where this service is already Internet accessible, such as with Outlook 365 and Gmail.

Revolution

MiCloud Flex Feature	Public OTT	PrivateOn-Net	Notes
Multicast Notifications	No	Yes	Multicast and XML Notifications do not work via MiVoice Border Gateway, nor across the Internet through NAT gateways and devices. This is ON-Net only
Live Streaming	No	Yes	Limited to On-Net connections
All	No	Yes	On-Net private connections only. The Mitel Paging Relay unicast to multicast functionality can be used to assist with multicast paging across WAN links, where required.

Availability and Resiliency

This chapter discuss the availability and resiliency mechanisms as they relate to the variouscomponents, applications, and services that are used to build the overall UC solution.

The following Availability Configurations are supported with Flex:

- Single Data Center - Non-Resilient
- Dual Data Center - Voice Service Resilient
- Disaster Recovery to secondary Data Center - Voice Service Resilient

Single Data Centre HA Resiliency

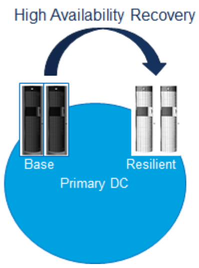

Dual Data Centre Resiliency

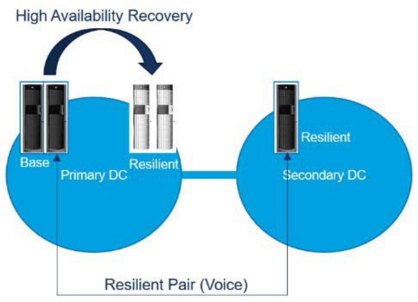

Dual Data Centre, Zerto Disaster Recovery

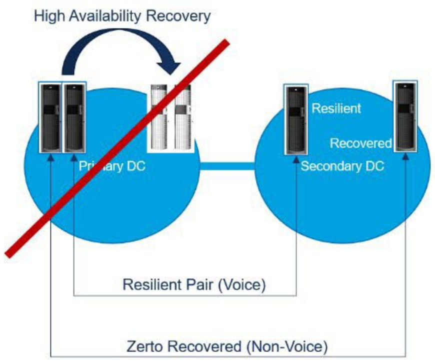

Certain applications and functions that are not user or voice critical, or business critical, may not be provided in a resilient configuration, for example some reporting may not be resilient. Voice services and items affecting voice services are configured as resilient, when necessary.

MiCloud Flex takes advantage of the VMware High Availability (HA) capability to provide improved service and application availability, even for single data centre deployments. See the section below on

VMware High Availability for deployment considerations for HA capabilities. Use of HA for non- resilient functions can quickly recover such applications and can closely mimic resilient levels of operation.

Even with VMware HA capability, improved availability is offered through use of dual data centers, in different regions, and redundant components. This is the default configuration for voice critical services.

When higher levels of service availability are required across all applications, these can be included ina Disaster Recovery as a Service (DRaaS) configuration. For MiCloud Flex, DRaaS is provided in a VMware environment through integration with the Zerto application. Further details can be found in the Dual Data Centre with Disaster Recovery section in Deployment and Platform Considerations section, below.

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

Flex deployments are based on a VMware infrastructure and follow the reservation settings as outlined in the *Virtual Appliance Deployment Solutions Guide*, available in the Mitel Doc Center.

VMware High Availability

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. VMware High Availability is a protection mechanism that is used to rapidly recover a virtual application during hardware or operating system failures.

All applications in MiCloud Flex are protected through VMware High Availability (HA). Where a Mitel virtual application has no native application resiliency capabilities, VMware HA is still used to provide a higher level of system availability, within the local data centre.

For those Mitel applications that offer native resiliency capabilities, VMware High Availability works in conjunction with the application's resiliency capabilities to further increase system availability. With VMware HA, the protected servers are geographically co-located within the same data centre.

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 applications are 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.

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

Application Considerations

MiVoice Business

The overall design of the call handling solution with the MiVoice Business must be designed to meet the customer requirements, which may be a simple office deployment, or may involve multiple offices and remote locations. This may involve a base pair of resilient MiVoice Business call controllers or may involve a larger deployment of multiple call control units. A hierarchical design simplifies management and call-routing configurations. A network that offers multiple call connection paths is inherently resilient but must also be programmed as Alternative Routes.

All MiVoice Business units must be clustered together to simplify call routing between primary and secondary controllers. All units within a cluster are aware of other controller status and location of end devices and can route calls accordingly. MiVoice Business units are also connected with IP-Trunks and programming of Alternative Router Selection (ARS).

SIP Trunking

When external SIP Trunks are used, the MiVoice Border Gateway provides the necessary Session Border Gateway (SBC) gateway functionality for the MiVoice Business. The MiVoice Border Gateways can be deployed in pairs for redundant operation of SIP Trunks.

Primary and secondary SIP trunk registration (from MiVoice Business Gateway) with primary and secondary Service Provider SBCs can be used to a common service provider to provide 100% fail-over in the event of a lost SIP trunking gateway or connection.

Only SIP Service providers that have been tested and recognized from the Mitel SIP Center of Excellence (CoE) shall be used. New service providers can obtain access to the approved list through the MSA programme.

Use of third-party SBCs as part of the customer deployment are not considered as these do not provide the inherent Application-Level Gateway (ALG) functions required with the Mitel and other third-party applications. Mitel solutions are designed and tested to work together,

PSTN Trunking

Direct PSTN connectivity can be provided through use of 3300 Appliance, or EX controller, units, when deployed as survivable gateways. Alternatively, SIP ATAs may also be used for minimal analogue connectivity.

Connection Provider Network Considerations

The connection provider will need to connect any private network infrastructure of the customer to the hosted service provider. To minimize the impact of any outage, and improve network availability, the following recommendations are suggested:

- Provide multiple connection per geographic location from the partner network to the hosted provider
- Provide multiple geographic access points to the hosted provider
- Consider use with a connection provider to increase geographic connection points, and also improve interconnection cost efficiency
- Ensure that the hosted provider includes a (backbone) connection between primary and secondary customer deployments in different geographically separated data centre, or different regions.

Customer Network Considerations

If a high availability UC solution is required, it is critical that the customer network also be designed to be resilient. This can include the following network architectures:

- Design the network in a hierarchical fashion, such as core, distribution, and access layers.
- Use of networking equipment and protocols that can be cross connected to full redundancy.
- Use of multiple DHCP servers, and use of DHCP forwarding between subnets

Application Overview

Mitel IP Desk Phones

The MiVoice Business solution allows individual IP phones to be configured and licensed forresilient operation.

The System Administration needs to define which phones require resiliency and which MiVoice Business controller pairs will be used to provide this capability. MiVoice Business units that provide primary/secondary pairs must be clustered together with IP-Trunk connectivity.

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 primaryMiVoice Business.

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

MiVoice Business Console

The MiVoice Business Console is a PC based application and Mitel IP softphone. This also supports resilient operation in the same manner as a Mitel IP Phone.

MiVoice Business

The MiVoice Business is deployed as a software solution when deployed in data centres in an appropriate hosted virtual environment. This includes the operating system and application, and only requires access to a virtual machine. The MiVoice Business is also available as a hardware appliance product such as 3300 controller, an EX controller or SMB Controller (SMBC), when deployed as a location based survivable gateway.

The MiVoice Business call control resiliency solution provides the ability to preserve telephony servicein the unlikely 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 appliance platforms.

For additional information related to resilient operation, see the *MiVoice Business Engineering Guidelines*, the *MiVoice Business Resiliency Engineering Guidelines* and the *Virtual ApplianceDeployment Solutions Guide*.

Resilient Call Routing

All MiVoice Business units in a cluster are aware of a phone's primary and secondary controller, and also the status of controllers within that cluster network. If a phone is hosted on a secondary controller because the primary is inaccessible, the controllers in the cluster recognize where to the route the call.

All systems in a cluster must be connected by IP trunks, so that calls can be routed to the correct controller. Alternative routes should also be programmed to assist in directing calls in an efficient manner. For larger systems, configure the MiVoice Business systems in a hierarchical manner to easy the management and programming of these routes rather than a total mesh network.

Where applications are geo-resilient between data-centers, the customer, or partner, must configure the appropriate IP routing across the hosted provider backbone for applications to communicate successfully between the data-centers. Alternatively, this connection may be made over the connection provider connectivity to the dual data-centers, although this is the least preferred option as it offers a lower level of availability.

Resilient Call Control Features

While in service on a secondary MiVoice Business, an IP phone retains call services. 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 a Survivable Gateway

The MiVoice Business may be deployed as a Survivable Gateway where the customer requires the ability to access the local PSTN from the customer premises. It can also provide secondary controller functions for local IP Phones, in the event that the connection to the hosted provider is unavailable.

A local office may also require connectivity to SIP trunks. In this event, an additional MiVoice Border Gateway will also be needed, as well as connectivity to a SIP SP that can be available in the event that the outage is not network related.

Typically, this requirement is met with the dedicated hardware appliances, see following section.

MiVoice Business Hardware Platforms

The MiVoice Business call control is also available on different hardware platforms, for on-premises deployments, including 3300 appliances, EX Controllers, SMB Controllers and Industry Standard Servers running Mitel Standard Linux (MSL). These hardware platforms will be deployed on customer premises, not within the hosted provider network, and most often used to provide local survivable gateway and PSTN connectivity. Refer to the *MiVoice Business Engineering Guidelines* for further platform details.

Mitel 3300 Appliances, SMBC and EX Controller

The 3300 appliances, SMB Controllers (SMBC) and EX Controllers are Mitel proprietary hardware platforms used for running the MiVoice Business software. These are used where legacy TDM connectivity is required, or where a customer wants a physical unit on-premises. The Mitel MiVoice Business appliances can also provide the following functionality:

- IP to PSTN gateway (analog and digital/PRI) (PRI not available on SMB Controller)
- Analog phone connectivity
- Embedded Voice mail and Auto-Attendant
- Local Music on Hold
- Local Ad-hoc conferencing
- The SMB Controller also offers Teleworker and SIP Trunk Service Provider connections via the integral MBG

These hardware units may be used as local survivable gateways in the event that the cloud solution is unavailable.

Industry Standard Servers (ISS) and Virtual Servers

MiVoice Business can be installed on Industry Standard Servers (ISS). ISS platforms offer higher call processing performance and faster fault recovery times than the 3300 appliances.

Deployments with Industry Standard Servers are IP connected, in the same manner as the hosted solution. Where there are legacy TDM requirements, an external gateway such as a Mitel EX controller or an Analog Terminal Adapter (ATA) must be used.

Connections to a SIP Trunk Service Provider will require an additional MBG, deployed locally for local survivability.

ISS deployed MiVoice Business can be used as local survivable nodes in the event that the cloud solution is unavailable.

MiVoice Border Gateway

MiVoice Border Gateway is available in variants:

- Standalone as a virtual machine in a hosted environment
- Standalone, on Industry Standard Server, on-premises with a local survivable gateway
- With MiCollab as a virtual machine in a hosted environment (typically for smaller deployments).
- Included with SMBC

The MiVoice Border Gateway ensures the deployment of secure internal and external workspaces, enabling remote workers seamless access to the voice capabilities of the office. It provides a gateway between the internal private network of the customer and the public network of the Internet. The MiVoice Border Gateway provided Application-Level Gateway functions for Mitel IP phones and SIP phones to call control applications, and also as a gateway to SIP Trunk Service Providers.

The MiVoice Border Gateway can provide web-proxy capability to a limited number of predefined applications within the solution. Different deployment options are available, depending upon the quantity of connections required. See MBG Engineering Guidelines for further details and also the Application sections below.

MiVoice Border Gateway units are clustered on both the internal and external interfaces. This ensures correct load-balancing and hand-offs in the situation where MiVBGs are operational but unreachable due to network issues. Units are also clustered for license sharing and load balancing for devices that support the Mitel IP Phone MiNet protocol. For SIP device connectivity and SIP-Trunk connectivity Border Gateways are typically clustered as primary/secondary pairs.

MiVoice Border Gateways that are deployed with a survivable gateway solution with MiVoice Business need to be clustered with the primary units in the hosted solution. Note that clustering limits may apply. Please contact your Mitel Account Team if the number in a cluster exceeds the limits in the *MBG Engineering Guidelines*.

Mitel IP Phones (MiNet based) maintain multiple IP addresses to multiple MiVoice Border Gateways. If an IP Phone loses connection to the MiVoice Border Gateway, and on failing to re-establish the connection, it attempts to connect to the next IP address in the list. When MiNet sets are configured for persistent resiliency lists, the resiliency list is retained through a power cycle.

MiVoice Border Gateway also supports a number of SIP devices. Resiliency for these phones is provided through DNS and lookup to multiple gateways.

MiVoice Border Gateway supports WebRTC connectivity. There is a selection for different WebRTC functionality, which is mutually exclusive. These settings are **WebRTC** and **WebRTC Pro**. The **WebRTC** setting supports basic level of WebRTC across multiple browsers. It is primarily used for anonymous user connections and MiCollab AWV. It provides basic call connectivity only. The **WebRTC Pro** setting is supported on Google Chrome only, but offers full softphone capability, including call recording and a number of Contact Center features. If both WebRTC and WebRTC Pro functions are required, then multiple MBGs will need to be deployed.

MiVoice Border Gateway supports Web Proxy connections to specific internal applications allowing access over the internet or public networks. There are limits to the number of connections available. See MBG Engineering Guidelines for further details.

The MiVoice Border Gateway can also provide connectivity to external SIP Trunk service providers. Typically, user gateways and trunk gateways are deployed separately as they have different resiliency mechanisms. For non-resilient, and smaller scale deployments, it is possible to deploy a common MiVoice Border Gateway for both user and trunk gateway functions, as long as a common IP address can be used for both gateway functions. In the case where a SIP Service Provider uses a private IP address, it may still be necessary to split the gateway into two separate functions.

The VMware High Availability (HA) solution is used to provide additional application-availability. 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 a MiVoice Border Gateway within the application suite. For smaller non-resilient UC solutions, this MiVoice Border Gateway can provide both the user and trunk gateways. However, for most deployments, the recommended configuration is to provide dedicated and standalone MiVoice Border Gateway units. In this situation, the MiVoice Border Gateway, within MiCollab, is still connected to the cluster, but with a primary function to share licenses and to allow configuration of the MiVoice Business Gateways from a central location in MiCollab. In this case, the MiCollab MiVoice Border Gateway is configured with a load sharing capability of 0%.

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

A MiCollab PC soft phone may be deployed on-net, or on the LAN, when this device is contained within the building. The softphone will register directly with the appropriate MiVoice Business. However, when the PC softphone is deployed on a more mobile device, such as a laptop, which may be more nomadic and out of the office, this needs to be deployed as a Teleworker device, even when located in the office. This is also a requirement if this softphone is part of a call recording solution.

The mobile device soft phone (Mobile phone, or PC) should always connect to the external MiVoice Border Gateway as a Teleworker, even if it is currently on the internal business network.

MiCollab Client

The MiCollab server must be operational for MiCollab Clients to be able to make calls in normal operation. MiCollab Client can be deployed as a dedicated application (thick client) or increasingly as a web browser (thin client). It can be used in both cases as a softphone. For Web-browser support, use of Google Chrome is recommended along with MBG Setting for “WebRTC Pro”. Be aware that this setting may impact MiCollab AWV, if being used, and deployment of multiple MBGs may be required.

In the event that the server is unavailable, the following options are available:

- The softphone client will not be able to search the directory and place calls from the directory. However, the phone will still maintain telephone connections and an ability to place calls manually.
- The PC/MAC without softphone relies on making calls on behalf of another phone. This servicewill not be available, and calls will have to be placed manually on the associated devices.
- Mobile phones will not be able to search the directory and make calls. However, the phone will still maintain telephone connections and the ability to place calls manually.
- Web browser based PC softphones will lose service as the browser is reliant on the server being available.

The MiCollab client can connect to MiCollab, when the device/client is on the internal network but needs to connect via the MiVoice Border Gateway when on an external network. This requires the use of split-DNS to ensure internal and externally connected devices connect to the MiCollab directly, or via the MiVoice Border Gateway, respectively. This connection requires a web-proxy connection via the MiVoice Border Gateway, along with any softphone connection on a remote PC.

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 asa core component of MiCollab. MiCollab Audio, Web, and Video Conferencing provides no native application resiliency; it relies on MiCollab service availability.

Split-DNS is needed to differentiate connections on-net direct to the MiCollab server from external connections via the MiVocie Border Gateway proxy. The external connection is also needed for guests to join the meetings.

Public Internet connection for AWV requires a separate public IP to be applied to this service. This is a separate IP address from that used by MiVoice Border Gateway and other MiCollab proxy connections.

MiCollab UM 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 from their PCs or telephones. MiCollab Unified Messaging also allows inbound callers to quickly find the person they need to talk with using a speech-enabled auto attendant and call routing functionality.

The MiCollab Unified Messaging application connects to both primary and secondary MiVoice Business controllers to provide resilient access and operation.

The MiCollab Unified Messaging application does not itself natively support resiliency of the application and follows any resilient operation of the MiCollab suite of applications, such as use of VMware HA and optional Disaster Recovery as a Service (DRaaS).

- MiVoice Business System Administration Tool Help
- Virtual Appliance Deployment Solutions Guide

MiCollab Advanced Messaging (MiCAM)

MiCollab Advanced Messaging is a next generation voice application offering Unified Messaging, Transcription, Speech-enabled Directory and Automated Attendant, Secure Voicemail and more.

MiCollab Advanced Messaging/Unified Messaging is Windows based Unified Messaging service. Users manage their messages from a workstation via a web-browser, e-mail application or via telephone calls. Options for storage location and e-mail integration are offered. Integration with public Cloud e-mail services is supported.

The solution consists of a core System Server, which manages the integration and user access, plus a number of outlying Call Servers to handle the incoming calls and user phone access. Messaging, Advances Speech Recognition and Text to Speech are options available on the Call Servers.

The central System Server is not resilient. However, the Call Servers can be deployed in a resilient manner and will continue to handle calls in a resilient operation mode. The system will synchronize data on recovery. Access to existing call recordings and playback may be limited during resilient operation. The level of restriction is dependent on the storage location of the recorded data, the access methods offered, and the level of integration with Cloud based e-mail services.

Mitel Open Integration Gateway (OIG)

Mitel Open Integration Gateway (OIG) 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. However, Mitel OIG supports inter-operation with MiVoice Business systems that are configured for resilient operation.

Mitel OIG application availability is provided by VMware HA and DRaaS.

MiContact Center Business

The MiContact Center Business application, with the exception of the IVR function, does not support resilient operation. MiContact Center Business uses VM-ware HA and optional DRaaS.

Resilient Contact Center IVR operation is available using multiple IVR servers, which can be deployed in both primary and secondary data centers. For smaller deployments, the MiContact Center Business server may host one of the IVR instances. However, the recommendation is to deploy multiple separate IVR servers. Each IVR instance is configured with active and redundant ports. In the case of one of the IVR servers being out of service, the redundant ports of the backup IVR carry the full traffic load.

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.

The IVR servers use a local data cache with a mirror of the main server configuration to continue operationand routing calls, if the MiContact Center Business server is unavailable.

MiContact Center also integrates with local Work Force Scheduling (WFS), or with the optional WorkForce Management application (WFM).

Web connections to the MiContact Center Business for agents are available for on-net connections and can be made available via the MiVoice Border Gateway web-proxy capability.

See the *MiContact Center Business Blue Print Guide*, on Mitel Doc Center, for further details on scaling, resilient and hierarchical contact centre deployments.

MiContact Center Speech (Advanced Speech Recognition and Text To Speech)

MiCC Speech encompasses three main features that integrate with the MiCC Business IVR units. They provide the following function:

- Advanced Speech Recognition (ASR)
- Standard Test To Speech (TTS) for message playback
- Enhanced Text To Speech (TTS) for message playback

Advanced Speech Recognition (ASR) may also be identified as “IVR Speech Recognizer”.

A number of voices (languages that may be regional, for example, UK English versus NA English), can be provided with Standard Text to Speech capability to more natural Enhanced Text To Speech. The improved resulting voice of Enhanced TTS sounds more natural, but also requires a separate and more powerful server to handle the processing.

The Advanced Speech Recognition (ASR) and Standard Text To Speech are two applications which may be provided on a common server. The Enhanced Text To Speech required a separate dedicated server.

The features are only associated with the IVR servers of MiCC-B at the primary location. Resilient operation will fall back to standard IVR operations at the secondary location without ASR and TTS advanced capability.

The ASR/TTS servers should be included with the Zerto grouping for disaster recovery along with MiCC-B.

MiContact Center Outbound

The MiCC Outbound application is an optional add on to the MiCC Business application. This provides a number of dialing campaigns ranging from dedicated outbound agents to using time of available agents and trunks to maintain call handling efficiency. As such, the scaling of the solution is tied into the scaling of the underlying MiCC.Business.

The call handling capacity of the SIP Trunk Provider should be determined prior to deploying this capability, as this system has the capacity to generate many unanswered calls. Consequently, agreement with the SIP Trunk provider regarding their usage policy must be in place prior to going live, otherwise trunk service for the whole solution may be impacted.

The MiCC Outbound application is not considered voice critical and is not delivered with a resilient solution. It is deployed as a single data center deployment.

Agents connect to the MiCC Outbound applications via a web interface. This is available directly for on-net agents, and via the MiVoice Border Gateway web proxy for remote agents.

Increased availability is offered through use of VMware HA and for increased service via the DRaaS.

Mitel Interaction Recording (MIR)

The solution can be deployed in a number of different resilient configurations. These will be covered in more details with the different deployment topologies, but essentially cover:

- Single Data Center - Non-Resilient
- Dual Data Center - Voice Service Resilient
- Disaster Recovery to Secondary Data Center - Voice Service Resilient

Screen recording is not resilient by default. Consult your Mitel Account Team if this is a requirement.

A single data center deployment relies on VMware HA for improved availability. This requires that any storage that needs to be recovered must be partitioned off-board the main server, through use of SAN. Separate disk partitions for application, database and storage are recommended. For larger deployments, and off-board storage via the Archive NAS server is recommended rather than a local disk partition.

For the Dual Data Center to provide resilient operation, there must be an existing connection between the networks within the two data centers. This allows recordings to be taken in parallel and remain synchronized with the core server for de-duplication at storage. The deployment uses parallelrecording to ensure that if a recording server fails, that the recording is not lost. This means that the primary and secondary recording servers are running active/active and each of these must have an active license both on the server and for the recording stream. See under the MIR License section forfurther details.

For the DRaaS service to work, there must be an existing network connection between the primary and secondary locations. The database and storage must be duplicated at both locations, so that on transition of the core server and database to the secondary location, the service can continue with minimal disruption. Instigation of Disaster Recovery involves a level of manual change-over and is not instantaneous.

MIR includes functionality that can control the stop/start recording functionality of the call. Ifthis is required, then additional Web servers are deployed to ensure resiliency between two data centers. The DNS needs to reference both webservers to enable the web-browser to locate the web servers.

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.

Agents, or users, connect to the Mitel Interaction Recording web server for access to playback recordings and to stop/start recordings in flight (Recording on Demand). Access to the web-server is available directly for on-net connections and also via the MiVoice Border Gateway we-proxy for remote agents, or users.

For more detailed information refer to *MIR (Mitel Interaction Recording)* in the section onApplications.

Speech Analysis (SA) and Azure Cognitive Services

Speech Analysis is an optional application integrated with Mitel Interaction Recording. For new installations Speech Analysis (Transcription) is now provided through integration with Azure cloud hosted Cognitive Services, Keyword spotting is carried out post-transcription.

The Azure Cognitive Services provides a speech to text transcription service for recorded conversations. Recordings are still maintained in the call recording storage location. Recorded files are transferred to Azure Cognitive services to be transcribed into plain text which is then added to the recording files. Multiple languages are supported. Key word spotting is provided by an advanced search capability within Mitel Interaction Recording.

Access to the web service is provided via the Mitel Interaction Recording web server.

MiCloud Business Analytics (MBA)

Mitel Business Analytics (MBA) is a Cloud-based service and relies on a local collection client, per customer, to collect information from the MiVoice Business SMDR information. For MiCloud Flex, the MBA collection is centralized through a multi-tenant service, provided by Mitel in the hosted management plane. Customers log into this cloud service for their individual call reports.

The collection service is not inherently resilient, although multiple clients may be deployed, for example, primary and secondary controller collection. If a client were to fail, the SMDR information collection will resume on client recovery.

The Mitel Business Analytics client, or Data Collector, is deployed as a virtual machine. Use of High Availability configuration is recommended, to maintain any locally collected data that has not be transferred to the cloud.

Mitel Performance Analytics (MPA)

The MPA client, or Probe, handles multiple connections to multiple devices and is scaled by additional units. The application is provided as a Cloud service, hosted in Amazon Web Services (AWS). Connection to the Internet is required to provide this functionality.

MPA is the recommended management access mechanism. Connection via a secured tunnel is managed from the Cloud service.

The MPA client, or MPA probe, is available as a virtual machine, where it may be deployed locally with a group of MiVoice Business controllers and associated applications.

For resilient operation it is recommended to deploy an additional MPA unit in the secondary location. To cover the situation of network outages for remote office locations, additional MPA clients can also be considered for inclusion at these locations as well.

Use of HA on the client virtual machines is recommended.

Workforce Management (WFM)

The Workforce Management (WFM) is provided as a Cloud service, hosted on Azure. This requires Internet connectivity from the customer network. The WFM Client integrates with the MiCC Business.

Initially deployed as a separate virtual machine, this service is now available integrated into the MiContact Center Large Server configuration. New deployments will use this integrated option.

This WFM client is associated with the MiCC Business server, which is not inherently resilient.

Use of High Availability on the client machine is recommended. The WFM client and MiCC should also be included with the other similar applications in the DRaaS configuration when Disaster Recovery is a requirement.

Revolution

Mitel Revolution is Mitel's next generation mass notification solution that was built for today's modern organization to improve network-wide communications, security, and emergency responsiveness across the organization.

With Mitel Revolution, organizations can facilitate enterprise-wide communication for virtually any communications need - including real-time and automated notification alerts for emergencies, large scale notification for routine (or non-critical) communications, facility-wide live overhead paging announcements, scheduling of bells/prerecorded announcements, and mobile-centric communications for mobile employees, students, or others who registered using its self-service portal.

If deploying this service, please contact your Mitel Account Team for deployment options.

MiCloud Edge SD-WAN

MiCloud Edge provides a means for customers to extend their private network to the cloud solution, using the Internet as the carrier network. This is an alternative to using private network providers, such as MPLS.

This requires the deployment of a SD-WAN router (or multiple for active/passive high availability) on the customer premises and private network. Mitel Cloud Edge will terminate the connection from these on-premises routers to customer dedicated public IP addresses within the MiCloud Edge Infrastructure. The termination will then apply the appropriate networking isolation, such as use of VLANs, to ensure customer isolation within the hosted service cloud.

Use of a resilient hosted solution requires that a network connection exists between the primary and secondary data-centers for correct operation and hand-off between the applications and gateways. This can be provided by a dedicated inter data-centre connection, or over a private network, such as MPLS. SD-WAN, and MiCloud Edge only provide connections to the data centres, and not between them. Therefore MiCloud Edge, and SD-WAN, are currently limited to deployments, for Flex, in UK and North America geographic regions.

Please contact your Mitel Account Team if you wish to deploy this capability, as it does require physical devices to be deployed on the customer network. Additional information can be found under MiCloud Edge and under Flex on Mitel Infochannel.

Third Party monitoring tools

It is possible that customers or partners may include network monitoring tools as part of the solution. The Mitel recommended implementation uses MPA, but customers, or partners, may decide to add in further monitoring tools. Be aware that network monitoring uses resources from the applications, and a high frequency of monitoring uses a higher level of resources. Mitel applications include real-time voice streaming which is delay sensitive. Unnecessary monitoring of system metrics may consume and divert resources with subsequent impacts to end the user voice quality experience. Use monitoring tools with due diligence and focus on those areas that are important, rather than a default measurement of every available metric.

Mitel products will continue to operate with reasonable passive metric collection, including external simple network probing and SNMP information gathering. Use of more aggressive and more intrusive 3rd party probing tools, including use of PowerShell, will consume system resources with a resulting performance impact. Aggressive or intrusive monitoring may result in voice quality degradation or may raise unnecessary alarms.

Should Customer and Product Support teams become involved in a deployment, Mitel may request that the customer temporarily reduce or remove active monitoring devices and services in order to isolate specific problems.

Anti-Virus and Intrusion Detection System software

It is recommended to protect deployments with anti-virus and/or intrusion detection software. Note that inclusion of this software may require system resources over and above those defined for the application and operating system. Consider these requirements as well, when deploying a virtual machine. Additional resources may be required for CPU, RAM and HDD/SSD storage. Consider any AV/IDS application settings, such as files and folders not to include in the check status. Some functions of Anti-virus may also impact performance of the system, especially with streaming. Consider when some of the background functions can be run, such as out of normal office hours.

Network Infrastructure

The Network Infrastructure covers many aspects of connectivity from the customer network to the hosted provider, to the Internet and to the PSTN. It also covers configuration aspects of the customer network and infrastructure, as well as the configuration and routing of information across a number of external networks. The configuration of the applications is also important, as these also play a part in ensuring that data is routed correctly.

The following diagram provides an overview of the different networks that can potentially be involved in making sure that the customer routes data to the correct location and also that the hosted service is correctly delivered to the end-customer.

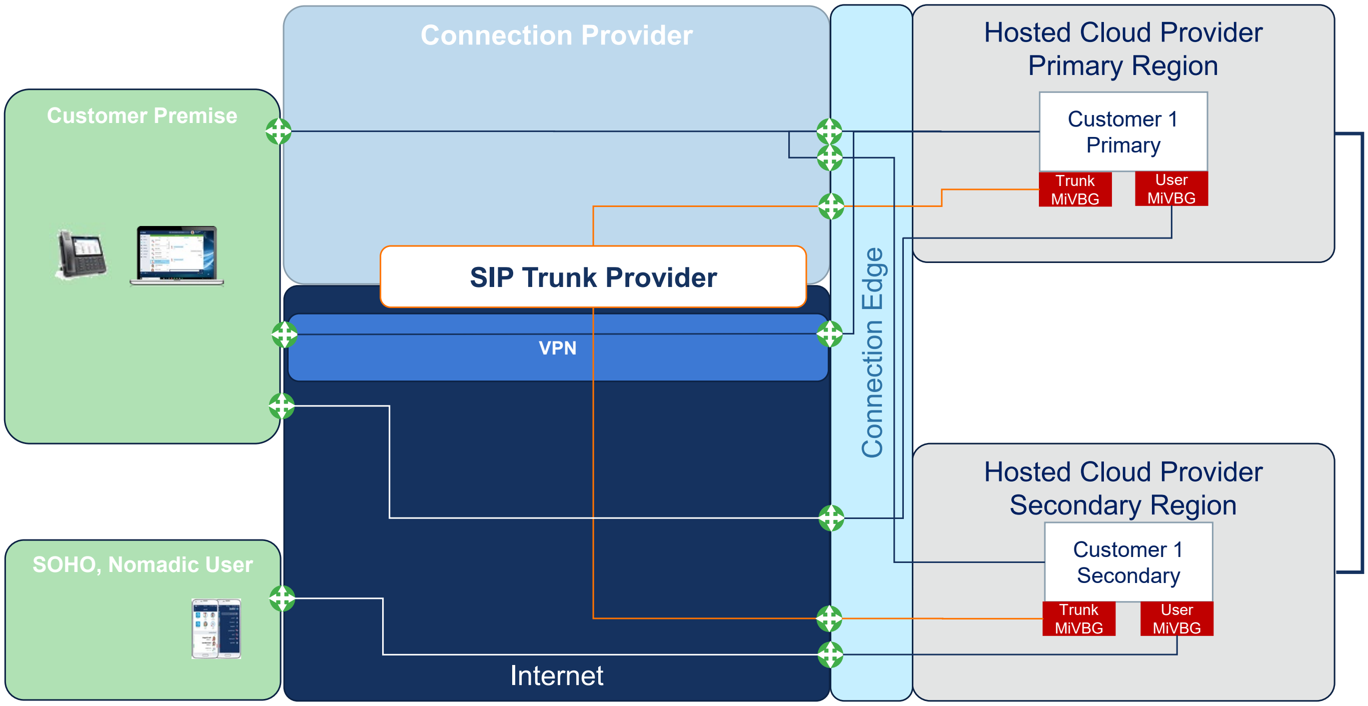

(Not all possible connections are shown to simplify the diagram, for example, the customer premises will connect over the Internet to both primary and secondary user MiVoice Border Gateways, and similarly the SOHO/Nomadic users will connect to both primary and secondary MiVoice Border Gateways.)

The connection provider may be associated with the partner or may be separate. This connection provider has the ability to connect to the hosted virtual networks via the Connection Edge in predefined global and physical data centres ("Cloud Exchange" or colocation facilities) and extend this private circuit to an end customer without resorting to the Internet.

The Connection Edge for Mitel Flex is provided by the Mitel Points of Presence at a number of defined Cloud Exchange facilities. VPN termination is provided over SD-WAN using MiCloud Edge application.

This diagram includes connectivity option for the following:

- Customer connected over a private network connection, using a connection provider, such as MPLS
- Customer connected over the public network, in an OTT configuration
- Customer connected over public network over VPN/SD-WAN
- SIP Trunk connectivity provided over a private network, such as MPLS
- SIP Trunk connectivity provided over a public network

Note that the diagram includes a Hosted Provider to Hosted Provider backbone connection. This may not be available in all geographic regions, and the capability to link the two data centres is then provided by the Connection Partner.
An end customer may also decide to deploy in a purely OTT connection, without use of a private network. As such they would look like a Small Office Home Office (SOHO) or nomadic users, accessing the solution entirely via the user MiVoice Border Gateway.

Although multiple customers may be hosted in the same common hosted data centre, use of segregation techniques, such as VLAN (Virtual Local Area Network), VRF (Virtual Routing and Forwarding, and VNET (Virtual network), ensure that customer networks are not connected, even when over-lapped and private IP addresses are used across common networks and connections.

Customer can only connect to each other via the PSTN either through the SIP Service provider connections, or via a local survivable gateway/SBC on-premises.

Customer Connected over Private Network

For this connection, the customer network, is connected via a dedicated router over the private network of the private network Connection Provider, or MPLS provider in this case. Such a private network will allow a business to cross connect multiple offices in different locations, as a private network. It can also provide multiple connections to the different hosted data centres, and thereby providing resiliency to the hosted solution.

The private network could also be created with SD-WAN but requires that the connections be terminated at some point as they exit into the Connection Edge or Hosted Provider. This can be a complex solution, and not discussed here. Contact your Mitel Account Team if this is a consideration.

There is no NAT (Network Address Translation) involved in this connection, and bidirectional access is possible between customer premises and the hosted solution. The hosted solution is in effect, another set of subnets to the customer.

Customer Connected over Public Network (OTT)

A customer may not have a dedicated and private network connection, or there may be users who are working remotely from home, or other locations. The option to connect to the hosted solution is therefore over the public Internet, or Over-the-Top (OTT)

Users in the office, or at remote locations are treated as Teleworkers and terminate any connections at the user MiVoice Border Gateway (MiVBG). The internal network of the hosted solution is not seen by these external devices as the MiVoice Border Gateway provides Session Border Control/Application-Level Gateway functions (SBC/ALG)

Dual hosted networks can be provided and are cross connected via the inter-region backbone connection, when available. This connection is a requirement for resilient OTT deployments, since an alternative MPLS connection is not available.

Users can register with the primary MiVoice Border Gateway, or if unsuccessful, register with the secondary MiVoice Border Gateway.

MiCloud Edge VPN/SD-WAN Network

MiCloud Edge provides a complete cloud-delivered WAN solution that optimizes business-critical voice and video application performance over any WAN network with enterprise-class network connectivity, deep visibility, and remote monitoring.

A customer may decide to use the MiCloud Edge hosted service to provide private network connectivity between a customer premises and the hosted cloud service. This will utilize SD-WAN VPN technology over the public Internet and will require deployment of physical units at the customer premises. This does not provide customer site to customer site connectivity.

Customers may also provide their own SD-WAN termination (BYOD) for customer termination as long as this is compatible with devices identified in the MSA program.

SIP Trunk Connectivity over Private Network

Typically, this deployment would apply to a private network Connection Provider that has connectivity to a SIP Service Provider, or where that private network provider is also the SIP Trunk service provider.

The operation over a private network is slightly different depending on whether the SIP service provider has a private or a public IP address.

Where the SIP service provider is using a private address, then the MiVoice Border Gateway is configured with a corresponding private address, and this is then routed back to the provider over a different VRF/VLAN back to the hosted SIP service provider. Many customers may be on this common connection.

Where the SIP service provider is using a public address, then the MiVoice Border Gateway must also be configured with a public IP address. If the SIP service provider IP address is public, but *not reachable* over the internet, then the MiVoice Border Gateway will need to be connected over a different (public) VRF/VLAN from the customer connection back to the hosted SIP service provider, in the same manner as per a private IP address.

See also below, the option of using a public IP address, which is Internet accessible

SIP Trunk Connectivity over Public Network

Typically, this deployment would apply where the SIP service provider is separate from the private network Connection Provider. However, it can still apply to a private network provider that offers SIP trunks via a public interface, such as via their own Internet Gateway.

Where the SIP service provider is using a public IP address, which is also publicly accessible, then the MiVoice Border Gateway must be configured with a public IP address. If the SIP service provider IP address is within the public range, and *reachable* over the Internet, then there are two routing options, which involve the use of the BGP protocol. The routing options are then to go over a public VRF/VLAN connection to the private network Connection Provider to the public SIP service provider, or to go via the hosted provider public Internet gateway and connect to the SIP service provider over the public network. BGP is be used to provide the shortest route preference.

Each of these connection options has different advantages. Using a public IP connection provides opportunities to use BGP and provide alternative routing paths, should the preferred path be unavailable.

The mechanism of routing via the private network Connection Provider maintains control of the connections to the SIP service provider and manages the costs of the data usage on the connection.

The mechanism of routing over the public Internet provides an alternative connection mechanism to the SIP service provider, or an ability to connect to any SIP service provider.

MiVoice Border Gateway as SIP Trunk Gateway

SIP trunk connections are established from the MiVoice Business call control to the SIP trunk provider using MiVoice Border Gateway as a SIP-aware firewall and proxy between internal and external networks. Functions provided by MiVoice Border Gateway include, but not limited to:

- 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 (Keypad Markup Language) notifications to support EHDUs (External Hot-Desk Users)
- 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 endpoints defined by their IP addresses and signaling ports. A trunk can have any number of “channels” each of which corresponds to an active bi-directional media stream.

A MiVoice Border Gateway must use a static IP address on the external interface to allow connections to and from the SIP service provider. Often the MiVoice Border gateway IP address is included as part of authentication and included in an “allow” list.

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 alternative 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.
- These methods of providing resilient trunk routing for outgoing calls are shown in the figure.

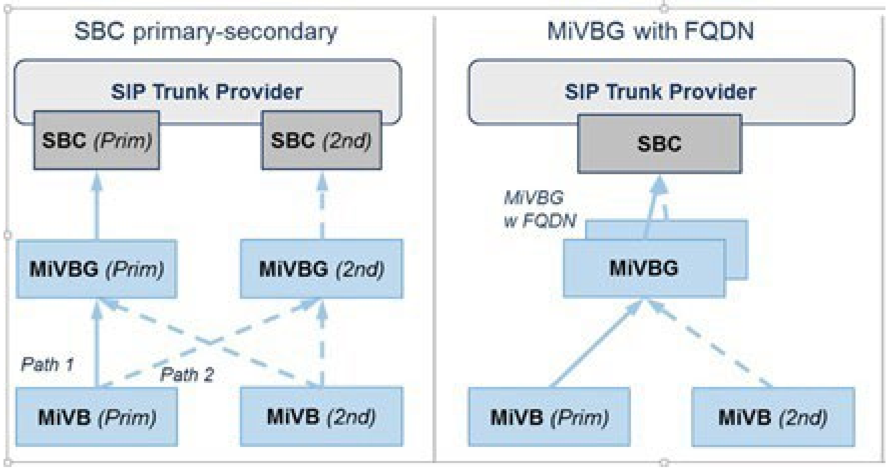

Incoming SIP call routing is determined by the SIP Trunk Service provider routing to MiVoice Border Gateways and by MiVoice Border Gateways routing to MiVoice Business. Typically, third-party session border controllers (SBCs) support primary and secondary routes. The SBC automatically detects trunk failures 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. the following figure shows resilient SIP trunk routing for incoming calls.

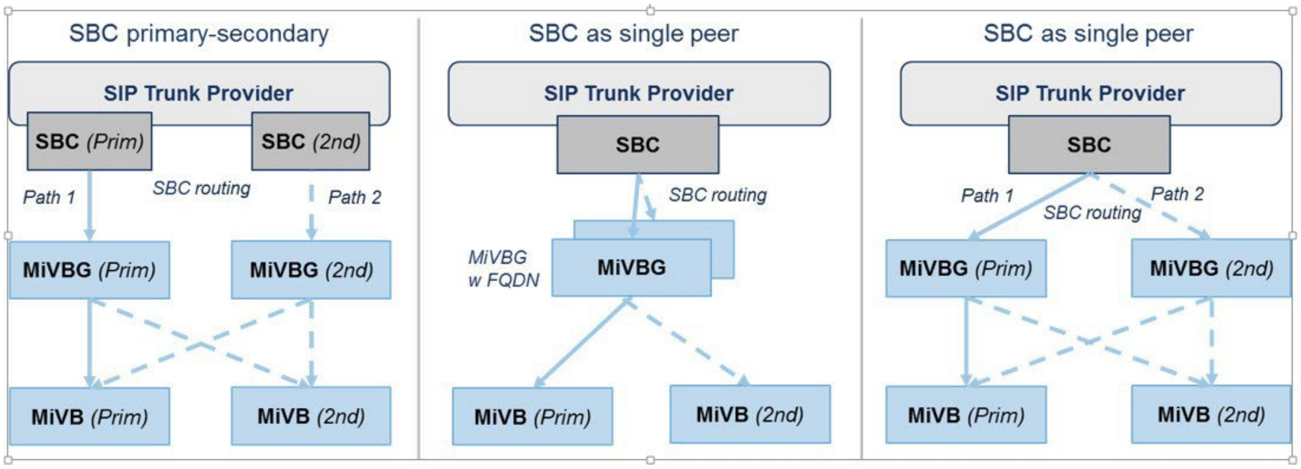

For smaller sites with shared trunking, one MiVoice Border Gateway can act as proxy for multiple MiVoice Business controllers, up to the active trunk capacity limit. For larger sites, multiple MiVoice Border Gateways and MiVoice Business controllers can be deployed for increased capacity. MiVoice Border Gateways support clustering 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.

MiVoice Border Gateway as Teleworker Gateway

Remote Teleworker users may connect to UC services using Mitel IP Phones, or SIP Phones, over public data networks through the MiVoice Border Gateway. This service implements a back-to-back user agent (B2BUA) to terminate connections externally and proxy these through to the internal network. This provides isolation of the connections, providing enhanced security, as well as acting as an Application-Level Gateway (ALG).

The MiVoice Border Gateway offers the minimum Teleworker service:

- Encrypted signaling and media streaming
- Adaptive jitter buffering and Packet Loss Concealment mechanisms
- G.729 compression to reduce bandwidth requirements (Restrictions on use of compression for Call Recording may apply)
- Secure Recording Connector functions for integration with call recording (see below)
- Web Proxy and Port forwarding to applications (see below)
- WebRTC (WebRTC cannot be recorded) and WebRTC Pro (call recording possible)

Use of the WebRTC Pro setting provides connections via Mitel Border Gateway as a SIP phone, providing the same capabilities, including call recording. This functionality is mutually exclusive with WebRTC, which cannot be recorded. WebRTC is the legacy setting and is used by certain applications including MiCollab AWV. Deployment of multiple MBGS may be required if both WebRTC and WebRTC Pro are required. See also MBG Engineering guidelines for scaling of WebRTC support as this also uses the Web Proxy capability

For larger sites, multiple MiVoice Border Gateways can be deployed for increased capacity. MiVoice Border Gateways support clustering with license sharing among cluster members. Load sharing and redirection is available for Mitel IP Phones (MiNet Based) across members of the cluster. SIP Phones use FQDN, DNS SRV Records and primary/secondary registration for resiliency.

Secure Recording Connector (SRC) and Call recording

Call recording requires that media be anchored through a MiVoice Border Gateway, using the Secure Recording Connector functionality, and connected to/from the call recording equipment. This is direct call recording.

Further details on the location of the SRC functionality within the network, and performance impacts are highlighted in the Application section.
See also *MBG Engineering Guidelines* for additional information.

Web Proxy and Port Forwarding

The MiVoice Border Gateway provides a Web Proxy as well as defined port forwarding for a number of hosted applications.

Further details can be found under the Application section.

See also *MBG Engineering Guidelines* for additional information.

MiVoice Border Gateway as Secure Recording Connector (SRC)

The MiVoice Border Gateway provides a number of functions, and one of these includes a function called Secure Recording Connector. This sits in the registration and media path of phones to allow a recording ‘tap’ of the media to call recording equipment (CRE). The connections to and from the SRC and CRE remain encrypted, hence the Secure part of the name.

The media is anchored to the Secure Recording Connector, and this is referred to as Direct Recording, and is the recommended deployment.

MiVoice Border Gateway units can be used to record Teleworker devices and also for Trunk Recording. However, to record IP Phones that are on the internal network, an additional on-net SRC is required. This is effectively another ‘Teleworker’ MiVoice Border Gateway, deployed in LAN mode. This allows SRC licenses to be utilized without additional teleworker licenses.

Web browsers connected through an MBG configured for “WebRTC” cannot be recorded. MBG configured for “WebRTC Pro” can be recorded as these are treated as SIP Teleworker devices.

The internal SRC should be combined with the MiVoice Border Gateways used for Teleworker in order to share licenses. Use of Zones is required to ensure recording licenses are shared, but Teleworker licenses remain with the externally facing Border Gateways.

Further details on the deployment can be found in *MBG Engineering Guidelines*, and under the Application section, below.

External Hot Desk User (EHDU) connectivity

Remote users may connect to UC services using PSTN phones, with tone dialing, 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 Mitel appliance 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 SIP Key Press Markup Language (KPML) to the SIP trunk proxy to provide notification of key press events used to manage call handlingand 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.

External Hot Desk Agents (EHDA) are External Hot Desk Users, with a different configuration and a higher traffic pattern.

Call recording

Call recording for EHDU calls is supported using Trunk recording.

Media must be anchored through a Trunk MiVoice Border Gateway, which is configured for SRC functionality and connects to/from Call Recording Equipment.

Note that Agent/User and Trunk recording are mutually exclusive, i.e., one or the other, not both.

Quality of Service

Network congestion due to high traffic levels can reduce the rate at which packets are transmitted, resulting in delays and increased delivery jitter. As network congestion increases, network switches and routers may be forced to discard packets rather than forward the information, i.e., packet loss. Judgement is required to determine which packets have priority and therefore saved, versus being discarded.

Discarded packets may result in broken audio or artifacts on video and collaborative sharing connections. Although these endpoints include algorithms to compensate, or conceal, lost information, there is a point where the level of packet loss cannot be accurately compensated.

The use of Quality of Service settings can provide the necessary priority marking on the information. It is then up to the network equipment to honor these. However, it also recognized that the Internet is made of a collection of networks, and therefore priority handling cannot be guaranteed. But, QoS markings are still useful within private networks, and especially with any private WAN connections, or network congestion points such as Wi-Fi gateways/base-stations.

Network assessment

Ideally an assessment of the network connections should be conducted prior to deployment of any VoIP or UC installation. 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 (green) row.

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

Network Quality of Service Settings

Two areas where priority mechanisms typically operate in the network are:

- Ethernet (Layer 2) in the LAN, through use of VLANs and packet tagging
- IP (Layer 3) at network routers using Differentiated Services Code Point (DSCP) values in the Type of Service field. This priority setting is carried across most networks, whereas VLAN and tagging information remains local. Many Layer 2 switches can now cross map Layer 3 information to Layer 2 settings or prioritize on Layer3 instead of Layer 2.

Other network technologies bring their own solution to priority marking and handling. One of these technologies included Wi-Fi, which has become ubiquitous enough to virtually replace wired connections, especially in smaller office environments.

The following table provides a view of IP, Ethernet and Wi-Fi priority settings

SERVICE CLASS	L3 PRIORITY	L2 PRIORITY	WMM ACCESS CATEGORY	WMM CATEGORY
Telephony (voice)	46 (EF)	6	AC_VO	Voice
Signaling	24 (CS3)	3	AC_BE	Best effort
Multimedia conferencing	34 (AF41)	4	AC_VI	Video
Standard	0 (DF)	0	AC_BK	Background

Bandwidth Considerations

Even when QoS mechanisms are employed in the network it does not alleviate the requirement to verify that there is sufficient network bandwidth to carry all of the expected traffic. QoS settings can improve packet handling during congestion periods, but if there is insufficient bandwidth, then congestion will be a common operating mode, along with poor quality audio and video.

Bandwidth consumption

In order to determine 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. Consider peak time traffic operations as this is when congestion is most likely to occur. The use of compression zones and IP networking may also have a bearing on traffic flow in parts of the network.

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

- Call traffic patterns
- Number of voice calls and whether they will be compressed
- 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.*

The bandwidth required for third-party applications should be available in the vendor 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 measuring bandwidth consumption, and this is another option for determining bandwidth utilization.

As a rule of thumb, the following bandwidths can be used for a rough estimation:

- Voice calls (concurrent): 100kbits/s (G.711) and 40kbits/s (G.729). Wideband CODECs consume network bandwidth between the two values. Use G.711 as the rule of thumb.
- Video calls (AWV, and video conference such as CloudLink MiTeam Meetings): 2.5Mbits/s
- Screen Recording: 1.5Mbits/s (at 15FPS)

These numbers do not replace more accurate calculations for required bandwidth but can give a quick estimation of the correct range to consider.

MiVoice Business Compression Zones and Bandwidth Management

Coders and Decoders (CODEC) convert analogue voice and video into a digital representation, and back to a user interface, for example, TV, Monitor, loudspeakers and microphones. The digital information can be easily transported long distances without degradation, across many different network types and protocols. This information is referred to as the payload within the transport protocols. Different CODECs can provide different sized payloads given the same input information. A reduction in payload is often referred to as compression.

Compression of audio calls can be used to reduce the bandwidth requirements across links with limited bandwidth, such as external WAN connections. However, it should also be noted that compression usually results in loss of information and fidelity of the information. There is therefore a choice in terms of quality, or cost. Certain applications, such as Speech Analysis and Speech To Text may require data in an uncompressed format to correctly transcribe the information.

The MiVoice Business includes 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 bandwidth 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 connections and links.
- Call Admission Control uses this information to decide whether a call should be completed or not.
- Emergency calls are always allowed through, even if it results in other users receiving degraded quality.

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

Speech-Analysis on call recordings is best achieved on data that has not been compressed. For best performance, with Speech Analysis, it is recommended that compression zones and CODEC filters are set to ensure that compression CODECs are not used up to the recording equipment and to the SRC units.

Network Infrastructure for IP Phones

IP phones require some basic networking infrastructure, in order to power up and register with an appropriate MiVoice Business call control, or MiVoice Border Gateway.

Some key components that need to be provisioned and configured include:

- DHCP services
- TFTP services
- DNS services
- Network Switch

DHCP

A phone initially requires an IP address to get started, and for most deployments this is easily achieved through use of a DHCP server. The phones can also be programmed manually, but this limits their ability to be moved within the network, but still possible. Most basic networks, including home office environments support a basic level of DHCP to provide an IP address.

However, other useful information can be obtained via DHCP including network configuration and location of other services, such as where to find the call control or TFTP server. Some of this information uses options within DHCP which may not be available on all DHCP servers.

Options for providing a DHCP server with added options include:

- Deploy a dedicated DHCP server on a local on-site connected Windows Server, or similar
- Use the DHCP server that’s included with the local survivable gateway

DHCP is typically limited to the local subnet, and so a cloud hosted DHCP server may be of little use to a remote office location. Therefore, DHCP services need to be provided locally at the customer premises. Use of DHCP Helper or forwarder can be used to enable DHCP to reach between subnets and VLANs.

TFTP

A TFTP server is needed by the Mitel IP Phones to download updates to the phone software. The phones retain an existing version of software over re-boot but require the TFTP service for updates.

A suitable TFTP server is available with MiVoice Business for internally connected phones. A suitable TFTP server is also available at the MiVoice Border Gateway for externally connected phones.

DNS

A DNS (Domain Name Server) is required to provide IP address information to a Fully Qualified Domain Name (FQDN) request from an end-device or application. Use of FQDN makes the deployment simpler for a user to manage, rather than IP addresses, and also provides the capability to provide multiple IP addresses and therefore provide application access resiliency.

Use of FQDN is recommended, although certain legacy phone types may not fully support this.

DNS services are provided both internally to the customer network, and also in the public domain, i.e., on the Internet. Use of Split-DNS may be needed for some applications, including the MiCollab Client application, for example, internally connected devices may connect directly to the MiCollab application, but devices on the public Internet will resolve the public MiVoice Border Gateway, which will the proxy/forward the connection the appropriate MiCollab application.

Multiple DNS records (A-records or SRV records) can also be used to provide resiliency by supplying multiple IP addresses, for example, to SIP Phones.

Network Switch

In many cases additional information can be provided to the phones, based on the network the phone is connected to, depending on the protocols available on that port. See the *MiVoice Business Engineering Guidelines* for further details

In summary, some information that can be obtained include:

- Power requirements
- Port identification (may be used with emergency call solutions)
- VLAN and QoS settings

Obtaining Network Parameters: Mitel IP Phone

Mitel IP Phones, have a number of different methods that can be used to obtain networking parameters such as VLAN and QoS information. The Mitel IP Phones obtain information from the different available sources, working through the list of options, starting at Priority Level 5 and decreasing until information is obtained.

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.
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 theLAN.
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.

Obtaining a Call Control IP Address

Nominally, a Mitel IP Phones, on an internal network, will obtain the IP address of the call control via DHCP, or be manually configured. However, this may not be the case for Teleworker phones, if the IP address is not previously manually programmed.

When a Mitel IP phone is first powered on in Teleworker mode, it attempts to find the IP address of the call control, or 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:

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

Often phones are not pre-configured before delivery, and customer DHCP servers (for example, home office) do not support the necessary DHCP options. In this case the phones revert to the externally programmed Redirect and Configuration Server (RCS). Note that MiCollab uses a different Redirection and Configuration (RC) server. The RCS is for Mitel IP Phones.

Redirect and Configuration Server (RCS)

The Mitel IP Phones can use the Cloud hosted RCS server to locate key connection information. This is typically used for initial Teleworker deployments, and allows deployment to OTT based customers, without resorting to an on-site visit. The phone will initially look for locally available information including statically programmed information, previously supplied information and DHCP information. If none of these are successful, the phones will look externally to the RCS service.

The partner or service provider offering the phone service needs to configure the RCS for the phone MAC and MiVoice Border Gateway.

Further details can be found under the *Redirection and Configuration Service (RCS) User guide* on Mitel Doc Center

LAN Power considerations

In most cases, the Mitel IP Phone will draw power via the Layer 2 Power-over-Ethernet (PoE) enabled network switch. Local power solutions are also available, when this configuration is not possible, for example, for a home office environment.

Use of central PoE network switches allows for central management of power provisioning and backup, as well as remote management and monitoring of the connection.

Additional information on PoE requirements can be found in the *Mitel IP-Sets Engineering Guidelines* and the L2 switch and router product documentation.

If connecting an Ethernet device at distances of more the 100 meters, or where alternative cabling (for example, non-CAT5) plant is available, the Mitel StreamLine solution might be appropriate. The Streamline is a long-haul Ethernet switch that can provide Ethernet connectivity with PoE over distances of up to 360 meters. For details refer to the Streamline documentation on Mitel Documentation Center.

Deployment Considerations

MiCloud Flex is hosted across global Mitel data centers that leverage tried and true, robust infrastructure virtualization from VMware. Mitel data centers provide a dedicated, secure, and highly reliable communications solution for every customer. Each customer's system can be deployed across geographically redundant Data Centers offering 5 9's protection for core voice services, with options for Disaster Recovery (DR) and Business Continuity (BC) protection for all applications. Where sufficient to meet customer needs, a more cost-effective single Data Center 4 9's solution leveraging VMware virtualization protection and recovery can be deployed.

The customers are deployed as a private hosted solution. This allows the deployment of a customer as a unique Over-the-Top (OTT) only solution or may be combined with a private hosted network as a private network extension of the customer network and address space. Each customer is isolated from other customers through use of Software Defined Networking, including use of VRF, VLAN and virtualization technology. From a hosted solution view it is also possible that multiple customers may be deployed with overlapped IP addresses on the same platform yet remain isolated.

Management access to the solution is from the customer network, or via dedicated network access portal. The Mitel Performance Analytics (MPA) client, or probe, provides this capability back to a cloud service. Access to the cloud service is authenticated by partner or customer, and a dedicated and secured tunnel is established with the on-net MPA client for access. As well as providing a secured tunnel access, MPA also collects key metrics on the performance of the solution, and this can also be accessed from a dashboard in the MPA Cloud service.

Further information can be obtained from the *MiCloud Flex General Information Guide* on InfoChannel.

Location of Flex Data Centers

The most current information is available in the *MiCloud Flex General Information Guide* on InfoChannel.

Solution Configurations

Solution deployment configurations include:

- Single Data Centre (with HA)
- Dual Data Centre (with HA)

Single Data Centre Deployment

In the case of the single data centre, this is simply an extension of the customer network. Most of the applications will fit into a single subnet and can be deployed as such. This minimizes any inter- region router configurations.

Deployments in a single data centre take advantage of any High Availability capabilities to minimize any outage and downtime. See the section below on High Availability for deployment considerations.

Teleworker and SIP Trunk connections can be delivered over the Internet into the Data Centre. Connections to an MPLS network can still take advantage of dual data links between hosted data centre and provider private network to minimize outages on this connection, for example, network upgrades, or accidental link outage.

IP addresses in this hosted subnet, or subnets, should not conflict with other subnets used by the customer.

Dual Data Centre Deployment

The Dual Data centre deployment uses two different and geographically separated data centres in different physical regions, i.e., not in the same city, but in different cities, or maybe east coast and west coast.

Connection between the two hosted data centres may be provided through a backbone connection of the infrastructure provider, or may require dual connections from a private network provider. The private network provider will have connections to each of these data centres, providing resiliency within the private network. Each data centre deployment will include a separate MiVoice Border Gateway and multiple public IP addresses for OTT resiliency.

Each of the data centres in a dual data centre deployment can be treated as a single data centre for High Availability purposes. This High Availability capability is not extended between the data centres. Voice application resiliency is provided through dual deployments in the two data centres. Disaster Recovery as a Service (DRaaS) powered by Zerto, extends between these two data centres for applications that require this optional capability.

DNS entries for OTT connections, and also internal DNS entries must include both primary and secondary application IP addresses as part of the FQDN lookup.

IP addresses in these hosted subnets should not conflict with other subnets used by the customer.

Connection Configurations

A number of connection options for the customer from the customer premises to the hosted cloud service are highlighted below:

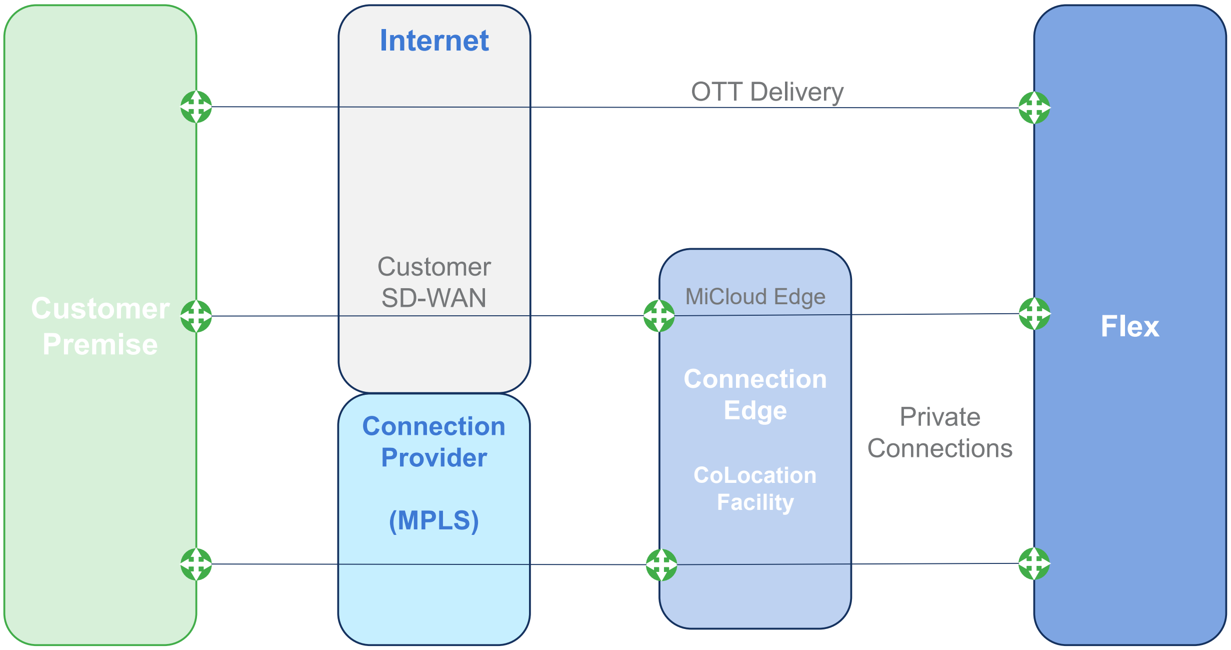

These include the following options:

- OTT Delivery (over the Internet)
- SD-WAN VPN connectivity via MiCloud Edge
- Private connection via a Connection Provider, such as MPLS

OTT Delivery

OTT connections are via the MiVoice Border Gateway. This gateway provides a number of Session Border Control (SBC) and Application-Level Gateway (ALG) functions. Specifically, it can be used to terminate connections for Teleworker devices. Teleworker devices include Mitel proprietary IP Phones, as well as SIP Clients. Additional functions can also be proxied and forwarded by the MiVoice Border Gateway, including connections to the Contact Centre solution and MiCollab Client and conference facilities. (A number of MiCollab features are also provided via Cloud services, on the Internet, such as CloudLink Chat and CloudLink MiTeam Meetings – these do not go via MiVoice Border Gateway). Note that direct connections over the Internet may not provide the full set of functionalities available to on-net users.

The MiVoice Border Gateway are provided with public IP addresses for connection to the Internet.

The MiVoice Border Gateway can be scaled by increasing the quantity of units and clustering them together. This will ensure correct hand-off and registration of the end devices to the appropriate and active gateway.

For smaller deployments the user and trunk MiVoice Border Gateway may be combined. For larger deployments, it becomes simpler to deploy and manage MiVoice Border Gateways as separate clusters or pairs for users and trunks respectively.

MiCloud Edge SD-WAN VPN

Customers can connect to the hosted solution through use of the MiCloud Edge SD-WAN VPN service. This is a connection over the public internet. SD-WAN provides the capabilities for connections via multiple internet providers for improved availability and load sharing. It is also a private network connection, or Virtual Private Network (VPN). The VPN provides a secured and private tunnel between customer and the hosted solution.

Use of SD-WAN requires the customer to deploy a local router with connections to an Internet service provider, or multiple providers. The termination of the connection is at the MiCloud Edge service, where customers are isolated from each other and VLAN connected to the hosted solution. The customer local SD-WAN router can be provided as part of the MiCloud Edge service, or may be provided by the customer (BYOD) as long as this is a supported device in the MSA program.

The advantage of this configuration is the ease and simplicity of setup. It does not rely on a connection partner, other than an Internet service provider, or providers.

The VPN connection has the effect of extending the customer network into the hosted services, being simply another subnet or subnets, within the customer network. This will allow the full suite of functionality to be provided for the solution. Just ensure that there is sufficient bandwidth allocated to the connections for the services being deployed.

Although a VPN can provide connection between two locations within a network, it has been found that when multiple VPNs are used to connect different locations within a business that not all routing conditions may be included. As a result, it may be possible for a HQ office to connect to a remote office, but not necessarily for the remote office to connect to the hosted provider via the HQ location and across two VPN links. In such a configuration with multiple offices and locations it may be necessary to deploy dedicated routers at all locations to the hosted provider. Alternatively, these remote offices can also be terminated at the MiCloud Edge, providing a single termination point for all remote sites.

Note that MiCloud Edge is used to connect a customer network to the hosted solution. It is not intended for deployment and connections between customer premises locations.

Customer Private Network Connections

The customer may connect to the hosted solution via a private network or Connection Provider. The terminations from this provider will be within a “Cloud Exchange” colocation facility, or Mitel PoP. This connection may use private network technologies such as MPLS to maintain the private network connectivity.

Advantages for the customer is that this can be treated as a separate connection, or subnet, to the existing infrastructure, with minimal disruption. The provider may offer a Service Level Agreement, which cannot be guaranteed over the public Internet. A connection partner may consolidate multiple customers to the Connection Edge, or Mitel PoP, and offer cost advantages due to bandwidth consolidation.

The private network maintains isolation within that network, and this isolation is extended into the hosted service through use of VRF and VLAN technologies.

SIP Trunk Service Provider Connections

SIP service providers may offer services via different methods. Typically, these methods include

- Via a dedicated MPLS circuit (either public or private address)
- Via a public Internet reachable connection

The first deployment with a dedicated circuit means that a particular VLAN or circuit needs to be accessible to all customers, or rather via each customer trunk MiVoice Border Gateway that is assigned to this function. Each of the gateways shares this common connection to the private network provider. MiVoice Border Gateways of multiple customers cannot communicate directly with each other, only via the SBC of the SIP trunk provider termination. The IP address used for this trunk MiVoice Border Gateway may be private, or they may be a non-Internet routable public IP address.

Although it is possible for a single MiVoice Border Gateway to be used for both user and trunk access, they must share a single IP address. With this private trunk configuration, the two functions must be separated onto different interfaces, and hence different groups of MiVoice Border gateways in the hosted customer deployment. For a small installation, the consideration of using a “three legged” MiVoice Border Gateway may arise. However, this requires a custom configuration to achieve this, and is not recommended for Flex deployments.

The second deployment is simply using the MiVoice Border Gateway as a trunking gateway with a public and Internet routable IP address. This connection can connect to any SIP Trunk Provider that offers a public connection via the Internet. For smaller deployments a single MiVoice Border Gateway may be possible, but this requires that a common IP address is used for both Users and Trunk connections. For larger deployment these two functions will typically be provided by different groups of MiVoice Business units.

Emergency calling and E911 Support

E911 is a requirement for North America, and associated with an emergency call being made, typically through dialing the emergency number ‘911’. Similar mechanisms may exist in other jurisdictions. Partners and service providers may need to provide additional DID numbers or location identifiers to identify users and phone devices to the emergency services. Partners and customers must ensure that any DID, or ID, and location information are maintained and updated to ensure that any emergency calls direct the emergency services to the correct location (US RAY BAUM’S Act)

Partners must ensure that the dialing plan allows an unfamiliar user, for example, someone not familiar with the office or customer system, to be able to dial only ‘911’ (or local jurisdiction emergency number) and still reach external emergency services, i.e., it must not be necessary to dial extra trunk digits, such as ‘9911’ in order to reach external emergency services. (US Kari’s Law)

Mitel has provided a bulletin covering relevant information, available to channel partners on InfoChannel.

Customers and partners may decide to deploy other emergency numbers to reach external emergency services, for example dialing ‘112’ or ‘999’ to also cover visitors to the facility. Dialing ‘911’ outside of North America, may not result in location identification information (CESID) being presented, if services in those regions are not configured to accept this information.

Emergency calls are presented from the MiVoice Business system to the SIP service provider. For E911, additional location ID information, known as a CESID, is presented to the SIP service provider to forward this to the Emergency Handling Service (aka PSAP).

Location and call-back number information is provided in advance to the Emergency Handling Service and associated with a particular CESID reference, rather than specific location information being transmitted on the emergency connection. This maintains user privacy, allows multiple users to use a common CESID, or even change CESID, and provides consistent information in a dire situation. The MiVoice Business contains a database of user, devices and associated CESID reference. The MiVoice Business provides dynamic updates as needed to this database to ensure location information is current for user and device.

Information about an Emergency call can also be forwarded to external applications such as the Mitel Revolution product, or to the Mitel Performance Analytics probe/client. This is triggered by SNMP Trap information from the MiVoice Business handling the call. The Mitel Revolution product or MPA can then inform specified users/managers in the deployment of the emergency call being made. This can be used to identify the user and may also allow these people to handle call back from the emergency services, and also direct any emergency services that arrive at the business front door. This application will work in all jurisdictions, and not limited to North America.

Updates to SIP SP

Information on CESID to location is provided via the SIP Service Provider. The granularity of location information is dependent on local jurisdiction requirements, the business installation, size and location of the building and users. The information required could simply be a civic address, leading to the business front door, or may include additional location information such as floor and region on the floor. Check the local jurisdictions for location granularity requirements. In the default case, the civic address of the building and a call-back number shall be provided.

Database in MiVoice Business

Moves within the network can be identified with the Mitel IP Phones using port information from the network Layer 2 switches. If a phone is moved, the CESID associated with the phone is updated based on the network information. If the network information is not recognized, an alarm is raised to the administrator to provide update information.

For SIP phones, the registration IP information can be used to identify zones within the building and CESID information can be applied to these zones. The zone granularity is dependent on the region assigned to the zone and how network sub-nets are assigned to physical location.

In order to use these capabilities, the partner and customer must correctly identify zones and network connection locations and populate the information into the appropriate databases, and to continue to update and maintain this information.

MiCollab Mobile Softphones

MiCollab Mobile Softphones are always deployed as Teleworker phones and are considered nomadic devices. Identified emergency calling numbers are provided by configuration within MiCollab and may also be recognized by certain operating systems (iOS and Android are known to support this). When an emergency number is recognized, the softphone will use the native phone dialer and calls will be made over the cellular network where triangulation techniques will identify the location of the caller. The partner must ensure that suitable emergency calling numbers are identified in MiCollab, and this feature enabled.

Teleworker Phones

Teleworker phones are nomadic devices and may also change location. However, most Teleworker office-based phones tend to stay in the same remote location.

For static, but external office phones, a CESID can be assigned to the phone and the location information updated via the SIP Service Provider. Just remember to update this information if the phone moves within a large building, and/or to a different civic address.

In the case where a CESID is assigned, but not programmed, the default office civic address location will be provided. This may not be the current location, and the user will need to confirm.

In the case where no CESID is assigned, the presented information is the Calling Party Number for regular calls on the system. Again, this may point to the civic address of the office building, unless this user is assigned a specific DID number, in which case this number will be presented to the emergency services. Again, the location information may not be current, and the user will need to confirm.

It is important that users update and maintain their location information when possible. It is the customer’s responsibility to have a process in place which will allow the information to be updated.

PC Softphone

The PC softphone is similar in operation to the MiCollab Mobile Softphones for laptops and tablets, except without the cellular network. For fixed location PCs, or workstations, and a fixed IP addresses within the office, a static CESID can be assigned to the device.

A mobile PC, or tablet, within the office building will pick up different IP addresses as it moves and reconnects to the system. The IP address can be associated with a zone and this zone associated with a CESID location, in the same manners as SIP phones (the PC softphone is a SIP phone).

PC softphones located out of the office will be treated as Teleworker phones. It is the customer’s responsibility to have a process in place which will allow the information to be updated.

DECT and In-House Wi-Fi

Typically, Wi-Fi covers a large area, but this can also be divided into zones and ranges of IP addresses. As a device appears on a different part of the Wi-Fi network, it will be pick up a different IP address, which can be associated with a zone, which can be associated with a CESID location. If this is not possible, then the default CESID needs to be applied to all Wi-Fi devices.

DECT uses a different wireless protocol and may roam between RFPs while still maintaining the same IP address of the primary base station. In this case, the IP address could be assigned to a zone and CESID, but typically will be assigned to the wider default CESID and civic address of the building.

Mitel Revolution

The Mitel Revolution optional application provides a way to centrally manage creating and sending notifications. This application can be used to send emergency and non-emergency notifications through a number of different mechanisms, such as Instant Messaging clients, SMS, Mitel Revolution Desktop Notification Client; Paging Relay, Analog Systems; IP Speakers, Message Boards; Social Media accounts; and more.

Emergency call trigger information is provided to the Revolution application through SNMP Traps from the MiVoice Business units in the solution.
Further information on this application can be found on Mitel Doc Center.
Please contact your Mitel Account Team if you wish to deploy this feature and service.

MiVoice Business Clustering

MiVoice Business units are typically clustered together in a deployment, to share database information and simplify routing and resiliency operations. Clustering allows scaling as well as hierarchical call handling in larger deployments. Primary and Secondary MiVoice Business must be clustered in order to share database information about the users, and correctly handle hand-off between the units.

For a large deployment, including use of multiple devices per user, it would be common to have multiple MiVoice Business in the primary location clustered together, and also clustered with a similar number of multiple MiVoice Business in the secondary location.

In a contact centre deployment, MiVoice Business may be clustered for resiliency of agents, and additional MiVoice Business may be deployed for functions such as trunk termination and queuing with IVR. See Mitel Doc Centre for *Contact Center Blueprints*, for further information.

In a mixed environment of contact centre agents and back-office staff, the MiVoice Business will typically be deployed based on their functional operation. However, these units can still be clustered together to simplify the numbering plan and transfer of calls between units.

If remote survivable gateways are used, these also need to be clustered into the MiVoice Business network, as these units will provide secondary operation in the event of loss of network connectivity.

For larger designs, MiVoice Business may also be used as dedicated call routers to simplify call routing between different systems. Such a deployment might be needed when there are multiple remote offices with multiple survivable gateways at these remote locations. Such a deployment allows scaling and handling of different network topologies. Such a system can also become complex to design and deploy. Use of Professional Services is recommended. Contact your Mitel Account Team, in such situations.

On-Premises MiVB and Survivable Gateway

It is possible to deploy a MiVoice Business 3300 Appliance, SMB Controller (SMBC) or EX Controller, as a local On-Premises Survivable Gateway. This appliance can provide alternative SIP Trunk connectivity (with an appropriate local MiVoice Border Gateway), PSTN connectivity via PRI/ISDN (not available on SMBC) or analogue trunks, as well as connectivity to local SIP ATA gateways and analogue POTS phones.

It is also possible to deploy a MiVoice Business solution on dedicated servers. This will provide connectivity for IP based phones but will not provide the PSTN/ISDN and analogue connectivity of the appliance unit. SIP trunk connectivity is still possible via an appropriate MiVoice Border Gateway.

The On-Premises MiVoice Business will typically be deployed as a secondary controller to the hosted MiVoice Business. In the event of a network failure, this controller will continue to handle calls for the local users. In the event of a network failure, it’s possible that external calls will still terminate at the hosted solution SIP Trunks. Calls should be forwarded on no answer to the hosted voice mail system. Some SIP SP providers can also be configured to divert on no answer to a tertiary, which could be the survivable gateway.

When a survivable MiVoice Business is deployed, it may be desirable to also include a MiVoice Border Gateway for Teleworker phones and also for SIP trunk terminations. In this case the MiVoice Border Gateway will be deployed on a dedicated server, or on a dedicated virtual server. For Teleworker phones, this MiVoice Border Gateway should be added to the MiVoice Border Gateway cluster to allow phones to have an alternative gateway. However, this also adds to the number of units in a cluster.

The local MiVoice Border Gateway can be configured as either primary or secondary for the user local devices, with the following operational outcome in the event of a lost network connection. Note that a lost private MPLS connection may include loss of Internet connectivity as well.

- Local MiVoice Border Gateway set as primary:
    - Internet still available: Phones will still connect as normal to the local network. Call handling will be limited to all devices of the user to outgoing and local calls.
    - Internet lost: Devices on the local network can make PSTN outgoing calls and local calls. Teleworker phones will rehome to the secondary hosted MiVoice Border Gateway and will have existing incoming and outgoing call handling. Calls to local users will be via PSTN.
- Local MiVoice Border Gateway set as secondary
    - Internet still available: Devices on local network can make outgoing and local calls. Teleworker phones will continue to handle calls as normal. Calls to local users will be via local trunks
    - Internet lost: Devices on local network can make PSTN outgoing calls and local calls. Teleworker phones will continue to handle calls as normal. Calls to local used will be via local PSTN trunks.

IP Addresses, certificates, DNS updates and domain names for the public facing MiVoice Border Gateway are provided by the customer.

Contact your Mitel Account Team when deploying local survivable controllers and gateway, and they can advise on the best configuration for the business requirements.

MiVoice Border Gateway Clustering

A number of MiVoice Border Gateway servers can be clustered together to provide load sharing for Mitel IP Phones (MiNet based teleworker) and also to provide primary/secondary availability options. When clustered, the MiVoice Border Gateway will share database information regarding registration and licenses, allowing devices to register and connect through multiple units.

The load sharing on the MiVoice Border Gateway is configurable and can be adjusted dynamically. This can be used to allow maintenance and upgrades to particular units, and load share to other units in the cluster. Consider this when deploying MiVoice Border Gateway, i.e., consider that N-1 units may need to take the complete load at some point. Load sharing is only applicable to Mitel desk phones using the MiNet signaling protocol.

There are both internal network and external network clustering connections between the different MiVoice Border Gateways. This is to ensure that in the event of a network failure, but still functioning MiVoice Border Gateway, that phones are directed to an appropriate unit that can still handle the call connections.

The MiCollab includes a MiVoice Border Gateway unit. For small deployments, this unit can be used to provide external connectivity, and reduce the need for an external server. However, this is limited in capacity, and also to a single data centre deployment. For larger systems and when resiliency is needed, the use of external MiVoice Border Gateway is required. The MiVoice Border Gateway on MiCollab, is still connected to the cluster of MiVoice Border Gateway, but with a load capacity of 0%. This unit is still used by MiCollab to provision users and share licenses within the cluster.

As well as providing external gateway functions, the MiVoice Border Gateway can also be used to provide Secure Recording Connector (SRC) functions. This feature is used to allow through connected devices to be connected to Call Recording Equipment (CRE). Additional streams are generated to the CRE. This adds additional load to the performance of the MiVoice Border Gateway. See *MBG Engineering Guidelines* for further details on performance considerations. Note that in a resilient call recording solution, there may be multiple recording streams and considerations.

For externally facing MiVoice Border Gateway, these units can also provide SRC functions for external devices. For internal devices, an internal MiVoice Border Gateway may be deployed. Internal devices are registered and connected via this internal MiVoice Border Gateway/Secure Recording Connector similar to Teleworker devices, but on the internal network. See Mitel Interaction Recording under the Applications section for further details. The MiVoice Border Gateway is deployed in LAN mode to minimize license requirements.

It is possible that on larger systems, especially with a large number of applications that the number of MiVoice Border Gateway in a cluster exceeds the limits identified in the MBG Engineering Guidelines. If this situation occurs, please contact your Mitel Account team, who can then advise on the solution and limits which may be applied.

See *MBG Engineering Guidelines* for additional details. Items for consideration in the hosted solution include:

- Number of SIP trunk routing rules
- WebRTC Connections/Setting. These units do not load share and cannot be call recorded. There is a limit per MiVoice Border Gateway, This limit is identified under the WebRTC section of the *MBG Engineering Guidelines* under the Application section
- WebRTC Pro Connections/Setting: These connections are treated as SIP Teleworker devices and consume licenses up to the purchased limit. The limits of Web Proxy connections apply. These calls can be recorded. However, WebRTC and WebRTC Pro settings are mutually exclusive on the same MBG unit. This may require multiple deployments if both functions are required.
- Web proxy and port-forwarding. Support is provided for Mitel specific Applications. See *MBG Engineering Guidelines* for further details.

MiCollab

MiCollab provides a number of collaboration and provisioning functions within the solution. A number of these are identified below:

- Provisioning tool for the solution, including license management for the MiVoice Business, and cluster, and also for clustered MiVoice Border Gateway, via the local attached MiVoice Border Gateway. MiCollab provisioning is also used to enable CloudLink functionality (see below)
- CloudLink chat: Once enabled via MiCollab, this feature is offered as an OTT solution for users.
- Audio, Web, Video collaboration – This collaboration tool is provided within MiCollab and is accessible for on-net and public connections (via MiVoice Border Gateway) – see also CloudLink MiTeam Meetings
- CloudLink MiTeam Meetings: This is a public cloud service collaboration tool providing global access for PC, and Smart-Phones, once provisioned via MiCollab. This premium service is provided to users with the appropriate licenses.
- Integration with Office365 and Azure AD to provide wider Single Sign on (SSO) capabilities, as well as authentication configuration and access to CloudLink services.
- NuPoint Voice Mail and advanced UM
- MiCollab Client (and presence information) and MiCollab SIP softphone (externally connected via MiVoice Border Gateway)
- MiCollab Web Client and web based PC softphone (externally connected via the MiVoice Border Gateway)

System Backups and Upgrades

With respect to backups and upgrades, this creates the following responsibility requirements:

Backups:

- Mitel maintains 7 days infrastructure snapshots, which are collected nightly. This snapshot does not allow direct access to the applications, nor customer data, but is a block capture of the complete virtual machine.
- Partners are responsible for their own application and customer data backups. This can be saved to a partner storage location, for example, local laptop, or on the customer site subject to customer agreement. Mitel does not maintain a backup server for this service.

Upgrades

Different methods apply depending on whether the applications are based on the MSL Operating System or Windows operating system. The different processes are outlined below. Refer to the particular application documentation for more detailed information.

- For MSL based application, the preferred method is to use the software Blade upgrade mechanism. Take a snapshot of the machine prior to upgrade. MiCollab includes a number of applications, and these are upgraded from the Application tab. On successful completion and verification of the upgrade, the snapshot can be deleted.

- For Windows based applications, start by taking a snapshot of the server. The server and application are then upgraded according to OS and application requirements. On successful completion and verification of the upgrade, the snapshot can be deleted.
- These processes are needed due to different levels of access available at the infrastructure and application layers. The infrastructure team should not be aware of customer data, and likewise the customer should not be aware of the infrastructure. This is part of the shared responsibility model.

System Traffic

The level of traffic, the number of users and their UC profiles drives much of the scaling calculations and eventual product deployment and unit quantities. User UC profiles drive the number of devices ina solution. The purchasing tools, including CPQ, include information to complete these calculations based on the user requirements and input.

Much of the relevant information is provided here, and also detailed under the corresponding product Engineering Guidelines.

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 different UC profiles, ranging from simple standalone devices, such as a lobby phone, through to users using multiple phones.

Further details on the different UC profiles, the number of associated devices, and functionality is provided in the General Information Guide.

Traffic and scaling calculations need to consider an average number of devices per user in order to scale the different services and functions

As a guideline, different servers have different defining limits, which typically fall into the following categories:

- Number of Users
- Number of devices
- Number of concurrent streams

For UC deployments, the average number of devices per user can range from 2-3 devices, with 2.75 being the expected normal. For calculations, a simple rule of thumb of 3 devices provides some levelof margin for future expansion. Typically, these devices would be:

- Desk phone (internal network)
- Mobile phone, registered via MBG
- Teleworker Phone, registered via MBG, either as a dedicated IP Phone, or as a PC softphone

Office phones are often deployed as hot desk enabled devices. Although there are no users assigned to these devices, they also require a device licenses, as they also consume a Directory Number (DN) from the system, i.e., this is also counted as a device license.

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.

High Volume Contact Centre deployments would consider:

- 27 Calls per hour (27CPH)
- Hold times 100-120 seconds per call

Note: Customers wishing to use different traffic rates, or assumptions, are advised to contact their Mitel Account Team to confirm the 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, typicallyin the morning and early afternoon. Traffic and resource occupancy can be expressed in Erlangs or Centum Call Seconds:

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

Traffic Summary

Despite all the complex calculations, these definitions result in the following simple rules of thumb:

- Assume 5 UC users to 1 trunk
- Assume 1 Agent to 1 trunk
- Assume 50% addition to trunks for contact centres when IVR is deployed, i.e., 1.5 trunks peragent

Note: Customers wishing to use different traffic rates, or assumptions, are advised to contact Mitel Professional Services in order to confirm their scaling and quantity of unit calculations.
With respect to backups and upgrades, this creates the following responsibility requirements:

Applications

The following sections provide more detailed information with deployment and scaling of the applications as part of the overall hosted solution.

Phone support

All currently available Mitel IP Phones are supported in the hosted deployment. However, as newer functionality is included into the solution, and as new installations are deployed, there will be increasing emphasis on use of FQDNs on internal and external DNS servers. Certain legacy phones may not be compatible with use of FQDNs, and care needs to be exercised in using these devices for newer installations or with newer features and applications that include FQDN.

The following phones are identified as working with FQDN (April 2021):

- 5304, 5312, 5320, 5324, 5320e, 5330e, 5340e, 5360 (Note *1)
- 6905, 6919, 6920, 6930, 6940, 6970
- MiVoice Business Console
- Wireless and DECT Phones
- SIP DECT
- RFP 12, 44, 45, 47, 48
- SIP DECT Handsets: 612d, 622d and 623d
    - IP DECT
    - Base Station
    - IP DECT 5613 and 5614
    - Single Cell
    - RFP12
    - Handset 112 DECT
- MiCollab Softphones including:
    - UC Endpoint SIP Softphone
    - Web based PC softphone (WebRTC based)
    - ACD Hot Desking SIP Softphone

Note *1: Earlier versions of 5300 phone, including 5302, 5330, 5340 (non-'e' version) and those withonly 10/100Mbps Ethernet connections cannot be upgraded to support FQDN. FQDN support is included in Phone Firmware load 6.5.0.128 and higher. MiVoice Business 9.1 SP1 includes this firmware load.

Support for FQDN in the 69xx range of phones is provided in MiNet release 1.5.2 and beyond.

It is recommended that all new installations use FQDN and only phones that support this. For legacy installations that are not using FQDN, then phones that support FQDN can be used along with phonesthat do not use FQDN. Some features that only support FQDN may not be usable, such as supplying multiple DNS A-records for resiliency purposes.

For further details on these devices refer to the 6900 Phones, Conference Phones, MiVoice Business Console and General IP-Phone on-line documentation. The *Mitel IP-Sets Engineering Guidelines* (under General IP-Phone) also provides information on the DHCP options that can be used to assist with network and registration configuration. For Teleworker phones, registration information can also be provided through the RCS server.

The SIP DECT 112 phones have limited capacity for FQDN length, such that the SIP username is limited to 32 characters, and the registrar to 64 characters. The SIP DECT 112 can resolve DNS+SRV records for resiliency, in the form of "_sip._udp.fqdn". "_sip._tcp.fqdn" and "_sip._tls.fqdn".The MBG, if used as DNS server, is not able to resolve these record requests, so an alternative service that can handle SRV records is required.

MiVoice Business Call Control

When deploying the MiVoice Business a number of limits for consideration for UC deployments include, but not limited to:

- Performance (Call handling)
- Number of Registered devices
- Application attachment
- EMEM channels and number of mailboxes, when used.

When each of these factors is considered, along with the number and type of devices per user, the number of MiVoice Business controllers that are required for the solution can be determined.

Engineering Guidelines for the MiVoice Business (MiVB) call control can be found in Mitel Doc Center. The resource settings for the MIVB for VMware can be found in the *Virtualization Deployment Guidelines*, also in Mitel Doc Center.

MiVoice Business Building Blocks:

The building blocks for MiVoice Business fall into two main categories:

- Virtual (and hosted) Deployments: MiVoice Business Virtual 250, 1500, 2500, or 5000 (see below)
- Appliance gateways for on-premises survivable gateway deployments: 3300 Appliances and EX Controller

A configuration option exists to expand the MiVB2500 virtual up to MiVB5000. Although this configuration can offer more capacity in certain areas, such as increased media streaming, this is not realized for typical UC deployments, where a single user has multiple devices. 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). Therefore, for most deployments, the MiVB2500 is the main building block for larger UC deployments.

MiCollab

The MiCollab platform is a collection of UC applications including, but not limited to:

- UM and Advanced UM using the NuPoint application
- Audio, Web and Video conferencing for meeting collaboration
- Presence information through MiCollab Client
- MiCollab softphone integration (for smart phones and PCs), including WebRTC/WebRTC Pro Softphone
- User and licenses provisioning (ULM)
- Initial configuration of users for CloudLink Chat and CloudLink MiTeam Meetings
- MiVoice Border Gateway for license provisioning. (This MBG may be used for smaller not-resilient configurations, but the standalone MBG is scaled for larger and resilient deployments)

The MiCollab application can scale up to 5000 users based on 2.75 devices per user, as a single deployment. A single MiCollab can support up to 7500 users with a reduced number of devices per user. If you require more than 5000 users, please contact your Mitel Account Team for further details.

- MiCollab may be associated with multiple MiVoice Business units
- Multiple MiVoice Border Gateways can be associated with a single MiCollab unit.
- A single MiVoice Border Gateway cannot be associated with multiple MiCollab units, only to a single MiCollab unit

The resource profiles for different MiCollab deployments can be found in the Virtual Application Deployment Guidelines, available on Mitel Doc Center.

Large deployments of MiCollab (2500 users and 5000 users) may experience slowness in registration or web access following a system outage or under heavy traffic conditions, e.g., contact centre usage. Where this response may become noticeable or troublesome to users, the MiCollab operational resource may need to be increased. Additional information on resource updates is provided in the Virtual Application Deployment Guidelines, available in Mitel Doc Center.

Registration location of the MiCollab softphones is also important as it impacts other applications in the solution. The softphones may be deployed as Mobile devices or PC devices. Since these devices are mobile, and may be connected on external networks, the recommendation is to deploy these as external Teleworker devices, even if on the local internal office network.

The PC softphone includes an option to register with either a Teleworker profile, or not. Where the PC softphone is registered as an external device, or where call recording is required, the Teleworker profile should be used. For internally connected PC softphone without call recording, either profile may be used:

PC Softphone setting	Not Call recorded	Call Recorded
Internal Network	Internal or Teleworker	Teleworker
External Network	Teleworker	Teleworker

MiCollab Client also supports a web-based softphone using WebRTC capability. This requires settings on MBG to be configured for support and also for web-proxy forwarding. The MBG should be set to “WebRTC Pro” to enable call recording, and other contact centre functions in conjunction with MICC-B Web Ignite. Note that MiCollab AWV utilizes the “WebRTC” setting on MBG and cannot be recorded. These two settings are mutually exclusive and may require the deployment of multiple MBG to achieve both in tandem.

Engineering Guidelines for the MiCollab can be found in Mitel Document Center.

MiCollab Advanced Messaging (MiCAM)

MiCollab Advanced Messaging brings next generation voice applications to the MiCloud Flex communications environment with Unified Messaging, Transcription, Speech-enabled Directory and Automated Attendant, Secure Voicemail and more.

MiCollab Advanced Messaging/Unified Messaging is Windows based Unified Messaging service. Users can manage their messages from a workstation via a web-browser or e-mail application or can manage messages with a telephone call.

The ability to manage multiple services from one client is known as Unified Messaging. This is offered to users that use their E-Mail client, and store messages on their E-Mail server.

There are two main licensing levels and configurations:

- Basic: The messages are stored locally with the system server and copies are provided to the E-Mail server
- Advanced: The messages are stored on the E-Mail server, which may be local, or Cloud based. Full synchronization of the files and status are maintained with the E-Mail server and the application.

In addition to simply accessing voice files, the use of the E-Mail client also allows messages to be managed through the client including notifications, playback on local workstation or on a phone, or even to forward a message.

The MiCAM solution is an alternative to MiCollab Unified Messaging (a.k.a. NuPoint UM).

The MiCollab Advanced Messaging consists of two main types of servers:

- System Server: Central server, which manages the “message store” and administration tasks
- Call Server: Servers where calls are actually processed and connected to the VoIP call control There may be multiple Call Servers to one System Server

The number of call servers required is determined by the required number of media ports, or channels. The number of media ports is in turn determined by the number of users in the solution and whether a resilient configuration is required. See the following section on Architecture and System Scaling for further details. The call servers register as SIP devices to the MiVoice Business Call Control and require SIP licenses on the MiVB.

Basic Configuration:

The following diagram highlights the different components and configurations for the Basic Configuration:

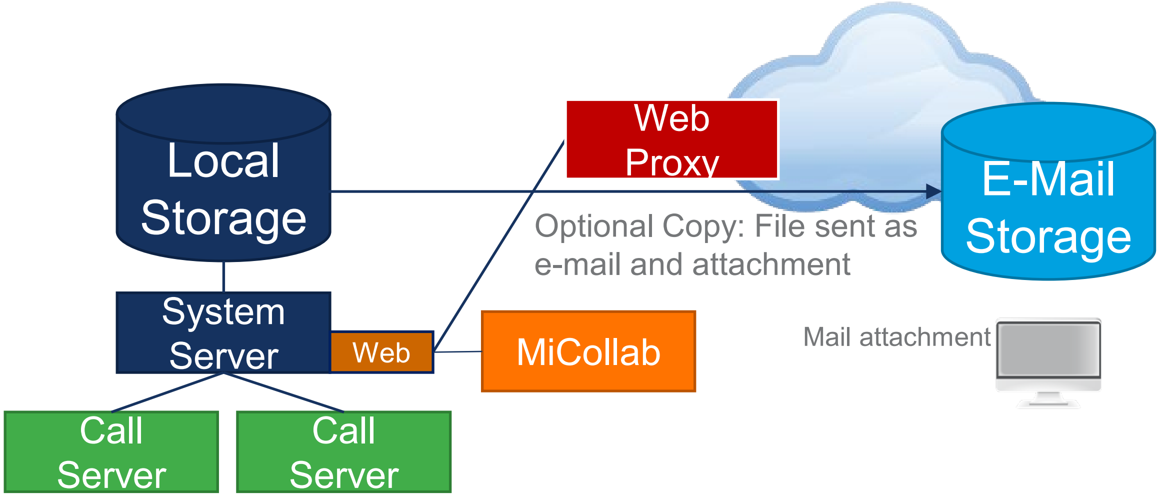

The MiCollab Advanced Messaging functionality is provided by the System Server and a number of Call Servers. The scaling of these is described in a following section.

The System Server provides a local Web Server capability. Users can access information on the messages and notifications via a web-client. The MiCollab Client (associated with Mitel MiCollab) visual voice mail can also access MICAM through this capability. Users who are Internet connected (OTT) will require a correctly configured Web-Proxy. Initial deployments may not offer this external capability

The messages are stored locally with the System Server. Optionally, copies of these messages can be forwarded to an E-Mail service, where the user can access these as an E-Mail and attachment. However, there is no data nor status synchronization with this mechanism. User may find that they have two messages and will need to delete two messages to clear them. There is no integration with the E-Mail client.

Advanced Configuration

The following diagram highlights the different components and configurations for the Advanced Configuration:

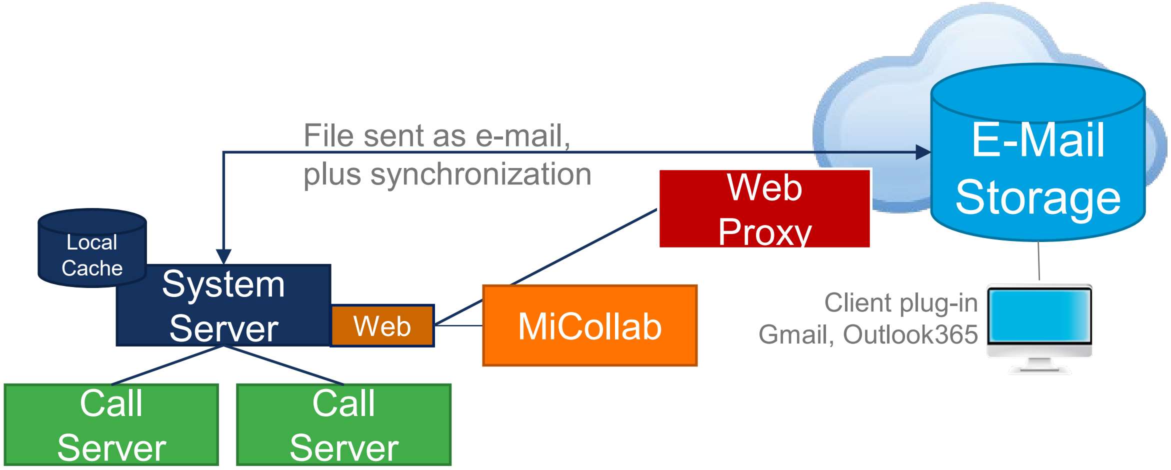

The MiCollab Advanced Messaging functionality is provided by the System Server and a number of Call Servers. The scaling of these is described in a following section.

The System Server provides a local Web Server capability. Users can access information on the messages and notifications via a web client. The MiCollab Client (associated with Mitel MiCollab) visual voice mail can also access MICAM through this capability. Users who are Internet connected (OTT) will require a correctly configured Web-Proxy. Initial deployments may not offer this external capability.

The messages are stored on the E-Mail server, with a local cache on the System Server provided for immediate access of more recent messages. Messages and message status are fully synchronized between the System Server and the E-Mail service. See MiCAM documentation on Mitel Doc Center for details in integration with different E-Mail platforms.

Users have the option to use web access, E-Mail client or phone service to access their messages.

Server Architecture and Scaling

There are two main functions that drive the scaling of the servers and types of servers in the deployment. The main server functions are:

- System Server: This is the core server with primary access and control of the application The System Server also maintains the licensing, database and base language files for ASR/TTS used on the Call Servers. It either stores or caches the recordings locally as well as communicating with an external E-Mail server. One System Server is deployed at the primary data centre.
- Call Server: The Call Servers provide the media handling services for the solution where messages are recorded as a file for transfer to the System Server for storage. There may be multiple Call Servers to a single System Server.

The configuration consists of one primary System Server with a number of attached Call Servers. Multiple Call Servers may be attached for both scaling and resilient operation in both primary and secondary data centres.

The different MICAM Server options are defined as:

- Call Server Type B
- Call Server Type C
- System Server Type B
- System Server Type C

The System Servers include additional HDD for file storage or caching. The Server Type C is a larger server compared to Server Type B and provides additional performance capacity.

In addition, the Call Servers may also provide Text-to-Speech and Advanced Speech Recognition Services. These features are licensed via the System Server. The Call Servers are scaled (Release 9.2 and higher) to handle deployment with and without ASR/TTS enabled.

The System Server is limited in the number of remote Call Servers that it can handle:

- System Server Type B remote server connection limit: 8 Servers
- System Server Type C remote server connection limit: 20 Servers

The Call Servers provide the following number of ports, or channels (with and without ASR/TTS)

Server Type	Available Quantity of Ports
Server Type B	48
Server Type C	96

A ratio of ~50:1 users to MiCAM ports is assumed in defining the number of users to Call Server quantities and definitions. Further details are in the following section.

Server to User Ratio

The following table outlines the server types to the number of users, including when ASR/TTS are included in the deployment:

Users	Call Server (Non-Resilient)	Call Server (Resilient)	System Server
2500	Type B x 1	Type B x 2	Type B x 1
5000	Type B x 2	Type B x 4	Type B x 1
7500	Type B x 3	Type B x 6	Type B x 1
10000	Type B x 4	Type B x 8	Type B x 1
12500	Type C x 3	Type C x 6	Type C x 1
15000	Type C x 3	Type C x 6	Type C x 1
17500	Type C x 4	Type C x 8	Type C x 1
20000	Type C x 4	Type C x 8	Type C x 1

Where ports, or additional ports, are defined instead of users, then these ports are multiplied by the user to port ratio to derive an effective quantity of users. Servers are then allocated based on this effective quantity of users value.

Advanced Speech Recognition and Text-to-Speech

The Advanced Speech Recognition (ASR) and Text-to-Speech (TTS) features are optional to the solution. Some deployment factors should be considered with ASR and TTS features:

- The number of ports required for ASR and TTS are defined separately from the number of MiCAM channels. These are an overlay service on the existing MiCAM ports and cannot exceed this number.
- The ASR and TTS voice files are stored and purchased together
- 1 voice file (ASR + TTS) requires 1G of RAM for operation
- The Mitel Server definitions include sufficient RAM to cater for the addition of ASR and TTS voices and operation
- Mitel CPQ will limit to a total of 4 additional voices (for a total of 5). 1 voice file is included.
- Call Server software should be at version Released 9.2 or higher

When the quantity of ASR/TTS ports is less than the quantity of MiCAM ports, ensure that the separate hunt groups are provisioned in the call server(s) to ensure that these valuable resources are not blocked due to base level MiCAM functions.

License Considerations

The number of licenses required is calculated within the CPQ sales tool. Factors to consider include:

- Server infrastructure and scaling is based on the quantity of (effective) users.

- Additional MiCAM ports can be purchased in addition to the number of ports calculated from the quantity of real users
- The quantity of effective users is derived from the quantity of real users plus the quantity of users associated with any additional MiCAM ports. Typically, the ratio of users per port is ~50.
- Any additional MiCAM ports should consider if the deployment is resilient, or not. The sales tool will double the ports for a resilient system.
- The number of ASR/TTS ports cannot exceed the number of MiCAM ports. ASR/TTS ports are purchased individually
- There is an upper limit of 20,000 effective users on the system
- MiCAM port licenses are bundled in packs of 32 and include 1 Call Server License. Additional port bundles can be purchased to increase the overall port count and also Call Server Licenses. For small deployments, typically less than 32 ports in total, ensure that sufficient port bundles are purchased to include sufficient Call Server Licenses, especially in a resilient configuration.
- For subscription licenses, the minimum number of real users on a system is 100, for a minimum contract period of 1 year.

Resiliency

The Call Servers handle the calls as directed from the MiVoice Business call control. These units will cache any recorded data until access is available to the System Server for storage. The Call Servers are therefore deployed in a resilient deployment, in the same manner as the MiVoice Business units are deployed between a primary and secondary data center.

The System Server is only deployed in the primary data center. The Call Servers are duplicated in both the primary and secondary data centers. This doubles the solution port count, the number of Call Server connections and results in a doubling of licenses, both for MICAM and also ASR/TTS, when enabled.

Messages being recorded are cached locally until complete. If the Call Servers are unable to reach the System Server, the message will remain cached locally. Once connection is reestablished with the System Server, the appropriate files are transferred for storage. During resilient operation, it may not be possible to retrieve or playback messages via the Call Servers as the System Server and messages may be unavailable. In the case of the Advanced configuration, message access may still be available through access to the E-Mail server.

All ports on both primary and secondary servers are fully licensed and accessible. For resilient operation, ensure that the secondary call servers include sufficient ports for MiCAM and ASR/TTS for this backup operation. During normal network operations and peak loads, it may be possible to offset any load through server sharing across both primary and secondary servers.

Use of Neverfail with the System Server is not supported. Local survivability is provided through use of High Availability and inter-data center Disaster Recovery as a Service (via Zerto).

Disaster Recovery (Zerto)

The Call Servers can already be deployed in a Resilient configuration and so are not included into the Disaster Recovery group. However, the System Server is not redundant, and therefore should be include into the Zerto Disaster Recovery as a Service Group.

Server Resource Definitions

The Mitel resource definitions for the Servers, when deployed as a cloud solution, may differ from that provided by the vendor. This is to account for the following:

- Local message storage, or cache, through use of increased HDD
- Additional RAM in order to support ASR/TTS functions and language support

Server Definition	vCPU	RAM (GBytes)	HDD App/OS (GBytes)
Call Server Type B	6	16	160
Call Server Type C	12	24	160
System Server Type B	6	16	500
System Server Type C	12	24	500

The definitions in the table above outline the resource requirements for VMware. Other platforms should look at these values as a minimum requirement and select server settings accordingly.

The deployment should partition the main functions such:

- Call Server (80G for OS and ASR/TTS + 80G for the MICAM Application and local cache)
- System Server (80G for OS and ASR/TTS + 420G for the MiCAM Application and Storage, of which the Application takes 80G)

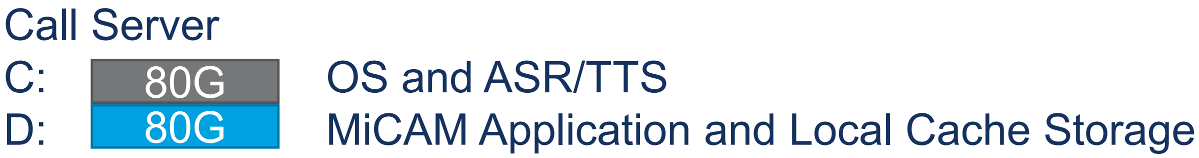

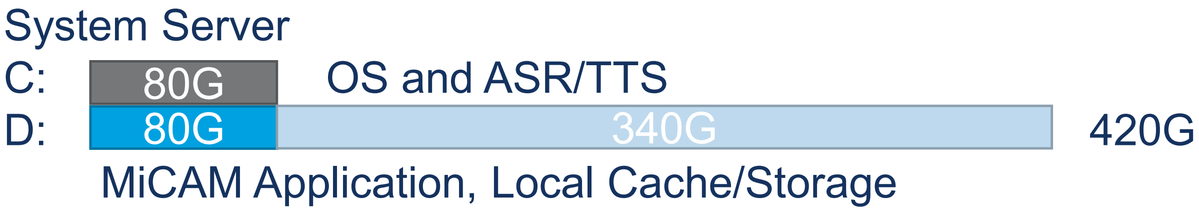

Both the MiCAM Application and Storage occupy the same drive partition. Available message storage within the partition is reduced by the space required for the application, and this results in a message storage capacity of around 340GBytes on the System Server.

Even though the System Server is not directly involved in the ASR and TTS actions, the necessary application, voice files and licensing are all managed from this server, and therefore storage space is required.

Message Storage Requirements (HDD Storage)

Message storing, or local caching, is required for the following servers, and needs to be configured:

- For the Call Server local cache when the System Server is inaccessible
- For the System Server message storage when this is the primary storage location
- For the System Server when messages are cached locally for faster access, even when the primary storage is with the E-Mail server.

For the local cache service, the recommended minimum storage capacity is 200MBytes.

Call Server Caching

The Call server is either Type B or Type C with either 2500 or 5000 users respectively. The local cache is needed to buffer messages that the user might receive but are unable to be transferred to the System Server due to inability to access this unit.

Based on an extended outage period and assuming a high message receipt rate (Office traffic rates), then a minimum local cache of ~5GBytes is recommended for Server Type C and 2.5GBytes for Server Type B. This is easily accommodated with the defined 80G partition.

System Server Storage and Caching

The System Server either Stores the messages locally as primary storage, or it caches the information for faster access as secondary storage. In either case, the same calculations apply as to the required capacity.

Based on a non-compressed CODEC of G.711, and the following assumptions, the table below identifies the minimum storage requirements for the different number of users:

- Message length: 30 seconds
- Stored messages per user at any time: Up to 15 minutes
- 1 year storage

Number of Users	Minimum Storage Capacity (GBytes)
250	2
500	4
1000	8
2500	18
5000	36
10000	72
15000	108
20000	144

Even for the largest system, the minimum storage requirement is covered by the server specifications, and no additional storage is required.

IP Ports and connectivity

User connectivity to the server can be provided via a Web server. This is available both to on-net users with full capability, and to OTT connected users with some reduced functionality. For external connectivity user functions and limited admin functions are provided.

Internal IP Ports

The following IP ports are needed for user connections to the service:

Port	TCP/UDP	Direction	Description
443	TCP	To Web Server from Customer LAN	Web Services

External IP Ports

The following IP ports are needed for user connections to the service via the MBG Web-Proxy service:

Port	TCP/UDP	Direction	Description
443	TCP	To Web Server from Customer LAN	Web Services

Features provided via the external service include:

- User and Admin Login
- Mailbox access, including ability to save, delete, forward and download
- Record greetings
- Out of Office notifications
- Report (admin)
- Auto-Attendant (admin)
- Some Application Settings (admin)

MiVoice Border Gateway

The MiVoice Border Gateway is the “front door” to the solution for any external and Teleworker devices. It provides Session Border Functions (SBC) for user devices, both Mitel proprietary and also SIP devices, as well as gateway functions to SIP Trunk Service Providers.

The primary function of the MiVoice Border Gateway is to terminate Internet user devices, either Mitel IP-Phone (MiNet based or SIP based, and proxy these through an application level gateways (ALG) to internal applications and other phone devices. The MiVoice Border Gateway includes the ability to securely terminate connections and also provide NAT capabilities, allowing users to deploy phones anywhere on the Internet without resorting to dedicated VPN connections.

The MiVoice Border Gateway also provides external terminations for other applications, including MiCC Business and MiCollab.

As a trunk gateway, the MiVoice Border Gateway terminates SIP Trunk calls to/from a Service Provider and can providerudimentary routing to the appropriate call control unit in the solution.

Some deployment limits that need to be considered for UC deployments on MiVoice Border Gateway, include, but not limited to:

- Performance (Call Handling)
- Number of concurrent media and UC connections through the MiVoice Border Gateway
- Number of registered users
- Use as a Secure Recording Connector (SRC) or use in addition to the normal MBG functions
- Number of Web proxy and WebRTC active connections
- Settings for WebRTC/WebRTC Pro
- Number of DID/DDI routing rules, when applies to a SIP trunk gateway

For a hosted deployment, this results in the two main MiVoice Border Gateway building blocks:

- MBG250
- MBG2500

For smaller deployments, the MBG250 can handle up to 125 connections. It can also be used as a single gateway for both users and SIP trunks. This is roughly equivalent to 50 teleworker connections to 50 trunk connections on the same external interface. For larger configurations use the MBG2500 configuration.

The MiVoice Border Gateway that is included with MiCollab, when deployed with a smaller single data centre deployment, has the same performance characteristics as the MBG250. This is a special deployment. Normally this integrated MiVoice Border Gateway is used to synchronize database updates and licensing for the external MiVoice Border and not deployed on the network edge.

MBG2500 can handle 2500 UC devices and 500 streams at nominal traffic levels. However, where the traffic per device is lower, such as when there are multiple devices per user, this MBG2500 can handle up to 5000 devices. Ensure that the appropriate server resources are applied. When this MBG is used as an SRC, remember to consider the additional SRC load, and in a resilient recording solution, to also add in this additional resilient component.

There are other configurations for MBG which provide different levels of scaling and functionality, but these two MBGs are the main building blocks that are deployed.

Remember that MiCollab softphones will register through the MiVoice Border Gateway as a device, even when these units are active on the business LAN. These devices are treated as Teleworker devices, to ensure that they can roam between multiple IP networks, and not just the local business network.

The MiVoice Border Gateway is the front door to the solution, for OTT connections. Pay particular attention to the IP Ports that are being used, and also those that will not be used, and ensure that any external firewall is configured appropriately. Misconfiguration of firewall IP Ports is often a source of one-way audio and other failed functionality.

The MiVoice Border Gateway can also be deployed as a standalone unit on dedicated hardware in a customer premises with a survivable gateway solution. Ensure that this unit is clustered with the appropriate teleworker or SIP Trunk MiVoice Border Gateway in the hosted data centre to ensure correct resilient operation and license sharing. This cluster connection should follow the cluster limits of the *MBG Engineering Guidelines*. In the event that this limit is exceeded please contact your Mitel Account Team.

See *MBG Engineering Guidelines* for further details.

MiVoice Border Gateway as Secure Recording Connector

The MiVoice Border Gateway can act as a Secure Recording Connector (SRC) for call recording applications. This allows secured connections from end devices, or trunks, to provide additional separate and secured recording streams to call recording equipment. These additional streams have a performance impact and impact the number of total streams that can be handled via a single MiVoice Border Gateway unit. Further details can be found in the *MBG Engineering Guidelines* in Mitel Doc Center.

In a parallel recording, or resilient recording, configuration this may result in multiple recording streams and these need to be considered accordingly, i.e., if there are 2 Call Recording Equipment (primary and secondary), then there are 2 recording streams, and associated additional performance requirements on the gateway.

Secure Recording Connector units can be deployed as dedicated units, such as for internal network connected devices, and also combined with the externalUser/Teleworker MiVoice Border Gateway. For simplification of management, scaling and deployment, the recommended topology is to use the external MiVoice Border Gateway to carry out both functions for Border Gateway andSecure Recording Connector. Additional Secure Recording Connector units are then deployed for internally connected phones.

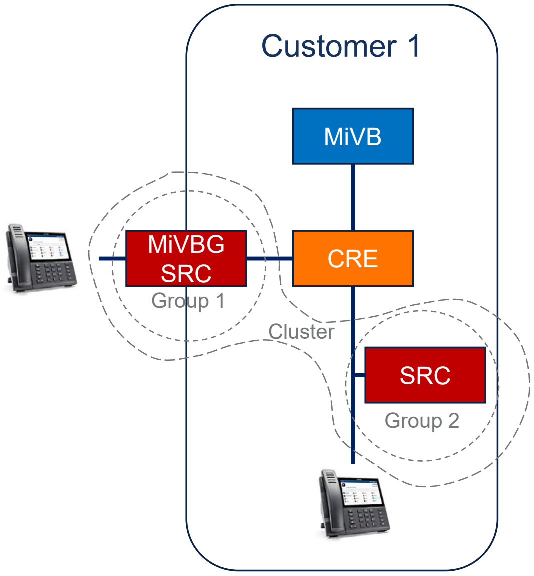

Where the SRC is using on conjunction with the Teleworker, or trunk, gateway, this will typically be deployed in server gateway mode.

For the internal SRC, this needs to be deployed in LAN mode. This allows this unit to provide SRC functions without the requirement for additional teleworker licenses.

When configured with a user gateway recording solution, it is recommended to join the SRC cluster with the user MiVoice Border Gateway cluster in order to simplify management of the licenses and sharing. It may be necessary to apply groups to the different set of units so that Teleworker licenses are only shared with the user MiVoice Border Gateways, and not with the SRC gateways.

See *MiVoice Border Gateway Engineering Guidelines* for further details.

Where the clustering results in more MiVoice Border Gateway units in a cluster than recommended in the *MiVoice Border Gateway Engineering Guidelines*, please contact your Mitel Account Team for further advice.

WebRTC

MiVoice Border Gateway can provide WebRTC gateway capability. With MBG Release 11.4 SP1, the WebRTC capability has been enhanced and provides two operating modes, allowing backwards compatibility. These modes are:

- WebRTC, and
- WebRTC Pro

Note that all WebRTC (Legacy WebRTC and WebRTC Pro) connections are treated as Teleworker connections and require appropriate Teleworker licenses. This includes any WebRTC softphone connection from on-premises private networks. An Internet connection is therefore needed.

Note that the operation of WebRTC settings and WebRTC Pro settings are different:

- WebRTC operates with multiple browsers. It provides anonymous operation and is used by AWV. It provides basic call handling features. It cannot be call recorded
- WebRTC Pro operates with Google Chrome. It requires teleworker registration. It provides feature compatibility with SIP device and can be call recorded.

For MiCloud Flex deployments, the default values as defined in the *MiVoice Border Gateway Engineering Guidelines* apply. At time of release, this limits the number of WebRTC connections, per user-facing MiVoice Border Gateway to the following values:

WebRTC Service	Concurrent Connections
Audio	124
Audio and Video	62

Additional MiVoice Border Gateways may be deployed to increase the capacity. Note, that for UC users, not all users are active at the same time. Therefore, these connections are similar to trunk connections, in that 124 connections may be used by ~600 UC Users.

WebRTC connections cannot pass via an SRC and therefore cannot be part of call recording.

WebRTC clients are not inherently resilient. Resilient operation requires the use of DNS and multiple A-Records. A failed connection will need to be re-established with the alternative gateway. Because DNS will often deliver multiple IP addresses in a round-robin fashion, in a resilient deployment, and during normal operation, the secondary MBG may also be added to the total supported connections, i.e., double capacity. The number of connections during resilient operation should be considered, if the number of supported end users is critical, i.e., halved.

MiCollab AWV uses the WebRTC setting. Note that WebRTC and WebRTC settings are mutually exclusive. If both WebRTC and WebRTC connections are required, multiple MBG will need to be deployed,

WebRTC Pro

When this mode is selected in the MiVoice Border Gateway, the WebRTC browser connections are treated as SIP devices, and the SIP phone limitations apply to the server. These are primarily audio-only connections.

Since the connections are treated as SIP Phones, they can participate in a number of call features, such as:

- Hot-desk

- Call Recording
- Local and External Resiliency
- Participate in a number of Contact Center agent features

This functionality is limited to Chrome browsers only. Although other browsers may appear to initially work, the operation cannot be guaranteed.

When the **WebRTC Pro** option is enabled on an MBG, the ability to support WebRTC **Anonymous Mode** on the same MBG is no longer available. Enabling this option impacts certain features of MiCollab AWV which use anonymous WebRTC. Specifically, the AWV web client option to **Join Audio** from the PC is not supported when the MBG is in WebRTC Pro mode. To connect audio via the PC the user must call into the AWV bridge from the MiCollab WebRTC softphone or make the call from a dedicated desk phone. Alternatively, deploy multiple MBG with different WebRTC settings.

Web Proxy and Port Forwarding

The MiVoice Border Gateway has a limit on the number of web-proxy connections that can be handled. A larger MiVoice Border Gateway is available for situations where a large number of web-proxy connections are needed. Scaling can be achieved by additional MiVoice Border Gateways, as needed, and different DNS references for the different applications.

Currently supported applications include

- MiCollab (Client and AWV)
- MiContact Center Business
- MiContact Center Outbound
- Mitel Interaction Recording
- MICAM
- MIR
- OIG

For deployments with high number of users and web-proxy connections, use of the MiVoice Border Gateway with extended web-proxy support is recommended. Post installation configuration may be required to enable the number of web-proxy connections up to 5000 from the base 500 quantity. See MBG Installation and Maintenance Guide for further details. Additional server resources are required to support this extended functionality. See the Server resource definitions in the Virtual Application Deployment Guidelines on Mitel Doc Center.

For smaller deployments, it may be simpler to combine all port forwarding through a common MiVoice Border Gateway. However, for larger deployments with multiple MiVoice Border Gateways, it may be more appropriate to designate certain functions (and URLs) to certain gateways, for example, separation of contact centre users and functions versus back-office teleworker users.

In addition to web-proxy connections, some applications also require dedicated port forwarding. These will be reference in the applicable application documentation.

Further details are available in the MBG Engineering Guidelines in Mitel Doc Center.

External IP Ports

The MiVoice Border Gateway is the “front door” for external devices and Teleworker devices, plus a number of UC applications. Each of these require access to specific server ports in order to be handled correctly. Take note of the IP Port information that is provided in the MBG Engineering Guidelines in Mitel Doc Center.

Mitel Open Integration Gateway (OIG)

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 an OIG 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:

- Mitel Open Integration Gateway must be deployed per customer, or two per customer if resiliency is required without using High Availability
- 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.

Details on server resources are included in the Virtual Application Deployment Guidelines, available on Mitel Doc Center

Web-Proxy

Web-Proxy capability is available via MBG to allow external user access and also third-party application access. Access is via IP Port 443.

This allows the following services to be accessed:

- Third party application access
- Salesforce and Google plugins
- API controls

- Server management and configuration

CloudLink Chat

CloudLink chat supersedes MiCollab chat, once configured. Configuration is via MiCollab.

This provides a unified mechanism to share chat to all users whether on premises or mobile user, or Teleworker. The service is provided in a public cloud (AWS) and is accessible from anywhere in the Globe that provides Internet access.

Chat is available via any current browser and via a dedicated client for PC/MAC, or mobile phone (iOSor Android).

CloudLink MiTeam Meetings

CloudLink MiTeam Meetings is an upgraded Unified Communication and Collaboration package from MiCollab AWV. Access to CloudLink MiTeam Meetings is configured via MiCollab.

This provides a unified collaboration and video/audio conferencing facility available to all users whether on premises or mobile user, or Teleworker. The service is provided in a public cloud (AWS) and is accessible from anywhere in the Globe that provides Internet access.

CloudLink MiTeam Meetings is available via any current browser that supports WebRTC and also via a dedicated client for PC/MAC, or mobile phone (iOS or Android).

Note that while MiTeam Meetings may be recorded within that application, this is totally independent from the recording solution provided by Mitel Interaction Recording and use with Secure Recording Connector.

Mitel Performance Analytics (MPA)

Mitel Performance Analytics is a fault and performance management system designed to provide users with fast actionable problem resolution so that optimal service quality levels are maintained for end customers. Mitel Performance Analytics provides real-time alerts, detailed reporting and ubiquitous accessibility with secure remote access.

The Mitel Performance Analytics consists of two components:

- MPA Cloud Service, hosted on AWS Cloud. Access to the service and to the customer probe is via a user portal to this service
- MPA Probe, or client This is a local agent that runs within the customer deployment and reports back to the cloud service.

Mitel Performance Analytics (MPA) provides access to a number of capabilities, including:

- Gathering and reporting of performance related information
- Management access from the Cloud service
- Backup Management

Performance related information includes a number of metrics including, but not limited to, voice quality, SIP Trunk utilization, up-time, system performance and availability, etc. These metrics provide an indication of infrastructure usage and stability. Metrics and thresholds can also trigger alarm notifications to identify when a system may be overloaded and require upgrade, or where there is unexpected behavior.

Management access is provided from the MPA Cloud Service through use of a virtual tunnel. The MPA Probe, part of the customer specific Flex on Google Cloud installation, connects back to the MPA Cloud service, and thereby establishes a secured connection. Logging into the MPA Cloud service provides access to the tunnel. Access to the MPA Cloud Service is secured through multi-factor authentication.

A Mitel Performance Analytics (MPA) management probe is typically deployed in both primary and secondary regional locations, and both can be accessed from the MPA Cloud portal. In the event connection is lost to one of the probes, access will still be available to the secondary unit. The MPA Cloud service is itself deployed on the highly available AWS Cloud service and does not impact and call handling nor voice streaming. An additional probe may also be deployed for on-premises deployments.

Additional information can be found under Mitel Performance Analytics Engineering Guidelines under Mitel Doc Center.

The suggested virtual machine resource profile is:

Product	vCPU	RAM	HDD
MPA Probe	2	4GBytes	50GBytes

This provides sufficient resource for information gathering from up to 50 units. A unit could be any application that can provide SNMP information as well as any network equipment that can equally provide similar SNMP information.

MiContact Center Business

MiContact Center Business (MiCC-B) is an enterprise grade, multi-channel and cloud-ready contact centre solution that works with the MiVoice Business platform.

For smaller contact centre deployments, the MiCC-B already includes an internal IVR capability. However, this has limited capacity to 40 ports. Anything requiring more than 40 ports should use external (aka Remote) IVR servers. A separate IVR is also needed for resilient operation.

In the subscription model, a number of IVR base licenses are included with the base MiCC-, and this should be sufficient to deal with a deployment up to 1000 agents.

The IVR servers are limited to 120 active ports. Two pairs of servers, one in the primary data centre and one in the secondary data center are needed for a resilient deployment. For example, if 480 IVR ports are needed, this requires 4 IVR servers primary, and 4 IVR servers secondary, for a resilient deployment. MiCC-Business is limited to 480 active IVR ports.

Agent connection to the web service is available directly for on-net agents, and via the MiVoice Border Gateway web-proxy for external agents.

During the handling of incoming calls, callers may be placed on hold in a queue, awaiting an available agent. Often callers are provided with Music-in-Queue, which is typically provided by the MiVoice Business or multiple MiVoice Business units handling the queue(s). Each caller is connected to a media channel that provides this music. Consider the total number of callers that are likely to be in queue at any one time and scale the MiVoice Business for queuing to ensure that there are adequate number of media channels available across the queuing MiVoice Business units. Further details on the number of media channels available for a particular MiVoice Business configuration are available in the *MiVB Engineering Guidelines*. Worst case scenario would be where all trunks are utilized with incoming callers and there are no available agents.

Agent connection to the web service is available directly for on-net agents, and via the MiVoice Border Gateway web-proxy for external agents.

For information about scaling MiContact Center Business, see the *MiContact Center Business Engineering Guidelines* and also *MiContact Center Business Blue Print Guide.*

MiContact Center Business can support 1000 concurrent agents, dependent on traffic requirements. For larger deployments, contact your Mitel Account Team.

MiContact Center includes a Workforce Scheduler (WFS) application. When deployed there are requirements on the resource profile to be used. Further details can be found in the MiCC Engineering Guidelines. In addition, an external SQL server must be provided to use this application.

- WFS licenses are named licenses with one license per agent
- WFS is deployed as part of the MiCC Business installation
- Mandatory 20 hours of professional services for training and consultative services are required per installation.

MiContact Center includes functionality for Multi-Media and also integration with the WFM Client. When either, or both, of these optional components are required, the MiCC Large resource profile for MiContact Center should be used. See the Virtual Application Deployment guidelines for virtual machine resource profiles.

MiContact Center Speech (Advanced Speech Recognition and Text to Speech)

The MiContact Center Speech services include voice functionality that supplements the existing functions of the MiContact Center Business IVR units. These provide additional functions for:

- Advanced Speech Recognition

- Standard Text to Speech
- Enhanced Text to Speech

Advanced Speech Recognition (ASR) may also be described as “IVR Speech Recognizer” in some documentation.

Text to Speech voices (regional variances on a language, such as UK English versus North American English) can be used as Standard and Enhanced versions. The Enhanced versions provide a more natural speech presentation, but also require additional processing.

The deployment architecture is highlighted in the diagram below:

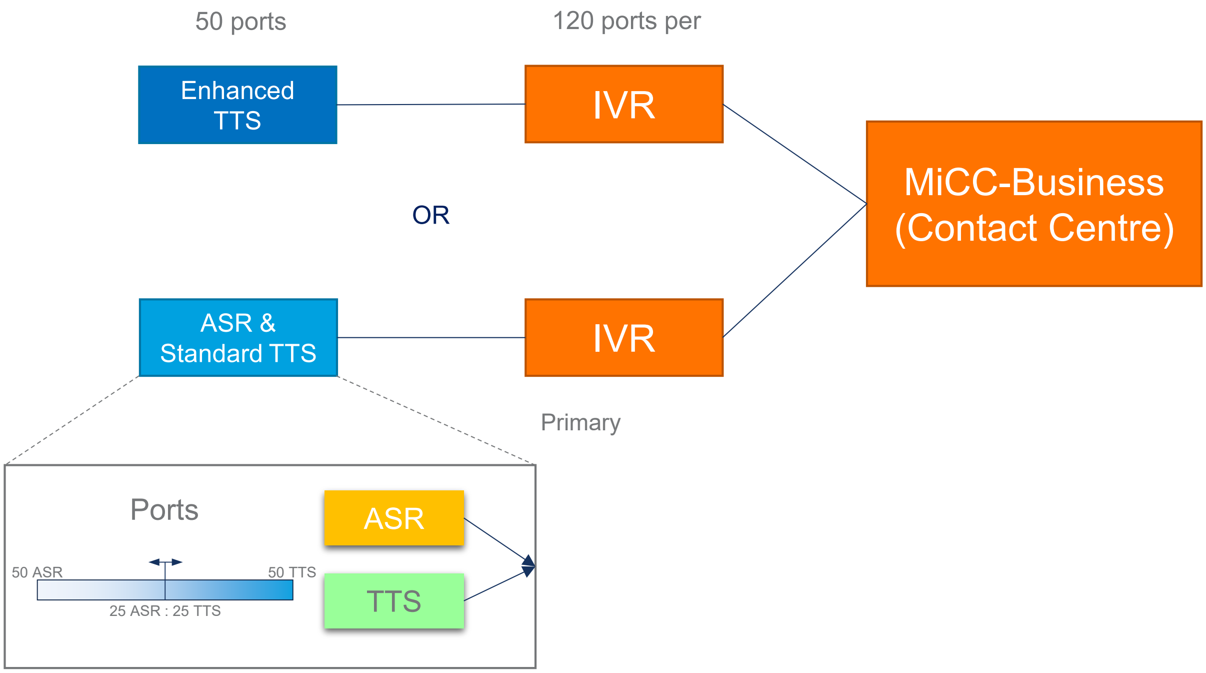

There are two server variations that can be deployed:

- Advanced Speech Recognition (ASR) and Standard Text to Speech (TTS) up to 50 ports total

OR

- Enhanced Text to Speech (TTS) up to 50 ports total

The combined ASR and Standard TTS server is limited by the MiContact Centre to a total of 50 ports. It is possible to select up to 50 ports of ASR, or 50 ports of TTS, or a combination of both, not exceeding a total of 50 ports.

Typically, ASR and TTS are deployed in a ratio to each other, and a suggested ratio is 4 ASR ports to 1 TTS port. However, it is possible to deploy with other ratios, such as 1:1.

The Enhanced TTS server only supports Enhanced TTS and cannot be combined with the other services.

The MiContact Center Business can only support 1 server connection, which results in the following deployment considerations:

- ASR and Standard TTS **OR** Enhanced TTS, not both
- This is a non-resilient solution

MiContact Centre Business queuing groups need to be defined for the different functions to ensure the calls are routed to the appropriate service and may be combined with more traditional IVR functionality. If resilient operation brings into service the secondary IVR units, the queueing should take into consideration that AST/TTS may not be accessible.

The AST/TTS servers can be deployed with the Zerto disaster recovery group, along with the MiCC-B, to ensure that the unit can be replicated to cover disaster recovery.

Deployment of the server requires that a static or Predefined MAC address is assigned to the primary interface (eth0) of the server. This can be managed through the vCenter application, once the virtual machine is deployed. This information is carried with the server through Vmotion and HA activities.

Should there be a need to disaster recover the server, the same static MAC address should be applied, or the server ‘re-hosted’ from within the application. There is a minimum time limit on the number of times the ‘re-hosting’ can be applied.

Voice Support

The servers are scaled to support up to 4 different voices. A voice is a variation on a language, so one language may require multiple voices, for example, UK English versus North American English would be two separate voices.

If more voices need to be supported, additional servers with different pre-loaded voices can be used. Ensure that routing is included in MiContact Center Business to route to the appropriate service and voice support.

Server Resources

The following resources are required for the different servers:

Server Description	Port Limit	vCPU	RAM	HDD
ASR and Standard TTS 60	60	2	12	780
ASR and Standard TTS 120	120	4	12	1240
Enhanced TTS 60	60	6	32	780
Enhanced TTS 120	120	12	64	780

Resilient Configuration

The IVR units can be deployed as resilient units in primary and secondary locations. In order to use the ASR and TTS services in this mode, additional servers are required to line up with the IVR units, i.e., duplication.

Deployment Scaling

Should the deployment start to get complex in terms of number of voices supported and where the number of ports per voice is not equally distributed, then please contact your Mitel Account Team who can calculate the number of servers required. The CPQ sales tool will provide an estimate of servers and types for most balanced system. Those with multiple voices in excess of 8 voices may require more manual calculation for server optimization.

Mitel Workforce Management (WFM)

The Mitel Workforce Management (WFM) product is an optional cloud service offering workforce-management and call handling analysis for contact centres. Customers access this information from the Cloud service. This requires a local client to be deployed with the contact centre to gather information.

The WFM client is available as an integrated option within the MiContact Center Large Server configuration. This integrated configuration will be used for all new deployments.

Where the WFM Client is already deployed as a separate virtual machine, for legacy deployments, the following virtual resource profile is recommended:

Product	vCPU	RAM	HDD
WFM	2	6GBytes	80GBytes

MiVoice Call Recording (MiV-CR)

MiVoice Call Recording is replaced with Mitel Interaction Recording (MIR).See Below.

MiCloud Business Analytics (MBA)

MiCloud Business Analytics is a cloud service which customer can optionally subscribe to. Customers can connect to the Cloud service to access their dashboards and reports on calls being handled.

MiCloud Flex utilizes a common collector node across all customers. Customer to customer data is isolated by appropriate VLAN and NAT techniques as well as multi-tenant isolation in the data collection node.

Please contact your Mitel Account Team for optimum deployment solution should you wish to deploy this feature.

MiCollab Client Deployment and Redirect Server (RS)

The MiCollab Redirect Server (separate from RCS) 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)
- MiCloud Management Portal 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.

Redirect and Configuration Service (RCS) server

The Mitel Redirection and Configuration Service (RCS) is a service that offers touchless deployment, firmware control, and branding of certain Mitel devices. Mitel’s RCS eases the issues Service Providersface with mass deployments. By entering the MAC address of a device into the Global RCS server, upon initial boot-up, the device can be loaded with the proper pre-determined firmware version, and then routed to its assigned server for configuration.

The RCS simplifies the Service Provider deployment of Teleworker phones for all topologies. It is usefulfor the Small Business topology, where all IP-Phones are deployed as Teleworker.

The use of RCS allows the service provider to identify the phones through their MAC access and theend-customer that these are delivered to. Typically for a Teleworker phone, a FQDN or gateway IP- Address must be programmed into the phone to obtain service. This information may be provided either statically in the phone or though defined vendor options in DHCP. This is the normal operationfor on premises deployments. Most Teleworker deployments do not provide such DHCP capabilities, and static programming by end-users may be error prone. If DHCP information or static information is lacking, then the phones are pre-programmed to seek the RCS service to receive updates and redirection information to the appropriate gateway. Use of RCS also eliminates the need to statically program devices by the service provider prior to delivery to the customer.

MiVoice Business Console

This section covers cloud hosted MiVoice Business Console installation and deployment guidelines.

The MiVoice Business Console (MiVB-Console) can be deployed on a customer’s private network and as a Teleworker device.

The two main functional connections and applications that are applicable to the MiVoice Business Console are:

- Call Signaling (MiNet) - to call controllers
- Presence information - to a dedicated MiCollab Client Presence Server

When multiple MiVoice Business Consoles are used in Teleworker mode, they must connect to a single common MiVoice Business Controller while using a common MiVoice Border Gateway.

The diagram below shows the possible connection scenarios for internally and externally (Teleworker) connected MiVoice Business Consoles:

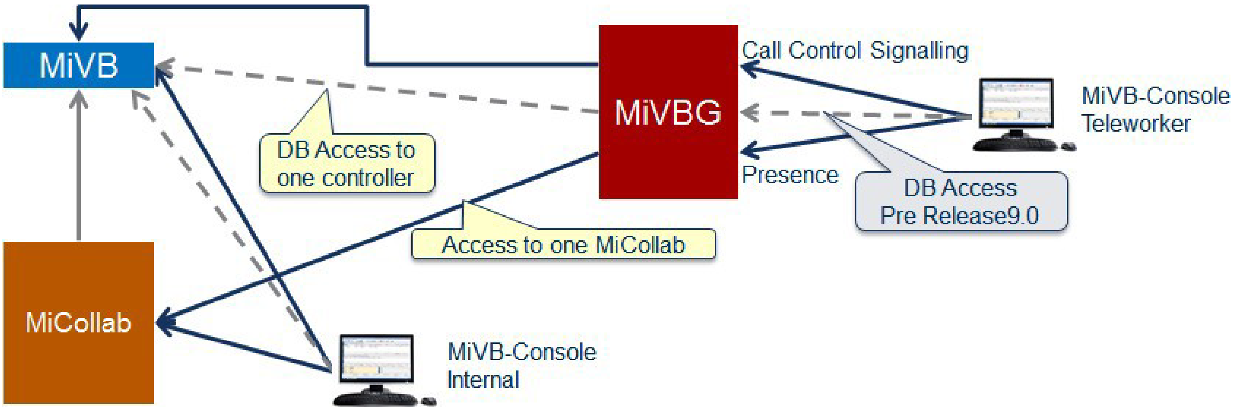

The internal MiVoice Business Console connects directly to the associated MiVoice Business Controller for Call Control Signaling. It connects directly to the associated MiCollab Client Presence Server. In addition, there is an optional connection to the MiVoice Business Controller for database access. The link is only required for backward compatibility to MiVoice Business prior to Release 9.0.

The external MiVoice Business Console connects to the Border Gateway of the customer. All connections go through this gateway, albeit to different access ports. The MiVoice Business Console is connected to the defined MiVoice Business Controller for Call Control Signaling. The connection for Presence is forwarded to the defined MiCollab Client Presence Server. This presence server is also linked to the MiVoice Business that the console registers for Call Handling, in order to maintain synchronization of the presence information for all users.

Avatars in a Cloud Deployment

Use of Avatars (user pictures) with Mitel desk phones was introduced with the 6900 phones. The Avatar files are stored on the MiCollab server and accessed via web services.

For internally connected phones, connection to the MiCollab for Avatar access is direct to the MiCollab server.

For Teleworker phones, connection for Avatars is required via a dedicated port on the MiVoice Border Gateway for forwarding to the defined MiCollab.

The server location of the Avatars is provided to the phones during initial startup through programming of a FQDN within the customer MiVoice Business (MiVB) Call Controller. The FQDN identifies the customer associated MiCollab Server where the Avatars are located. The FQDN must be resolvable within the internal hosted network of the service provider, and at remote external locations for Teleworker phones. This requires use of Split-DNS.

The deployment configuration is highlighted in the diagram below:

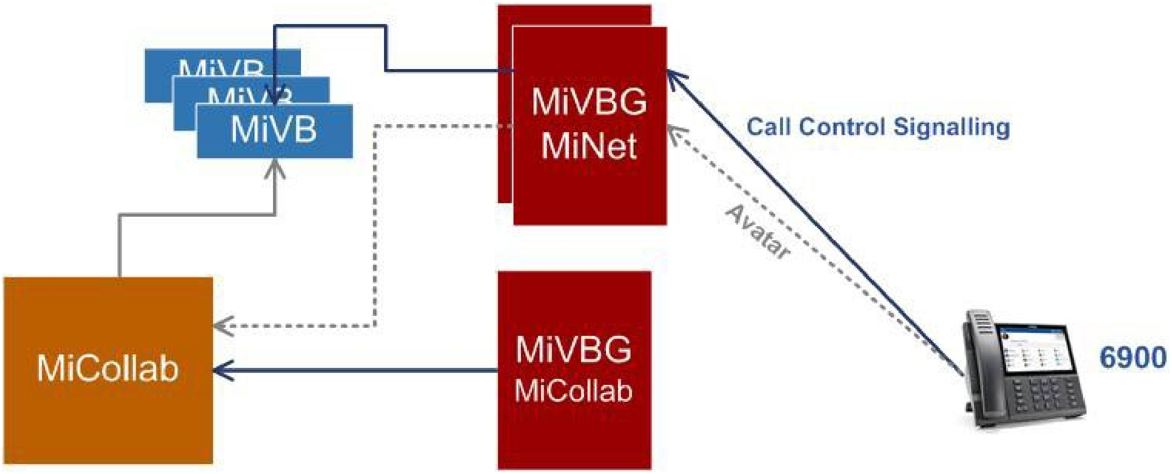

Mitel Interaction Recording (MIR)

Mitel Workforce Optimization (WFO) is a suite of products to improve and integrate call handling and business processes. It is primarily targeted to Contact Centre deployments but may be used outside of thatarena. The Mitel Interaction Recording (MIR) is a suite of products, within WFO, that provide Call Recording and Screen Recording capabilities. Speech analysis of the voice recording can optionally be included and provides an ability to keyword spot or to provide full transcription of the conversations. The sections below will outline the key topologies for Call Recording and Screen Recording. Speech Analytics (Transcription/Keyword) services are covered in a subsequent section.

The services are intended to scale from 1 to 1000 agents, or up to 5000 users. The service can be provided as an all-in-one server (with additional external storage as needed) in a single data centre to multi-server deployments across geo-redundant data centres.

Call and Screen recording may be used for multiple reasons, but typically these are used as part of contact centre deployments to record any discussions or transactions between end-customers and agents. These recordings may be for legal reasons, or recourse. They may also be used as part of training of new agents,and screen recording and playback are often used for this purpose.

During monetary transactions, or for privacy reasons, it may be necessary to be able to stop and start a recording.The provision of this stop/start ability (also known as Recording on Demand (RoD)) has some implications on the network deployment and servers that are required. These are detailed below.

The customer must also consider the retention policy for the information. All recordings require storage, and the longer these are stored, the more storage capacity is needed, along with associated cost. Typically call recordings are retained for 1 year and screen recordings for 3 months. Alternative retention periods are also possible. Unless stated, any calculations in the following sections use these 1 year and 3 months as defaults.

Although agents can be recorded on OTT connections, a number of other functions are only available via direct on-net connections. As a result, the solution is targeted towards customers with private network connections. It is not suitable for a pure OTT delivery model.

Architecture

The following sections look at the architecture and servers needed to deploy the MIR solution.

Key functional components of the solution

There are number of key operational components to the solution. These functions are grouped or identified against particular server configurations, and these servers are then scaled as required. At the lower end of the scaling this could result in a single all-in-one server, up to multi-server configurations for larger and geo-resilient configurations. The core functions of the solution are:

Component Functions	**Description**
Recording Module(RM)	Real time recording application for calls/voice and screen
CTI Connect (CTI)	Connects to the call control to determine call status and when to record specific calls
Replay Server	Allows agents or supervisors to playback recordings, audio or screen
Web / API Server	Web interface to the core and other services including playback
Enterprise Core	This is the main application and is the central hub for all the servers. Management of the solution isvia the unit
Database	Where information on the recordings is maintained.
Network Attached Storage (NAS)	This is a separate function and is the location of the storage media. This may be a separate disk partition with the Enterprise Core or may be a separate network drive or storage server. Typically, this will use different storage, maybe with restricted access, compared to the other server components, being mainly a write and forget media, with minimal playback.
Recording Control(RC)	Management and control of the recording modules. One RC may control multiple RM

The software provided for the solution includes all of these components. However, only certain components are enabled on different server configurations, and it is these servers which form the basic building blocks of the solution. It is important to distinguish which specific components are associated with which server, especially where this may impact functionality when resilient operation is in practice.

Function to Server Mapping

The servers are built into the following configurations:

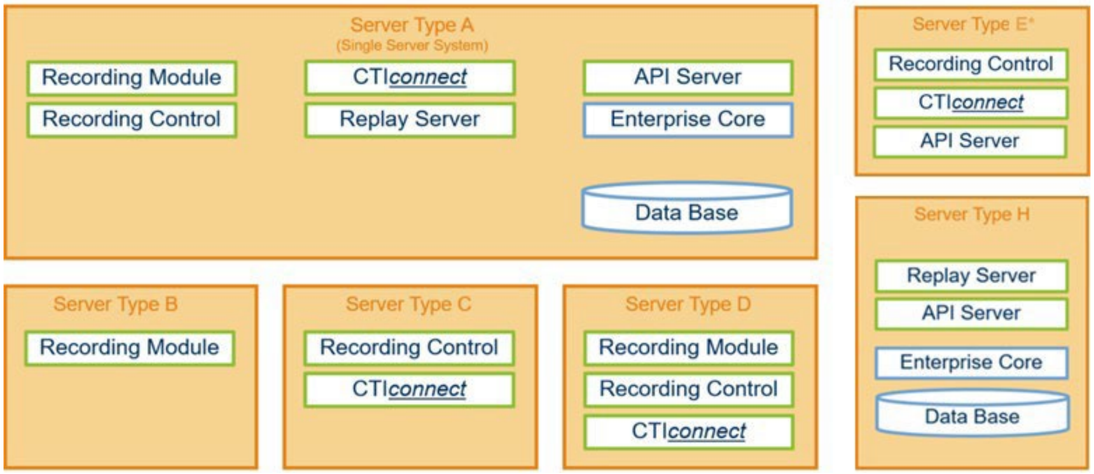

Note*: Server E* is identified with a 'star' to differentiate from the standard Server 'E' This is because server E installation includes several components, some of which are not utilized in this deployment, specifically the CTI component. See the table below.

These server configurations provide the necessary building blocks for the hosted solutions. The table, below, identifies the different component functions to the different server configurations.

Server Name	Description	Recording Control	Recording Module	CTI	Playback	API/Web	Core	Post-Compression	Database
ServerA	All-in-One	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ServerB	Recording Module		Yes						
ServerC	CTI and Recording Control	Yes		Yes					
ServerD	Recording and CTI (B+C)	Yes	Yes	Yes					
ServerE*	CTI, API, Recording Control	Yes		(Yes)		Yes			

ServerH	Core, Playback, API, Database				Yes	Yes	Yes	Yes	Yes
	Realtime			**Call Control**					

Other than the All-in-One Server A, all other servers are either targeted to Realtime or Call Handling functions, not both. This allows the recording components to scale as needed while minimizing the scaling requirement on the other servers. Post-compression of recordings is managed and performed via the Core server (Server A or Server H).

Server E* is an optional server and is deployed when there is a need to provide resilient "Recording on Demand" control via the web interface. This server brings in an additional web server (API) and recording control (Recording Control). The CTI component is not used but is part of this server due to the installation process. Use of primary and secondary web servers also necessitates the need for web service load sharing. More details are provided below. It does not play an active role with Core functions, nor Playback.

The Server E* is a derivation of the ASC Server E. They are not the same. Installation information is provided on how to create this server in the *Deployment Guidelines*. This server is only needed for the secondary backup on the Recording on Demand feature.

Architecture Scaling

Scaling and deployment across different data centers follows the basic server patterns highlighted below.

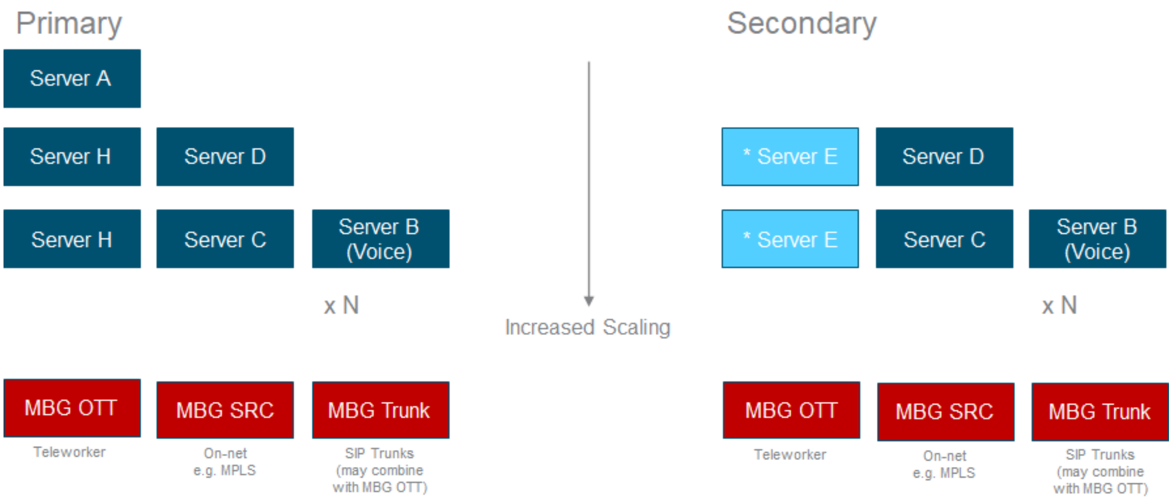

Database is integral to Server A or Server H

*Note: Optional server – needed for resilient Recording on Demand (stop/start recording)

Server E* is an optional server to provide resilient "Recording on Demand" control. This control is needed to start and stop recordings on demand, such as when dealing with a call with PCI information, or a call that requires privacy, for example, to comply with GDPR.

Resiliency Options

The solution can be deployed in a number of different resilient configurations. These will be covered in more details with the different deployment topologies, but essentially cover:

- Single Data Centre - Non-Resilient (All-in-One)
- Dual Data Centre - Voice Service Resilient (Parallel Recording)
- Disaster Recovery to secondary Data Centre - Voice Service Resilient (Parallel Recording)

Call Recording is deployed in a resilient dual data centre configuration by default, Screen recording is only deployed at the primary location and relies on High Availability. Consult with your Mitel Account Team if there is a requirement for resilient screen recording.Considerations include the bandwidth needed between different data centres and the workload for de- duplication of the recordings. Resilient recording utilizes multiple recording servers with the Parallel Recording topology. A single Core server is still used in this deployment.

The single data centre deployments rely on High Availability to recover any lost server. This requires that any storage that needs to be recovered must be partitioned off-board the main server, through a SAN connection. This ensures that should a physical host server fail, and the virtual machine is recovered on an alternative physical host machine, that the stored data is still accessible. This applies especially to the database (PostgreSQL) on the main server, and also any storage of the recordings, if these are associated with the Core server. Separate partitions for application, database and storage are recommended. The same off-board SAN storage also applies to the off-board storage via the NAS server. This ensures thatthe archive server and data can also be HA recovered.

In order for the Dual Data Centre to provide resilient operation, there must be an existing connection between the networks within the two data centres. This allows recordings to be taken in parallel and remain synchronized with the core server for de-duplication at storage. The deployment uses parallel recording to ensure that if a recording server fails, that the recording is not lost. This means that the primary and secondary recording servers are running active/active and each of these must have an active license both on the server and for the recording stream. See under the License section for further details. The Active/Passive architecture is not used with this deployment, as there is possibility of lost calls if a recording server fails.

In order for the DRaaS Disaster Recovery scenario to work, there must be an existing network connection between the primary and secondary locations. The database and storage must be duplicated at both locations, so that on transition of the core server and database to the secondary location, the service can continue withminimal disruption. Disaster recovery involves a level of manual change-over and is not instantaneous.

A requirement for dual data-centre resiliency (dual DC and Zerto) AND the Recording on Demand feature (stop/start) requires that dual Web servers are available. If there is a single web server and this fails, control of the call is lost using the web browser. The web server is part of the Core server functionality (Server A or H). In the secondary data centrer, the backup secondary web server is provided as part of Server E*. This is the main function of this server, even though other features are also available. The DNS (internal and external) reference to the web servers (plural) needs to contain IP address to enable the web browser to reach either web server.

Some Default Settings

In order to correctly scale the size and quantity of the servers, some inputs on number of users, or agents, and expected usage is required. This also impacts the number of call records and therefore the expected size of the database, and also the recorded storage. Some defaults used to scale the solution arehighlighted below:

Traffic Patterns

There are three main traffic patterns for:

- UC Users, either trunk or user recording
- Agents, with agent recording
- Agents with trunk recording
- Screen Recording

To scale the number of recording channels some assumptions are made and highlighted in the table below. An agent is expected to be busy from 75% to 100% of the time, and so 1 agent requires 1 recording channel. All other recording channels are against this reference. UC Users are, during the peak hours, roughly 5:1 as busy as an agent, and this is reflected in the expected traffic patterns of 6CPH for a UC user compared to 30CPH for an agent. Therefore 1 concurrent channel is needed for 5 UC users. Trunks are also equally busy as agents, and so if this is trunk recording, rather than agent, or user recording, then trunk recording requires 1 concurrent channel per trunk. However, a contact centre will often deploy IVR queuing gateways as virtual agents, and often this can be as much as 50% of the number of agents, for example, . for a 100 agent deployment, it may be expected to deploy 50 IVR and therefore a total of 150 trunks (100 direct to agents, 50 queued on the IVR) In a contact centre deployment, it's therefore important to understand if this is agent recording or trunk recording, as the number of concurrent channels needed maybe different.

In terms of calculating storage capacity, the contact centre agents are typically busy for the full shift, typically 8 hours, and so this 8-hour shift is used to calculate the storage. For a UC user deployment, it is expected that peak traffic will occur at two points in the day, once in the morning, for 1 hour and once in the afternoon, for 1 hour. For the remainder of the shift there is still background traffic, but this only accounts for another hour of storage in the morning and an additional hour in the afternoon. As a result, a UC user solution is only storing the equivalent of 4 hours of storage per shift, compared the contact centre deployment running for a full 8 hours per shift. Therefore, although the number of recording channels are determined by the peak traffic rates, the storage is determined by the effective number of busy hours in the day.

Function	Ratio
Users to Agents	5:1, for example, 1 agent = 1 concurrent channel ; 5 users = 1 concurrent channel
IVR to Agents	1:2, or 50%, for example, 50IVR to 100 Agents (= 150 Trunks)
Trunk Recording	1:1, for example, 1 trunk = 1 concurrent channel

Typical call rates are:

Station	Traffic
Agent	30CPH prorated to 75%
User	6CPH prorated to 75%
Active Agent effective shift duration	8 hours
Active User effective shift duration	4 hours

Screen Recording is only associated with a user or an agent. It cannot be associated with a trunk. So even if the solution is configured for trunk recording, the screen recording is only associated with the agents. Therefore, in the 100-agent system, with trunk recording, the servers need to handle 150 concurrent audio calls, but only 100 screen recordings. It is not expected that UC users will need to screen record. Screen recording is also high bandwidth and typically associated with a managed network. Screen recording is therefore limited to on-net configurations, such as physically in the same building or routed or switched across the corporate network. It is not intended for direct use over the Internet where bandwidth and delays cannot be guaranteed. Screen recording call rates reflect the agent call rates.

Solution Inputs

Some default inputs include:

- Deployment and overall configuration options
- Audio defaults
- Screen defaults

These are identified below.

Deployment Options

The following tables provide items for consideration when deploying an MIR solution, and these line up with entries in the Sales Engineering Tool for Mitel Interaction Recording. The same values are used in CPQ.

Deployment Options	Value	Notes
Compression of Audio	Yes	Data in storage is compressed to reduce the storage requirement. However,the initial recording will not use compression. The recorded file will be post- compressed depending on certain server configurations: Post-compress after a pre-defined storage time, IF Keyword Spotting, orTranscription Services are NOT included in the solution, or Post-compress after the file has been processed by the Keyword Spotting, orTranscription Services. Keyword Spotting and Transcription Services work best on un-compressed voice recordings. Compression licenses are needed for this service to reduce storage
Concurrent Audio Channels	See Note	This is not the number of licenses but is related to the number of streams being recorded at the same time. For UC users, this is typically at a ratio of 5:1 based on users, i.e., for 1000 users, 200 concurrent streams are needed.For Agents and Agent Recording, the values are the same, for example, for 100 agents, 100 concurrent streams are needed. For Agents and Trunk Recording, this value also includes the number of front-end IVRs as well as the number of agents, for example, for 100 agents with a 50% IVR ratio, then this would result in a requirement for 150 concurrent streams
Concurrent Screen Channels	See Note (Max: 1000)	This value is based on the number of active agents that require screen recording at the same time. This is not related to the quantity of screen recording licenses which may differ depending on the deployment and license bundles. For example, if there are 300 named agents, but for the work shift there are only 200 active agents,then this value would be 200 (active), and not 300 (named). Screen recordingis are not normally associated with UC users, and therefore only considered with agent and contact centre deployments.
Replay Channels	1 per 25 agents (Max: 40)	The number of concurrent replay channels is typically defined by the number of agent supervisors. As part of the deployment sizing, the ratio of 1 supervisor to 25 agents is considered. The solution is limited to 40 concurrentreplay channels. A screen replay and audio replay count independently, i.e., ifa conversation includes both audio and screen recording, this counts as 2 channels. Solutions require more than 40 concurrent replay channels require approval. Please contact you Mitel Account Team, if this is the situation,
Resilient Deployment	No/Failover/ Parallel/Flex	The solution can be provided as a single data centre in which case this is NOT resilient, other than any High Availability service provided by the underlying infrastructure. When a resilient solution is required, the default is to provide the necessary components to provide audio call recording and control. Resilient Screen Recording is not provided by default and requires approval for such an installation. Please contact your Mitel Account Team, if this is a requirement. The core server, including playback controls are not resilient, as these are considered lower priority non-real-timeactivities.Select the option for ‘Flex’ for resilient deployments (default parallel recording).
Data Centre	Single/Dual	A resilient deployment will involve duplication of resources, including duplication of the data centres with geographic distance between them, i.e., not in the same data centre, or city. Default is therefore Dual.
Resilient Screen Recording	Yes/No	By default, resilient screen recording is NOT provided. A requirement to deploy a resilient screen recording solution requires approval. Please contact your Mitel Account Team, if this is a requirement.
Vmware Zerto Disaster Recovery	Yes/No	This setting is only applicable to VMware deployments where Zerto Disaster Recovery is in use. In this situation, additional Core servers and storage will be selected so that this can be duplicated into the secondary data centre Zerto DR Group.
Recording on Demand	Yes/No	This provides the ability to start and stop recording on demand, such as whencredit card information is presented during a call, or when due to privacy concerns a client request not to be recorded. This function is provided with the base configuration, including in a single data centre deployment. In a resilient (audio) configuration, this drives a requirement for a backup control server to allow calls to remain with control, even if the primary control is not available.

Audio Defaults

There are a number of defaults which generate storage requirements, but also the number of records needed within the database server. Therefore a few long calls versus a large number of short calls could result in a similar amount of recording storage but be quite different in the number of database records and meta-data that needs to be recorded. The system is scaled based on a higher number of short duration calls. The calls are also recorded from both the agent and the client and recorded as dual channels within the same audio file. This is defined as a 'stereo' recording, offering each party of the conversation in either ear. Some defaults include:

Audio Defaults	Value	Notes
Average Call Duration	2 Minutes	This is a high call rate and will drive up the requirements on the database server. This results in a call rate of around 30CPH (Calls per Hour)
Agent Activity	75%	Although we like to push agents to work for 60 minutes within an hour, in practice time is needed between calls to record information and get ready for the next call. Typically, agents handle calls for around 75%, with backup activities in the remaining 25%. This results in a call rate of around 23-24CPH.
Hours per day	8 hours	A shift is typically 8 hours long, and the agents are kept busy for the entire shift! For UC Users, there are only 2 busy hours in the day, and the remaining 6 hours are background traffic. The peak traffic rate, andstreaming infrastructure is determined by the peak rate, but the average call rate for a UC User will be around 4 hours in a single shift determines the database and storage requirements.
Mono/Stereo Recording	Stereo	Both sides of the conversation are recorded and stored as a stereo audio file. Select ‘Stereo’ only
Recording Retention	1 year	Typically, audio recordings require to be retained for at least one year. This is often a legal requirement, although longer retention periods are also possible

Screen Defaults

There are a number of defaults that apply to screen recording. Screen recording can require large amounts of data, especially given the resolution and number of screens that an agent may be using. The following are some defaults used to scale the solution:

Screen Defaults	Value	Notes
Screen Resolution	1920 x 1080	Other resolutions are possible, and these can adjust the amount of data that needs to be recorded. It's nice to have a large and high-resolution screen, but it has an impact on both the data recorded and also any bandwidth needed to transfer the recorded data from the client machine.
Screens per user	1	It is assumed that if the agent is screen recording, then they will have at least 1 screen. Multiple screens can be used, but multiply the data to be transferred and stored
Recording Time	2 Minutes	It is assumed that screen recording of a conversation and transaction willfollow the same traffic patterns as the audio component, i.e., if the audio call increases in duration, so will the associated Screen Recording
Frames Per Second	15 FPS	For an active screen, such as with video, or lots of screen pops or selections, a higher frame rate is needed, and typically this will be 15FPS. If the agent is filling out a relatively static form, such as a purchase request, with minimal screen updates or window selections, then a lower FPS may be selected: • Moving Video: 15 Frames Per Second (FPS) • Talking Head (Interview Style): 10 FPS • Form Filling: 5 FPS
Recording Retention	3 months	Due to the quantity of data to be recorded, and the requirements for playback, the default period for screen retention is 3 months. Lower, or higher values may be selected. Screen recording is the dominant driver for storage capacity, when used.

Note that screen recording is not supported via MiVoice Border Gateway and is only available via on-net connections.

Recording Storage Platform

This setting is used to identify the underlying infrastructure platform being used as well as the storage technology. A number of platforms and operating systems have underlying limits on the maximum addressing range of the storage. As a result of large storage requirements, such a with screen recording, it may be necessary to provide multiple storage partitions, and servers.

As well as platform specific storage solutions, there are also generic settings (default) for simple calculations, and also settings for on-premises selection

Recording Storage Platform	Value	Notes
Selection	Generic (default), AWS, Azure, Google, iLAND, VMware and on-premises	Determines the upper addressing range of the storage solution and adjusts the storage capacity and number of servers accordingly.
Additional Storage (User defined)	0	This allows a customer to adjust/increase the amount of storage available above that automatically calculated

Solution Outputs

From the defined default inputs, a number of outputs can be derived. The CPQ tool considers these values in determining the amount of storage and server resources that are required. Some key outputs include:

Key Output	Value	Notes
Audio Storage	12GBytes	1 Agent for 1 year (Stereo Recording and compressed data)
Screen Storage	246GBytes	1 Agent for 3 months
Database Audio	1.6 to 2GBytes	1 Agent for 1 year Exact value decreases with increased number of agents due to overhead)
Database Screen	0.4 to 2GBytes	1 Agent for 3 months (Exact value decreases with increased number of agents due to overhead)
Audio Network Bandwidth	100kbits/s	This is based on uncompressed data per call. For a voice call this is bi-directional bandwidth, i.e. both up and down links
Screen Network Bandwidth	169kbits/s	This is based on screen defaults of 1 screen per agent, 15 Frames per Second and a resolution of 1920 x 1080

Number of Records Audio*	46800	This is based on 1 agent making 180 calls per day for 260 days per year (note that an agent may be working for 240 days in a year, but the business may be operational for 260 days)
Number of Records Screen*	11880	This is based on 1 agent making 180 calls per day for 66 days (3 months)

***Note**: Where a screen recording and audio recording are linked, these may be combined into a single record. However, it is also possible to have screen recordings independent of audio recordings, for example as training videos. CPQ considers the situation where records are independent to cover this higher usage situation.

MIR Application Boundaries and Deployment Considerations

There are a number of internal limits and deployment boundaries that need to be considered at both the Call and Screen Recording application level and also at the solution level. A number of these boundaries are highlighted below:

In terms of scaling and deployment, the deployed servers follow the architecture flow as shown in the diagram above, that is:

- Server A (non-resilient)
- Server H+D (non-resilient and resilient)
- Server H+B+C (non-resilient and resilient)

Audio call recording and associated controls are deployed as resilient. Screen recording is delivered as non-resilient. Core functions and non-real-time functions are delivered as non-resilient

Some deployment considerations are shown in the diagram and described in the sections below:

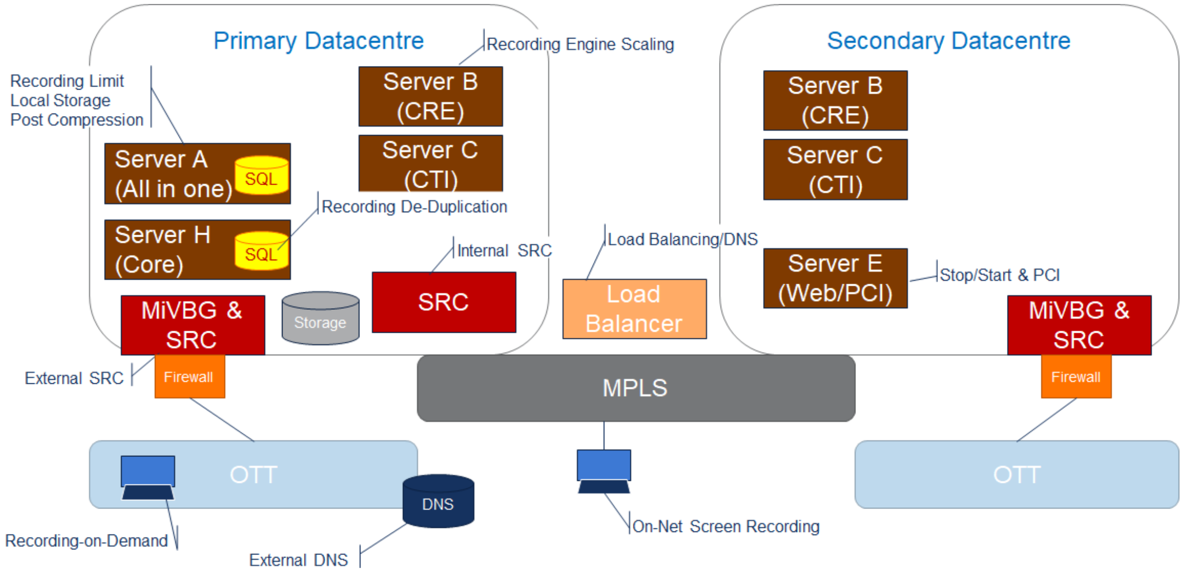

NAS Server

There are also boundaries that identify when an off-board NAS storage is recommended. The off-board NAS may also be identified as the ‘Archive’ server. The advantage of this external server is that it provides a means for the recording servers (Server B and component of Server D) direct access to the storage, without having to access this through a core server intermediary. It also allows recording servers in both primary and secondary network locations to access this across the network and be independent of the core servers. A customer provided NAS server is always needed when storage is provided by the customer or partner at an on-premises location. It is the customer’s or partner’s responsibility to install and maintain this server and storage.

Both of the Core Servers (Server A and Server H) can provide local storage, when possible. However, as the size of this storage increases the use of an external NAS server simplifies the recording server to storage connections and removes load from the core servers.

When recording storage exceeds 700GBytes, an external NAS server is required. Although CPQ will identify when an NAS server is needed, if there is an existing NAS server, this may also be used, as long as it meets the minimum resource requirements of the NAS server.

The data storage should be identified as a separate partition, stored off-board the core server, or via the external NAS server. This provides abilities to monitor and manage the size of the storage, without impacting other data. It also allows for a separate storage partition to enable HA recovery of the server.

PostgreSQL Storage

Although the PostgreSQL database is associated with the Core servers, the actual storage should be off board the virtual machine and should have its own logical partition. The size of the database is also dependent on the number of calls, which is also related to the number of users and traffic from those users.

The database for one deployment may differ from that of another deployment. The database needs to be sized accordingly, and managed, i.e., records can be deleted as they age out past the retention period, or alarms can provide alerts if the storage is exceeding expectations. A separate partition off-board also allows the database to be managed without contention with other storage and allows the core server to be recovered via High Availability operations and still maintain the data during that transition.

High Availability Requirements

High Availability guidelines relating to the deployment and type of storage should be followed. Primarily the biggest impact is the location of the data that needs to be recovered. In the case of the Core server, the primary storage is associated with the Application, the database and the call/screen storage. These should all be on separate partitions and managed accordingly, with their own thresholds and alarms. The application and operating system are more associated with the local drive and will be recovered with the virtual machine.

The CTI and Recording Control of Server C (and Server D) also includes an operational database. This again should be identified with a separate partition and off-board storage.

Server A Boundaries

In a non-resilient deployment, Server A (All-in-one) will be deployed when possible, rather than moving to the multi-server solution. A resilient deployment automatically starts with a multi-server solution.
The Server A includes all functions, including local storage, the database, playback, web interface, post-compression and real-time recording and controls. It can typically handle up to:

- 700 audio channels, OR
- 100 screen channels (or a ratio of the two)
- Up to 700G of local storage

The server A is available in different scaling versions with different resource requirements to cater for the smaller to the larger deployments. Beyond these limits, additional servers are needed, or a multi-server arrangement. Typically, the Server A may gain an additional 'Archive' server for storage before internal limits warrant the multi-server deployment

Playback Boundaries

The audio and screen playback functions are provided by the Core server (either Server A or Server H). An audio playback counts as one session, and a screen recording also counts as one playback session. Combined playback therefore counts as two sessions. Typically, playbacks are used for training and also by supervisors. The ratio of 1 supervisor to 25 agents is used as the default, and therefore a 1000 agent deployment will expect to have up to 40 playback sessions.

The number of playback sessions are limited to 40 concurrent sessions. Requirements for more than 40 playback sessions requires approval. Please contact your Mitel Account Team, if this is a requirement.

Recording Engine Boundaries (Server B/D)

Once the multi-server deployment is needed, Server H provides the core non-realtime functions of database, playback, web interface and post compression. Real time call handling is provided by server D (a combined CTI, Recording Control and Recording Engine), or for larger deployments a combination of Server C (CTI and Recording Control) and B (Recording Engine). The recording engine can typically handle:

- Up to 1000 audio channels, OR
- Up to 100 screen channels

For smaller deployments, a Server D or Server B may combine both audio and screen recordings. However, as the systems scales to larger sizes, separate Server B recording engines are needed, AND the

Server B are grouped by function, i.e., a number of server B are allocated for voice recording, and a number of Server B are allocated to screen recording. It is important that these servers are identified within the Recording Control (Server C) and also within the recording (screen) clients. This ensures that only the correct B servers are used, and screen recording is not allocated to audio recording, and vice versa. This becomes important for the following reasons:

- The MiVoice Border Gateway/SRC will only route to the appropriate audio recording servers
- Screen Recording is by default a non-resilient deployment, and if screen recording uses the audio recorders, this may impact the availability of these servers to audio recording. Often audio recording is associated with legal requirements, whereas screen recording is less so, if not at all.

When deploying a server B or Server D, it is recommended to have multiple storage partitions to cater for:

- Applications (recording engine and CTI (for server D))
- Cached recording

The server definitions include HDD storage that covers both of these partitions. Further details are provided under the Server Configurations section, below.

CTI Boundaries (Server C/D)

The CTI server includes a local transitory database as well as the CTI application. There is sufficient disk drive defined for the CTI server to cover both of these functions in a single partition. The same applies to Server D when this includes both Server C and Server B functions. It is recommended to store this data in a SAN environment to allow correct operation and recovery of the database through High Availability operations.

Disk Partitions

A number of the servers carry out multiple functions. Each of these functions requires a particular amount of disk storage. To reduce the possibility of services expanding storage beyond limits and to aid in alarming near capacity usage, it is recommended on some servers to define specific functions to a defined storage partition.

In certain cloud platforms, multiple virtual disks may be deployed, especially with database servers where a different memory storage type may offer operational improvements. Typically, in this case the operating system and application would exist on one drive with the C: partition, with the other virtual disk being partitioned with the other functions, for example, D:, E:, etc. The table below will help to identify those different partitions.

Server Definition	Application (C:)	Database** (D:)	Media Cache/Call Pool (E:)	Recording Storage **(F:)	Notes
Server A	Yes	Yes	Yes	Yes	This is an all-in-one server. Database and Storage will be defined in CPQ
Server B	Yes		Yes		The additional cache for the recording storageis defined in the Server Configurations section, below
Server C	Yes				Although there is a Database with this application, they can be on the same partition
Server D	Yes		Yes		The application (B+C) and database of ServerC can be combined to the same partition. Theadditional cache for the recording storage is defined in the Server Configurations section, below
Server E*	Yes				(optional) This is primarily a webserver
Server H	Yes	Yes		Yes	This is a core server, without local storage
Server NAS/Archive	Yes			Yes	The Archive server is primarily a storage (NAS) unit. The storage requirements are defined in CPQ

> ****Note**: The required amount of storage is dependent on the deployment, number of users, usage andretention period. This information is provided through CPQ. The application allocation is defined in theServer Configurations, below.
>
> Although a drive partition identifier is suggested, there is no requirement to use these specific values (i.e. C:, D: and E:). The available partition identifiers may differ depending on other services on the server, for example, fthere is a CD/DVD drive, or similar services.

The server definition provides sufficient HDD storage for the Operating System, the application, the PostgreSQL application and media cache. The media cache can be extended by adding additional memory, if required. The media cache occupies a separate partition. This is also identified as “Call Pool” as it used to temporarily hold call recording in a local storage pool.

The Database requires a separate partition, even if this is stored on the local server. This partition is independent from the OS/Application drive. The size of the database is also determined by the size of the deployment and any retention periods. This is calculated through CPQ.

The Storage, also described as Extended Storage, is used to store the recordings. This can be an additional partition and drive on the main server, if sufficiently small, or could be connection to a NAS/Archive server, or a connection to a file store location. The use of a file store location is especially convenient in cloud deployments, and suitable options include Cool Storage, Cold Storage and S3 storage, depending on the cloud provider. Consideration for the amount of read accesses to the data also needs to be considered. These long terms storage locations offer a good price point for data that is written but rarely read. If particular files are read often, then it may be better to create a copy of these on a separate drive. An example would be where screen recordings are accessed often for training.

The following recommendations apply to how the partitions should be applied for a single All-in-One Server A:

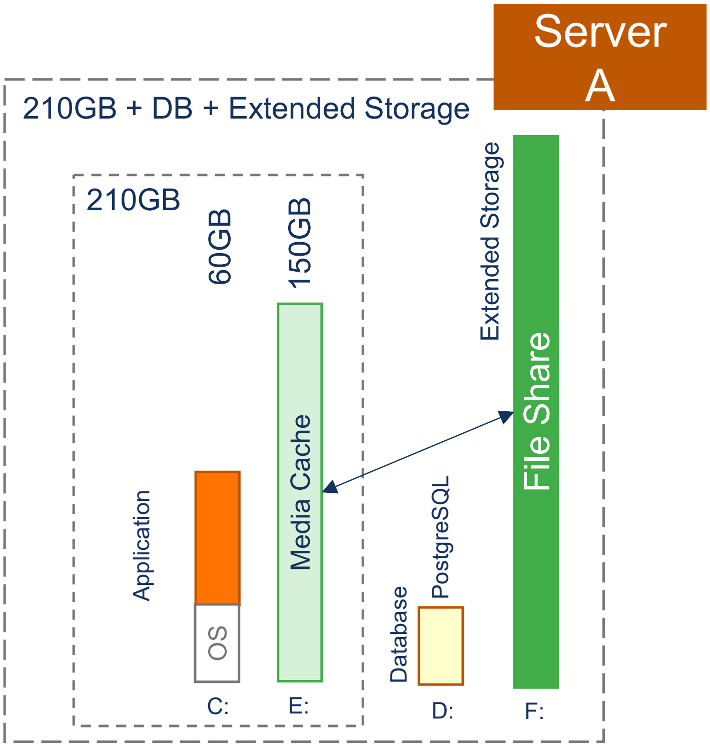

Typically, the server is specified with a minimum HDD requirement. Additional storage may be applied or may be provided in certain cloud environments. The specification covers enough storage for the following:

The recommended storage partitions for Server A are as follows:

Disk Partition Server A	Usage	Size
C:	Operating system and Application	60GB minimum (Included in server specification)
D:	Database Storage	User defined, additional
E:	Media Cache (a.k.a. Storage or Call Pool, which is separate from Extended Storage – see under F:)	150GB or 200GB (Depending on server definition)
F:	Call Recording Storage, also referred to as Extended Storage	User defined, additional

The media cache (or Call Pool) is used for a number of activities, but primarily:

- Temporary Storage before transferring to Extended Storage (long term and archive)
- Speech Post-Compression
- Allowing immediate playback of a recent call
- Transfer to Speech Analysis and returned file, prior to transfer to Extended Storage (long terms and archive)

For the minimum HDD specification this provides storage of information for at least 1 office/working day and allows sufficient time to complete a full day of speech recordings via the Speech Analysis. More drive can be added to increase this time further.

The Extended storage maybe realized as additional HDD on the server, where this storage is small enough and unlikely to grow. It may also be realized as external file share for certain cloud deployments. Something under 700G is considered acceptable. Beyond the 700G guidelines, use of an external Archive/NAS server is recommended, as shown in the multi-server configuration, below:

The following recommendations apply to the disk partitions for a multi-server deployment:

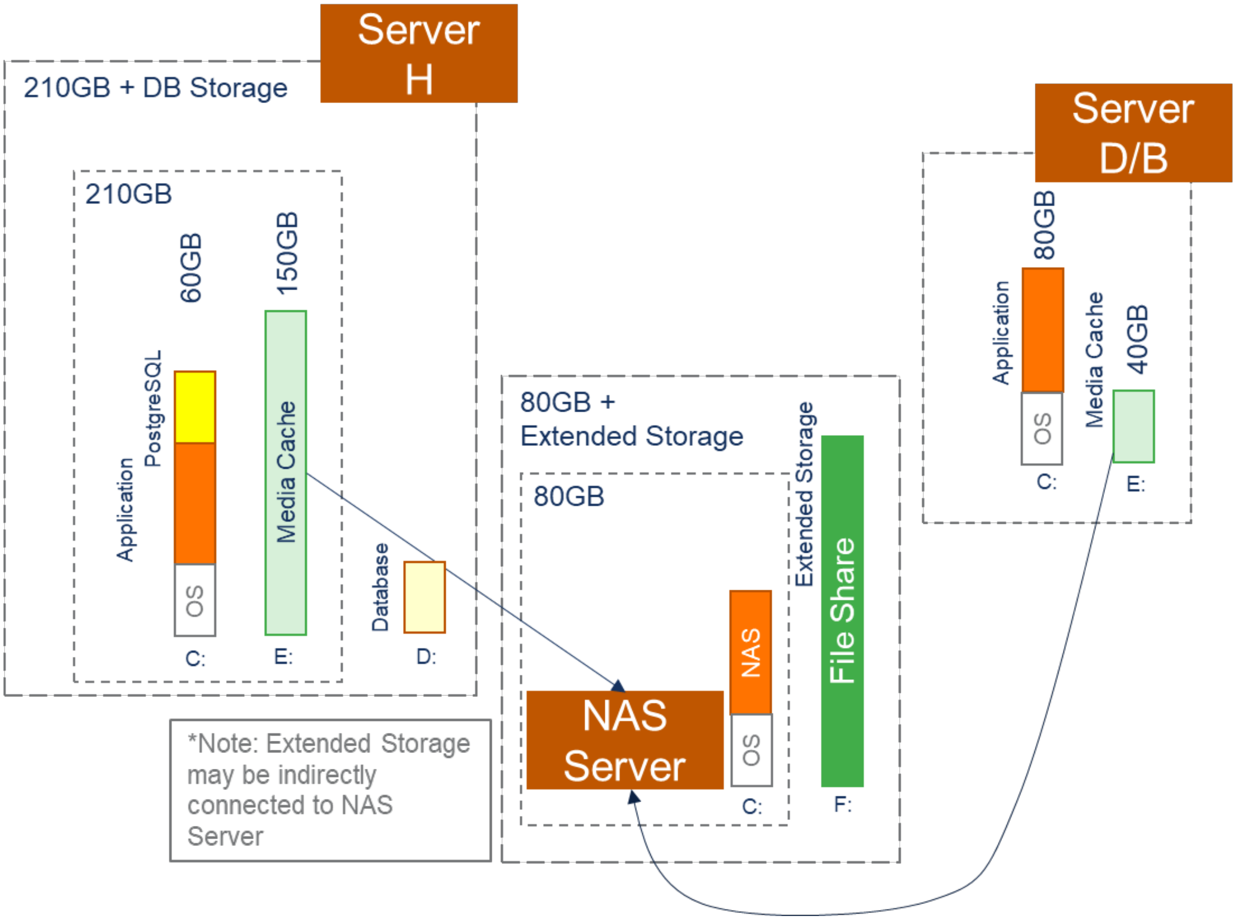

Note: The Archive Server is the NAS Storage Server.

Server H is similar in operation to the Server A (All-in-One) except that recording is carried out on a different server. This server still requires the media cache as this will still need to carry out a number of post recording tasks. Although the separate Data HDD is identified as HDD, and will work as HDD, the recommendation for these larger servers is to use SSD where possible.

The recording servers do not require a database server, so no additional partition is required for this.

Although for smaller deployments it is possible to apply the Extended Storage to the core Server H, this becomes a bottleneck for larger systems. Use of a separate Archive server is recommended, with the Extended Storage attached to this. The Core server and recording servers then connect to the Archive/NAS server rather than the file share.

The Server H uses the same OS HDD profile as the Server A and provided the same minimum 1 office/working day of cache. Additional storage may be included.

The recording server includes additional cache storage for temporary hold, prior to being transferred to the archive server. This provides up to 2 hours of media storage, and is scaled with the recording server:

The recommended storage partitions in a multi-server are as follows:

Disk Partition Server B or D	Usage	Size
C:	Operating system and Application	80GB minimum (Included in server specification)
D:	Not defined	N/A
E:	Media Cache/Call Pool	Small: 40GB Medium: 80GB Large: 120GB (included in server specification)
F:	Not defined	N/A

Disk Partition Server H	Usage	Size
C:	Operating system and Application	60GB minimum (Included in server specification)
D:	Database Storage	User defined, additional
E:	Media Cache/Call Pool	150GB or 200GB (Depending on server definition)
F:	Call Recording Storage, also referred to as Extended Storage (Only if NAS Archive server is *not* used)	User defined, additional

Disk Partition	Usage	Size

NAS / Archive Server		
C:	Operating system and Application	80GB minimum (Included in server specification)
D:	Not defined	N/A
E:	Not defined	N/A
F:	Call Recording Storage, also referred to as Extended Storage	User defined, additional

Other logical partition definitions may be used, However, these names provide consistency within the deployment, and therefore recommended.

Call and Screen Recording Storage Location

Call and Screen Recordings are typically stored within the hosted solution data centre. This has the advantages of low latency, an ability to up-scale the storage capacity on demand, as well as optimizing costs.

However, it is possible for a customer to provide on-premises storage by simply mapping a storage drive that can be accessed from the hosted solution, as long as the following caveats are adhered to:

- The on-premises location is connected to the hosted data centre with a private network connection
- The latency of this private connection between Recording Server(s) (Server A, B, or D), to the storage server is no more than 50mseconds
- That a separate NAS Archive server is always deployed on-premises, even when the storage requirement is below 700G (which could normally be accommodated directly on the Core Server)

When storage is provided on-premises it is the customer's or partner's responsibility to ensure that adequate capacity is provided and to provide ongoing maintenance and updates of this component.

Storage Boundaries

For larger deployments it is possible that the amount of storage required for the call and screen recordings will hit storage address boundaries. In this case it may not be possible to deploy the required storage as a single drive. For example, Windows will limit at 64TBytes, VMware at 62TBytes, and a service provider may provide further limits (iLAND is limited to 30TBytes).

The restriction is that any recording server (Server A, B or D) can point to a single storage drive. Recording servers may point directly to the recording storage or may point to the core server (A or H) which will forward to the appropriate storage. The recording servers also need to connect to the core server (A or H) to ensure that file locations are also included into the database for future recovery.

As an example, if an installation is too large for a single drive partition, it could be split into two drives, for example, one for audio and one for screen. Multiple recording servers can point to the same recording storage. Or, there may be a requirement to add more storage partitions/drives:

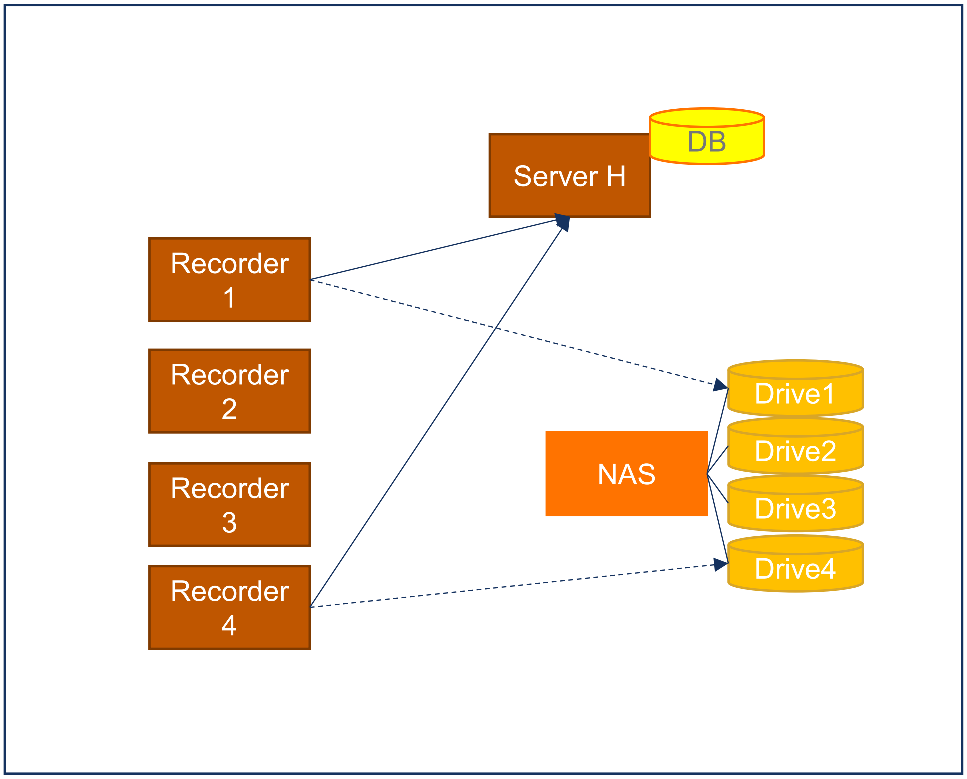

In this example there are 4 recorders and each points to a specific drive associated with this recorder. Each recorder still connects to the core server (H) in order to update the database.

Care is need at installation to ensure adequate load balancing and retention policy settings to limit the issue where one of these drives is oversubscribed, yet other drives are under-utilized. Screen recording, for example, will typically distribute call recording in a round-robin fashion across the group of available servers.

Consider also the case of a resilient deployment, where parallel recorders will point the same storage drive.

Call Recording CODEC support

The call recording solution supports the following audio CODECs:

- G.711 (A-Law and μ-Law), G.729, G.722, OPUS, SILK, ACELP

  Note that G.722.1 and G.722 are not compatible. Mitel phones will default to use the G.722.1 wideband CODEC for internal calls, and since this is not a supported call recording CODEC, calls may fail to record if this CODEC is selected. Calls to and from Trunks typically use G.711.

  In order to record internal calls, the MiVoice Business needs to disable the G.722.1 CODEC. This can be filtered in the Codec Settings form with G.722.1 filter set to ‘Yes’. This removes the G.722.1 CODEC from negotiation. Wideband CODEC operation for non-recorded calls will be via G.722.

  Currently the MiVoice Border Gateway does not support G.722 as a recording CODEC, in which case calls that are recorded will revert to G.711 during the setup negotiation.

This filter setting is a MiVoice Business cluster wide setting and will apply to all MiVoice Business and users in the solution.

Audio Compression and Post Compression settings

Transcription Services and Keyword Spotting work best with non-compressed audio signals. Therefore, the solution should ensure that trunk connections and agents are connected using non-compressed (G.711) links. This may require configuration of the MIVB zones to limit CODEC options or to group agents within a zone that will not by default force use of the compression CODEC. The network aspect of this is that any WAN links must provide sufficient bandwidth and QoS controls for the number of users and agents. Don't forget that agents can be active 100% of the time, so don't count on Erlang to offer reduction recommendations in network bandwidth!

Audio recordings are initially stored in non-compressed format. The Core servers (Server A or Server H) will then perform post-compression to these files, along with updating meta-data on the database. There are a number of settings that need to be configured for this:

- Post-compression can be defined to occur after a set period of time after the file has been saved. This option is the preferred option when Transcription Services or Keyword Spotting are NOT part of the solution
- Following analysis for Transcription Services or Keyword spotting. Given that these additional services may take a number of hours to process files through the day, this setting is preferred to ensure that files are not post-compressed in advance of the transcription/keyword analysis.

Compression licenses are needed for this service to reduce storage.

For MIR Release 7.0 and 7.1, these capabilities are unavailable when using Azure Cognitive Services. As a result, recording storage will remain uncompressed, and require additional storage space/medium.

Screen Recording Boundaries

Screen recording is primarily targeted to contact centre and use with agents. It is not intended for general UC user deployments. Screen Recording can only be included in a solution configured for Agent Recording. It cannot be included with a solution operating as Trunk Recording.

Screen Recording requires a high bandwidth connection per screen, and with multiple agents and screens requires an even higher bandwidth and streaming capacity at the central concentration point. This concentration of this data places additional resource demands on firewalls and border gateways. Due to the impacts on the infrastructure, screen recording is restricted to on-net, or LAN connected devices. For users at remote locations, use of a dedicated VPN is required.

Data Retention Settings and Storage Alarms

In deploying a solution, the default retention periods are stated, and from this the amount of database storage and recording storage are derived. These defaults are 1 year for audio recordings and 3 months for screen recordings. Storage limits are defined with the initial installation of the solution and cannot be exceeded without additional configuration and associated costs. It is therefore important to ensure that files that are older than the defined retention period are aged out gracefully. The default file retention period can be defined within the application, and it is highly recommended to set these values to prevent overrun and possible blocking of storage of new files. Individual files that need different retention periods can be adjusted afterwards, but it is expected that the majority of files will take the default settings.

It is also highly recommended that alarms are configured within the application to provide notifications when the storage (database and recording storage) are nearing their capacity. In the event that storage is being used at a higher rate than expected, this will allow time to investigate and adjust policies, review the stored data, or extend the storage, as required.

SRC Locations and Load Adjustment

In order to record the audio signals, the payload data must first be decrypted (cipher to plain text) and re-encrypted with the call recording solution encryption keys. This function takes place within the Secure Recording Connector (SRC) function of the MiVoice Border Gateway.

There are number of locations for the SRC:

- At the external gateway boundary for Teleworker devices (OTT Users)
- At the external gateway boundary for Trunk connections
- Internally for LAN connected devices (LAN SRC)

For agent, or user, recording, both the External OTT User gateway and the internal LAN SRC are needed. For trunk recording, only the SRC at the Trunk gateway is needed. It is important to understand which deployment is being installed, either agent or trunk recording, but not both. There are some limitations on devices that can be recorded:

- MiVoice Business Console can only be recorded with Trunk recording
- Teleworker devices include 53xx and 69xx MiNet phones, SIP phones and MiCollab Softphones.
- WebRTC is not included
- WebRTC Pro can be recorded as it behaves in the same manner as SIP Teleworker device

PC Softphones and Mobile softphones should register as Teleworker devices, even if deployed on an internal network and use the External SRC, since these are potentially mobile devices.

For a resilient deployment, there will be multiple streams from the SRC to both the primary and secondary recorders. This double load should be considered when scaling the SRC and MiVoice Border Gateway. Typically, the SRC will add an additional 50% load on the number of streams. With dual streams, this additional load increases to 100%. Therefore, if the MiVBG nominally handles up to 500 streams, when deployed with an SRC the system load is increased to 100% (existing streams) + 50% (primary connection) + 50% (secondary connection), or 200% in total, or 250 streams. CPQ will take this load into account when defining the solution.

In some situations, such as large-scale resilient deployments, it may be necessary to increase the number of MiVBG/SRC in a cluster group, and this value may exceed the recommendation in the *MBG Engineering Guidelines*. Should this occur, please contact your Mitel Account Team for consultation before proceeding.

Direct or Indirect Recording

Agents, or users, are only recorded using the '**Direct**' method, i.e., the devices must connect and streamvia an SRC. This is especially true for OTT connected devices, where they must transition the external gateway before accessing the LAN connected application. Trunk recording is by deployment a 'Direct' recording.

Indirect recording, without SRC, is not supported for this solution.

Recording on Demand (Stop/Start control)

An optional feature is the ability to stop and start the recording dynamically. This is especially important in a situation where credit card information (Payment Card Industry) must be provided, or where there are privacy concerns, for example, the client may object to recording (for example, . GDPR), or the content of the call may contain sensitive information (for example, . Public Health Information). In these situations, it will be necessary to prevent recording of the call, or to stop and re-start this. This is provided by the 'Recording on Demand' control and can be accessed via the application web-portal, which is included with the Core Server (ServerA or Server H).

Connection to the Web Server is directly available for agents, or users, that are on-net.

For a non-resilient solution there is only one core deployment, and the web server is available to provide this function. If the web server alone fails, then control for this function will also fail. Calls may continue to be recorded. Agents will need to be aware of this.

For a resilient solution, and additional secondary server is required, and this is provided with the secondary components. This server, Server E*, provides the additional web services. In order to provide this function, the web browser must be able to handle dual IP address information from DNS, and the DNS servers must be provided with dual/multiple IP addresses to the two web browsers.

Automatic Recording on Demand control is available. This uses screen recording capabilities to identify screen locations being accessed to trigger the recording control. This feature requires screen recording/scan licenses, even if the screen is not being recorded, as well as use of a thick client on the agent workstation. Because this is using screen services and the thick client, this capability is only available via on-net connections. Please contact your Mitel Account Team, should you wish to use this feature.

Recording File De-duplication

When calls are recorded in parallel (primary and secondary), there is a level of coherence between the stored information and the associated meta-data. This allows the de-duplication function in the core server to recognize duplicate information and remove this, thereby saving recording space.

When screen recordings and audio are maintained as separate files, there will be two audio files, which will be de- duplicated, and a video file (non-resilient recording). There will thus be two files remaining of a conversation and screen actions. The defined storage considers this situation. However, where the audio is tightly coupled with the video, such as when using audio from the PC, there will remain a difference in the recorded files; one will contain audio, the other audio and screen. Both of these files will be different and therefore not de-duplicated. This may result in increased storage due to the additional audio file. The storage calculation does not consider this situation. However, in 3 months, the combined screen and audio file will reach the retention period and be aged out, leaving the audio only file. If it is expected to have tightly coupled audio and screen recording, additional storage capacity for audio for the video retention period may need to be considered, i.e., increased.

External Firewall Port Forwarding

A number of services, especially those for agents, are provided through a web portal, or connection. Currently these are only available for on-net connections, i.e. via MPLS or SD-WAN connections.

Server Definitions

There are a number of server variants, as well as resource requirements. A cross reference list of these to the SKU and Ordering description is provided in the Server Configuration section below.

The recording servers (Server B and Server D) include sufficient cache in the hard drive to support up to 2 hours of recordings that can't immediately be transferred to the storage location. If it is necessary to increase the amount of storage for a longer duration, use the following formulae:

- 58MBytes per user per hour to calculate the additional voice storage. For example, if the system has 100agents and an additional 2 hours is required (for a total of 4 hours), then add a further 12GBytes of HDD storage: (58MBytes x 100 x 2 = 11.6GBytes, rounded up to 12GBytes)

- 600Mbytes per user per hour to calculate the additional screen storage. For example, if the system has100 screen recordings and an additional 2 hours is required (for a total of 4 hours), the add a further 120GBytes of HDD storage: (600MBytes x 100 x 2 = 12000MBytes, or 120GBytes)

Note that at the maximum recorder limits, there are 10 audio (stereo) recording channels to 1 screen channel.

Note that these calculations are based on the default traffic rates identified earlier.

Private Network and OTT considerations

Agent recording is available for both OTT and on-net connections via the appropriate Secure Recording Connector unit.

However, a number of other functions, including Recording on Demand (stop/start) and playback (web and thick client), are only available for on-net connections.

This deployment requires on-net connections, even if not used by remote agents. The application is therefore only suitable for deployments with private network connections. It is not suitable for a purely OTT solution, where there is no access to the hosted servers, and therefore access to the other functions.

Deployment Resiliency Options

There are a number of deployments with varying levels of resiliency and disaster recovery. These are highlighted below

- Single Data Centre with High Availability. This ensures that a failed server can be recovered asquickly as possible
- Dual Data Centre (geographically distanced, i.e., not in the same city) with a resilient deployment. This ensures that if an outage hits a region, or access to the entire data centre, such as a core router outage or upgrade, that there is a secondary and active solution in a geographically dispersed location
- Disaster Recovery. This provides resilient backup for voice related activities, much as the Dual Data Centre offers. However, this ensures recovery of the primary data centre and functionality in the eventof a major outage at the primary location, for example natural disaster. The switchover to a disaster recovery solution is not instantaneous. It replicates key data (storage and database) from the primary location such that a working system can be quickly recovered and brought online. A non-disaster solution would potentially lose all data, depending on location and frequency of backups and time to recover that data.

These four solutions are covered in the different topologies highlighted below.

Guidelines for High Availability should be followed. Any data that needs to be stored should be resident on a separate NAS or SAN drive with a separate partition. This ensures that this persistent data can be replicated, for example, to another data centre, andcan be accessed from any physical server where the applications may be active. The same rules apply for VMotion. Separate storage partitions for the database and the storage are strongly encouraged for:

- Ability to provide alarms on amount of data being used and whether it is reaching limits
- Storage may be on a different less costly, and less accessible storage media than the database whichis constantly accessed.
- Server A and H include a database and storage that needs to be recoverable
- The archive server includes recorded data that needs to be recoverable
- Server C includes an operational database that needs to be recoverable

Feature Support

While all of the features of MIR are available to devices that are on-net, there are limits to the features that are extended out via the Internet. Feature support andcomparison of on-net and OTT connections can be found under the *Feature Comparison (Private and Public Networks)* section, earlier.

Licenses

A summary of the licenses needed for the deployment can be found in the MiCloud Flex General Information Guide (GIG) and/or Ordering Guide in InfoChannel.

For the dual data-centre resilient configurations, the deployment is configured for parallel recording. This requires some additional licenses for the resilient components to become active. Refer to the MIR Ordering Guide for a more complete overview of licensing.

Use CPQ to calculate the requirements.

Server Configurations

The following server configurations are used in the MiCloud Flex MIR deployments:

SKU	SKU Description	Server Description	Server Size	vCPU	RAM	HDD (GB)	Recording Cache (included GB)
54011357	vWFO MIR Srvr Small All-in-One Infra	Server A Small	Small	4	16	210	
54011358	vWFO MIR Srvr Medium All-in-One Infra	Server A Medium	Medium	6	16	210	
54011359	vWFO MIR Srvr Large All-in-One Infra	Server A Large	Large	12	32	260	
54011363	vWFO MIR Server B Audio Small Infra	Server B Audio Small	Small	4	4	120	Includes 40GB
54011364	vWFO MIR Server B Audio Medium Infra	Server B Audio Medium	Medium	6	6	160	Includes 80GB
54011365	vWFO MIR Server B Audio Large Infra	Server B Audio Large	Large	8	8	200	Includes 120GB
54011366	vWFO MIR Server B Screen Small Infra	Server B Screen Small	Small	4	4	120	Includes 40GB
54011367	vWFO MIR Server B Screen Medium Infra	Server B Screen Medium	Medium	6	6	160	Includes 80GB
54011368	vWFO MIR Server B Large Infra	Server B Screen Large	Large	8	8	200	Includes 120GB
54011369	vWFO MIR Server C Small Infra	Server C Small	Small	4	4	100	
54011370	vWFO MIR Server C Medium Infra	Server C Medium	Medium	6	6	100	
54011371	vWFO MIR Server C Large Infra	Server C Large	Large	8	8	100	
54011372	vWFO MIR Server D Small Infra	Server D Small	Small	4	4	120	Includes 40GB
54011373	vWFO MIR Server D Medium Infra	Server D Medium	Medium	6	6	160	Includes 80GB
54011374	vWFO MIR Server D Large Infra	Server D Large	Large	8	8	200	Includes 120GB

SKU	SKU Description	Server Description	Server Size	vCPU	RAM	HDD (GB)	Recording Cache (included GB)
54011360	vWFO MIR Server ESmall	Server E* Small	Small	4	4	80	
54011361	vWFO MIR Server EMedium	Server E* Medium	Medium	6	6	80	
54011362	vWFO MIR Server ELarge	Server E* Large	Large	8	8	80	
54011375	vWFO MIR Server HSmall Infra	Server H Small	Small	4	8	210	
54011376	vWFO MIR Server HMedium Infra	Server H Medium	Medium	6	16	210	
54011377	vWFO MIR Server HLarge Infra	Server H Large	Large	12	32	260	
54011381	vWFO MIR Server	Archive Small	Small	4	4	80	

Server E* is a subset of the ASC provided Server E, to provide additional web services. See thedeployment guide on how to achieve this configuration.

The Recording Servers (Server B and Server D) include additional Cache HDD Storage to allow up to 2 hours of call recording or screen recording in the even that these servers cannot contact the central storage server. Additional storage can be added to increase the cache time, along the lines of:

- Recording Server Small: 40GB = 2 hours of storage
- Recording Server Medium: 80GB = 2 hours of storage
- Recording Server Large: 120GB = 2 hours of storage

Note that audio when cached is recorded in stereo and non-compressed. The cache takes this into consideration.

The servers to be deployed are calculated in CPQ and consider the different configurations and input requirements.

Deployment Considerations or Best-Practices

Although CPQ and many of the rules above define the MIR application and installation, some additional factors need to be considered, as these may adjust any defined deployment. The following should be taken into account:

- Server Resource specifications are *minimum application values.* Include a margin for use, maintenance and growth
- Application and Operating System updates may take longer if there is insufficient RAM or hard drive available.

- Consider additional RAM for Windows updates, above the minimum application recommendations
    - Servers B, C and D – minimum 6G
    - Server H – minimum 12G
    - MiCC SQL Server minimum 8-12GB (where integrated with MiCC-B)
- Include additional 1-2GB of RAM to allow Anti-virus tools to operate
- Ensure that Windows virtual memory setting is configured for ‘automatic’ and there is sufficient HDD available
- Apply Windows updates frequently to reduce disk space consumption of pending updates
- Ensure that Anti-Virus settings are updated to prevent false positive alerts and blocking of file transfer between servers and application software updates
- Consider separation of video and audio recording into separate recording servers to cater for future expansion
- Take advantage of any Cloud Tools to automatically update the Windows OS
- Use integrated monitor tools to provide alerts of application and system issues
- Use external monitoring tools to provide alerts such as Mitel Performance Analytics (MPA)

Deployment Configurations

The deployment configurations are calculated within CPQ, based on the provided inputs.

The following tables provide some deployment examples for different configurations. In the event of discrepancies in this information and that provided by CPQ, CPQ will take precedence. Storage vales are based on default recording retention periods, and default traffic rates. Changes to these values will result in a different capacity requirement.

Single Data Centre

In this deployment, all servers are deployed in a single data centre. The deployments will follow the basic deployment patterns below:

Deployment Overview

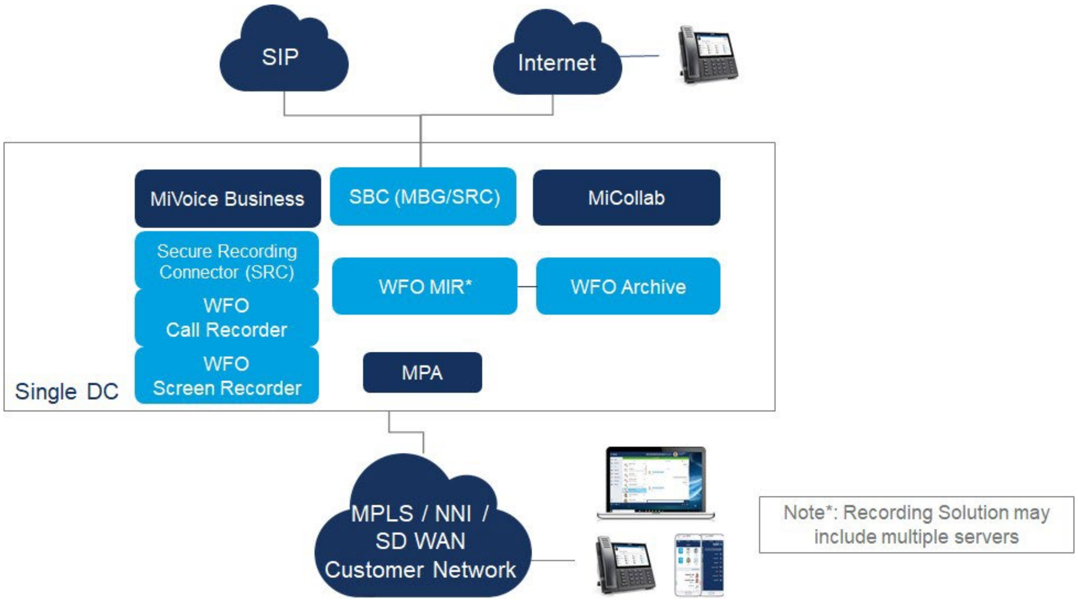
SIP
Internet
MiVoice Business
SBC (MBG/SRC)
MiCollab
Secure Recording Connector (SRC)
WFO MIR*
WFO Archive
WFO Call Recorder
WFO Screen Recorder
MPA
Single DC
MPLS / NNI / SD WAN Customer Network
Note*: Recording Solution may include multiple servers

Dual Data Centre

In this deployment, servers are deployed in both a primary and secondary data centre. Specifically recording servers and the optional web server are deployed into the secondary data centre.

The deployments will follow the basic deployment patterns below:

Deployment Overview

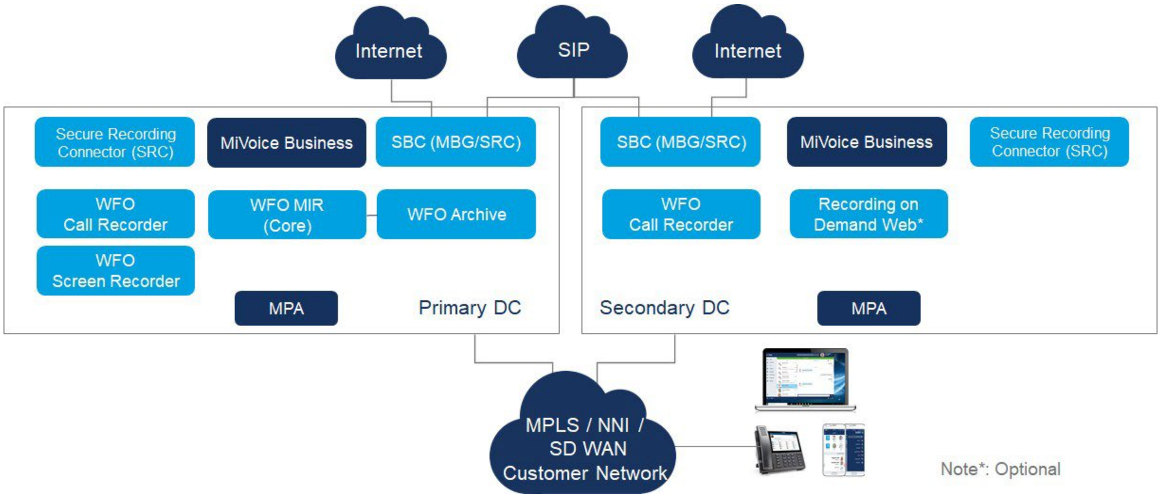

Dual Data Centre with Zerto Disaster Recovery

In this deployment, servers are deployed in both a primary and secondary data centre. Specifically recording servers and the optional web server are deployed into the secondary data centre. However, for servers and storage that is not resilient, Zerto DR will replicate these units into secondary data centre from the primary data centre. This applies to the Core server (Server A or H) and associated database and recording storage. Even if the non-resilient servers are not active, the reservations for the server still need to be applied to ensure that resources are available when needed. The storage data between the primary and secondary data centres also needs to be replicated so that it is immediately available when the secondary Zerto backed server is brought into service.

The deployments will follow the basic deployment patterns below:

Deployment Overview

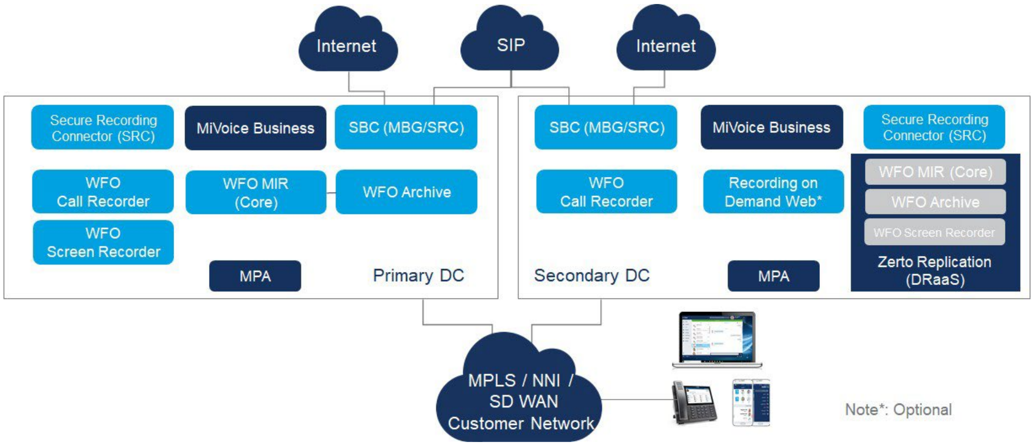

IP Ports needed in the Solution

The hosted solution connects to both the customer network and also to the Internet with the MiVoice Border Gateway as the Session Border Controller. Even though the hosted solution is connected to the customer network, there may still exist a firewall or router with limiting rules applied. There will be an external firewall and also configurations that are needed on the MiVoice Border Gateway.

Internal IP Ports

These are IP Ports that are on the customer private network. The following ports need to be opened, in order to ensure correct operation of the solution to and from the customer network:

Port	TCP/UDP	Direction	Description
443	TCP	To Web Server from Customer LAN	Web Services
2601	TCP	From Application to on-site MiCC	On-site/On-Premises Contact Centre Application
4000	TCP	To Playback Client service fromCustomer LAN	Playback Pro Services from Client application
4040	TCP	To Playback Web service from Customer LAN	Playback Web Services
4711	TCP	To Playback Client service fromCustomer LAN	Playback Pro Services from Client application
4498, 4499	TCP	To Screen Recording from Customer Screen Recording Client	Screen Recording from PC Client for recording and CTI control
6810	TCP	From Application to on-site MBG/SRC	On-site/On-Premises MBG/SRC gateway
20000 to 23999	UDP	To/From Application and MBG/SRC	Streaming from on-site MBG/SRC (bi-directional)

External IP Ports

External IP Ports are primarily for Web-Proxy connections and recording playback. The following IP Ports need to be enabled at the firewall and on the MBG:

Port	TCP/UDP	Direction	Description
443	TCP	To Web Server from Customer WAN	Web Services
4040	TCP	To Playback Web service from Customer WAN	PowerPlay Web Playback

Some of the features provided include:

- User and Admin login

- Quality Management and monitoring (Admin)
- Reports and dashboards
- Web Playback (requires port 4040)
- Web Stop/Start recording
- API services

Speech Analysis Transcription Services and Keyword Spotting

Speech analysis of the recorded audio files is possible. This allows the files to be fully transcribed from audio into text, which can be printed as well as searched. A slightly reduced capability of keyword spotting is also available, where the presence of certain keywords is identified in order to simplify future searches. For details refer to the section on Speech Analysis

As noted earlier, Speech Analysis works better with non-compressed audio, and post-compression of thefiles, after analysis, should be used.

Training Material and Professional Services

The deployment and integration of the MIR solution is quite complex. A good understanding of the product and essential components is a requirement to a successful deployment. It is strongly recommended to take training on the product and how to integrate this, and to also contact Professional Services to assist. Training is available through the Mitel Learning Management System (LMS).

Alternative Deployments

Can't seem to get a deployment to meet your requirements? Maybe you've exceeded some of the boundaries? Not sure where to go to next? Please contact your Mitel Account Team.

Speech Analysis using Azure Transcription Cognitive Services

The Azure Cloud API Service offers the capability to access a number of data analysis services. One of these, called Transcription Cognitive Services, offers the capability of transcribing speech information into plain text. The service is offered in a number of different languages and dialects, or voices, such as US and UK English.

The MIR recording solution integrates via a local connector with the Azure Cognitive Services to provide this transcription service. This runs in batch mode, using post recording analysis, rather than real-time direct streaming. The output from the Azure Cognitive Services is a simple text file, or files, including timing meta-data, which is then attached to the recording file. MIR also includes an advanced search engine (Release 7.0), which can search through these transcribed files for specific keywords to provide a faster search option.

Further details on the Azure Cognitive Service can be found here:

https://azure.microsoft.com/en-us/services/cognitive-services/#overview

The service being used is the “Speech to text” transcription service.

In order to use this service, the customer is also required to have an account on Azure and to set up the necessary services for connection from MIR, including public DNS entries.

Transcription and Keyword Spotting

Transcription services are provided through the Azure Transcription Cognitive Services. This operates in ‘batch’ mode and is essentially non-real time. Keyword spotting is carried out following Transcription using an advanced search engine included in Mitel Interaction Recording 7.0 onwards.

Connection to Azure Cognitive Services

The MIR server connects to the Azure Cognitive Services via an API service. This requires a connection over the Internet between the customer installation and the customer defined Azure region. Public FQDNs will be needed to identify the service and for inclusion in the MIR configuration, and therefore available prior to MIR installation.

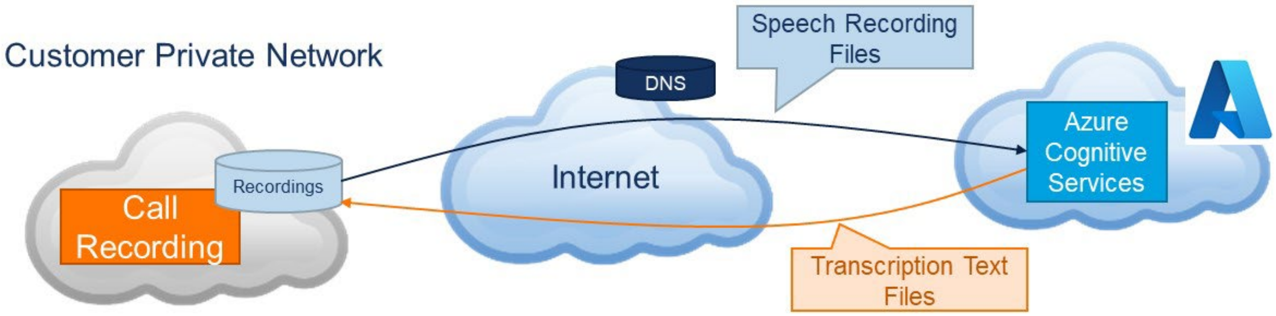

Connection and Limit Considerations

There are certain considerations to be aware in the deployment, with the key ones outlined below:

Maximum Number of Agents

It is assumed that the recordings to be transcribed originate from agents, with the following traffic parameters:

- 30 Calls Per Hour per agent, with a hold time of 120 seconds

- In a single shift of 8 hours, about 6 hours will be call recorded (75%)

Based on factors described in sections below:

Maximum number of agents = 200

Number of concurrent analysis sessions

The Azure Cloud API Service has a limit of 100 concurrent API sessions per customer. A recorded conversation results in two files to be processed per agent. Each file being processed consumes an API session. This results in a limit of 50 concurrent conversations that can be processed. The same limit also applies to the MIR Core server.

Transcription Rate

This is a measure of the time it takes a single session to convert an audio file into text. This results in the following limits:

24 hours of recording are transcribed in 24 hours per session. Two sessions are needed per agent conversation.

Analysis time may improve on this value when there are larger periods of silence in the recordings, but this cannot be guaranteed or quantified, in advance.

File transfer

The files should be recorded and stored on MIR in G.711 (8 bits) for best quality and analysis. However, Azure works with linear PCM (16bits). The files are converted locally within MIR to this format before transmission to Azure. This also has the impact of doubling the payload data size and therefore egress bandwidth from the customer installation.

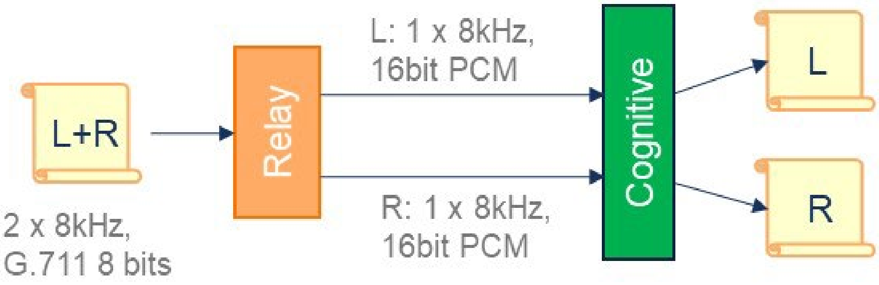

The recorded conversation files are split into left and right (both sides of the conversation) channels. The conversations are analysed based on the two sides of conversation resulting in two analysis sessions in parallel. The results of the transcription are returned as basic text files (2 files per conversation), including timestamp information. This is then appended to the MIR storage as additional meta-data to enable faster searching.

Bandwidth and Data Egress

These limits impact the infrastructure and may result in additional costs to the installation. The Transcription Analysis requires that data is transferred from the customer to Azure (egress at customer) and after analysis, from Azure to the customer (egress at Azure). The bandwidth is the minimum value required to ensure timely transfer of this data.

The amount of data transferred may also need to be considered at the egress points.

The following per agent values are based on call handling rate for agents, highlighted earlier:

Bytes per month per agent to Azure	16	GBytes
Bytes per month per agent from Azure	0.020	GBytes
Bandwidth per agent	300	kbits/s

Each agent is involved in a conversation, and this results in two files (designated left and right on a stereo recording) and 2 API/upload sessions per agent per recorded conversation file.

The bandwidth requirement is primarily over the Internet and is *additional* to any other bandwidth considered for operation of the installation, for example, bandwidth provisioned for Teleworker users.

Network Latency

Network latency between the customer and the assigned Azure region determine the maximum data-rate that can be transferred. It is recommended that the round trip delay be determined prior to installation. The round trip delay (for example using 'ping') should be less than 100mseconds.

Round Trip Delay customer to/from Azure less than 100milli-seconds.

Cost Considerations

Some cost considerations of the solution include, but not limited to:

- Setting up an Azure account
- Egress bandwidth and data charges from the customer location
- Azure Cognitive Service usage charges (discounts can be applied for guaranteed service use and quantity of files to be processed)
- MIR to Azure connection licenses
- Egress bandwidth and data charges from Azure to the customer

MiCC Outbound

The MiCC Outbound application integrates with the MiVoice Contact Centre Business (MiCC-B) and MiVoice Business(MiVB) to provide an outgoing dialing capability. There are multiple dialing campaigns provided, which provide an automated mechanism to create outgoing calls for agents.

MiCC Outbound has multiple power dialing capabilities and it is the partner's responsibility to ensure that the SIP trunk services associated with the outbound dialer service meet the requirements of the application and that they do not exceed the SIP providers usage policy (or similar).

Please contact your Mitel Account Team if there is a requirement to deploy this solution and application, as the installation requires professional services assistance.

Architecture

The solution consists of the following key components:

- Dialler Application
- NVP Dialling Engine
- Web Server
- Reporting Server (SSRS)
- SQL Database (and replication)

These components and servers can be deployed as separate units or combined for smaller deployments.

The SQL servers are kept separate from the other servers to minimize license costs. The reporting service is associated with database and provided on that server. For larger deployments, the database is replicated to allow the reporting service to work with a non-live database. These servers are defined as:

- SQLSSRS = SQL Database + SSRS (Reporting Service)
- SQLPlus = SQL Database plus Replicated database + SSRS (Reporting Service)

The key inputs that drive the quantity and type of servers, are simply the number of active agents. The deployment configurations are shown in the section below. The maximum number of agents supported with the defined architectures is 500.

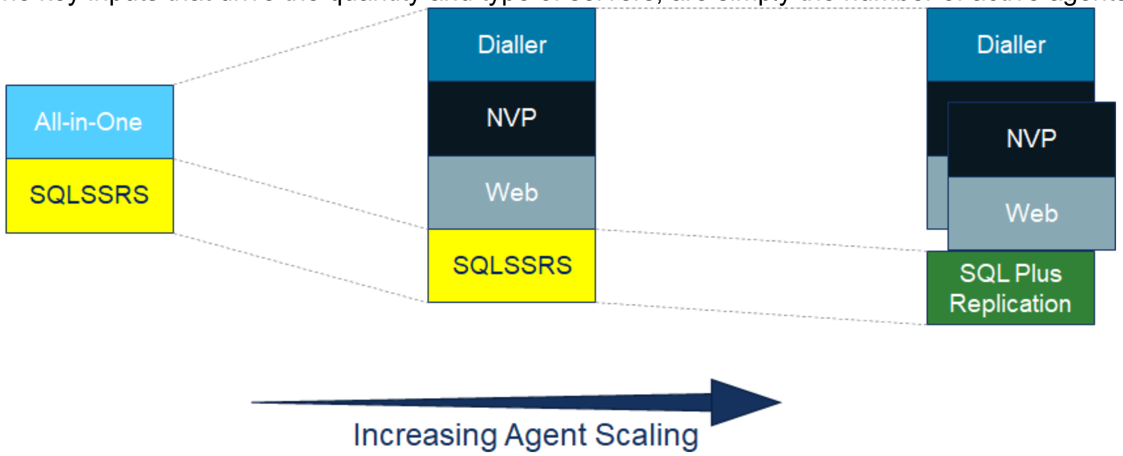

Scaling is achieved by scaling the base servers, and at larger sizes by providing multiples of NVP and Web servers.

Application Boundaries and Deployment Considerations

There are a number of internal limits that define the server sizing and quantities. This information is included in CPQ. The deployment configuration section outlines the outcome of these limits.

Defaults

It is assumed that the agents already exist, when this solution is deployed in a contact centre installation. When integrated with the MiCC Business application, it is assumed that the trunk capacity already exists for the agents, since they are using this in preference to handling incoming calls. There may be a requirement to dedicate certain trunk lines for incoming only and outgoing only depending on the deployment type. If trunk call recording is used, then this split in trunk usage will have to be considered, for example, there are usually more trunks on a system than agents for incoming calls due to the availability of IVR (virtual agents). For outgoing calls, trunks only need to be allocated per agent. If the agents are also not 100% busy making outgoing calls, then it is also possible to scale back the number of outgoing trunks as these can be shared with other agents dealing with outgoing calls. For a simple analysis, it is therefore assumed that trunks andagents already exist, and existing numbers can be applied.

Trunk to Agent Dialling Ra*tios*

The MiVoice Business has a soft limit of 2000 SIP trunks. The preferred deployment option (see below under Dialler Configurations) is Option 2 using SIP trunks to connect to the agents and also to external SIP trunks. Although the application can generate up to 8 calls per agent, the practical upper limit is around 4 calls per agent, or less. A ratio of 2:1 is recommended

In a steady state condition, each active agent will consume up to 3 SIP Trunk licenses. However, in order to carry out predictive dialing for these agents, additional SIP Trunk licenses are needed. Ensure that these are added to the Sales tool output if this function is required.

The following diagram indicates the call and media flow, the media loopback within the MiCC Outbound application, and where the SIP Trunk licenses are consumed:

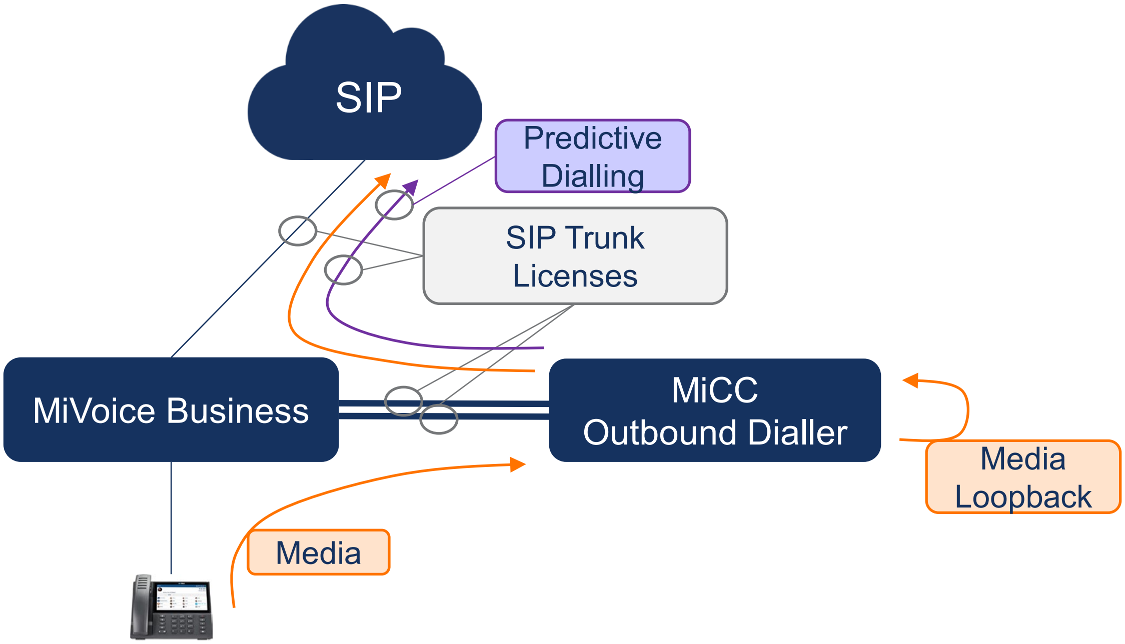

The following table indicates the number of SIP Trunk Licenses required for different number of agents, against the recommended dialing ratio of 2:1 (2 trunks dialing per agent, or 1 trunk active and completing with a secondary trunk in dialing phase). By default, each agent is assigned three free SIP Trunk licenses within the Sales Tool (CPQ). Additional SIP Trunk licenses are required to include the additional connections for predictive dialing:

Quantity of Agents	Predictive Dialing Ratio	Quantity of SIP Trunk Licenses Required	SIP Trunk Licenses	Additional SIP Trunk

			Included with CPQ	**Licenses Required**
100	2:1	400	300	100
200	2:1	800	600	200
300	2:1	1200	900	300
400	2:1	1600	1200	400
500	2:1	2000	1500	500

Should a different predictive dialing ratio be required, contact your Mitel Account Team.

Maximum Number of Agents

The current limit on the number of agents is 500. If a larger solution is required contact your Mitel Account Team for assistance.

Trunk allocation

Outbound dialing is achieved via SIP trunks to the MiVB, which then creates an outgoing call to the SIP service provider. As noted earlier, trunks are assigned on a 1:1 basis of agents to trunks for streaming and therefore the number of active trunks should be less for an outgoing scenario, than an incoming scenario.

However, with predictive dialing, the number of calls in call-setup may exceed the number of configured trunks. Ensure that the SIP Service provider can support more trunks for call setup than call completion, or the number of outgoing trunks may need to be configured to a different value from the incoming trunks. See also the section above on Trunk to Agent Dialing ratios.

Agent Connectivity

Once an agent logs in to the MICC Outbound, via the Web Application, the MiCC Outbound will generate a call to that agent, and that call will remain in place until the agent logs out. As calls are generated from MiCC Outbound to external trunks, call connection is achieved by routing the media of the agent to the active trunk call within the MiCC Outbound application. The MiCC Outbound application provides the media streaming termination for the external trunks and also the internal agent connection. The application will loop these two connections together. There is no direct streaming from the agent phone to the SIP trunk.

CRM Integration

MiCC Outbound can integrate to a number of CRM applications including on-premises and external Internet connected applications such as Salesforce.

Do Not Call

Identified lists of “do-not-call” can be integrated into the solution as additional files.

Call Recording

MiCC Outbound can integrate with the MIR Call Recording solution. The application will link to the Web API connection on the core server (Server A or Server H). This connection enables meta-data transfer related to the call to be recorded, as well as providing delineation between calls, when the agent is on a continuous connection to MiVB. This ensures that different calls from the same agent are identified, and the recordings separated, even though the agent is on a continuous connection.

The Call Recording that is associated with the MiCC Outbound is not used with this deployment, the server resource configurations do not include this function, nor are the possible call storage and sizes defined. Refer to Mitel Interactive Recording for Call Recording.

Resiliency and Disaster Recovery

The MiCC Outbound application is not considered voice critical and is not delivered with a resilient solution. It is deployed as a single data centre deployment.

The deployment uses the HA capabilities for rapid recovery within a single data centre. Where a deployment requires additional levels of availability this should be deployed with the DRaaS configuration between two data centres.

Dialling Campaigns

The MiCC Outbound Dialler provides the following campaigns:

- Preview: Agent requests the next CRM record and classify the calls
- Progressive: Agent is idle for a period before system "pops" and dials the next CRM record
- Power: As the agent becomes free, the system dials the next CRM record
- Predictive: System pre-dials a CRM record in the expected availability of the agent

The dialer can also monitor calls to determine if these are answered by a real person, versus a dead-call, or a FAX machine, or even an answering machine. Calls are dialed by the application and then transferred to the agent to complete.

MiCC Outbound Infrastructure Configurations

The predictive dialer can be configured in the following ways:

1. Integrated with MiVoice Business using MiTai. In this mode the dialer connects to the MiVoice Business as a number of IP5020 phones and dials out to the trunks and also to the agents using these phones. The MiCC Outbound does not include trusted device licenses, and therefore requires twice the number of devices licenses than agents, as the dialer is effectively creating two calls per connection and routing the media locally. This option is available for legacy deployments and existing on-premises deployments. This is not a preferred option for new deployments.
2. Integrated with MiVoice Business via SIP trunks. In this mode the dialer terminates trunk calls from the agent, and also generates tandem SIP trunks calls via the MiVoice Business. This requires a number of SIP trunk licenses (zero cost available for this application). This requires three times the number of licenses as concurrent agent licenses. Advanced dialer licenses are also required. This is the preferred option for all new deployments, including cloud deployments.

For cloud deployments, Option 2) is the preferred deployment model, and the sales tools (CPQ) are geared up around this deployment

Licenses

A summary of the licenses needed for the deployment can be found in the Flex 5.1 General Information Guide (GIG) and/or Ordering Guide on Doc Center.

Licenses are deployed on a concurrent active agent basis.

Microsoft SQL licenses are required for the SQL servers, and the Standard SQL server is recommended. See Microsoft, or Microsoft dealer for specific pricing and configurations.

Microsoft Windows Versions

The applications can run on Windows Server 2016 and Windows Server 2019. The later version is preferred, and also includes a level of anti-virus (Defender), not provided by default in Server 2016.

Server Configurations

The following server configurations are used to scale the deployment from 1 agent to 500 agents. The table below highlights the different server types and the scaling of those servers.

As part of the server definitions, the combined value for the HDD storage is defined. This storage is used for different functions. The storage needs to be partitioned for the different functions and also to ensure that storage from one function doesn't bleed into other partitioned. Apply storage alarms to the partitions. In order to take advantage of the HA capability of the underlying infrastructure, the storage should not be associated with the physical server, but rather with a networked storage, such as SAN drive. This is also a requirement for Disaster Recovery, where the stored data will be replicated to the secondary data centre, such as via Zerto.

In certain cloud platforms, multiple virtual disks may be deployed, especially with SQL servers where a different memory storage type may offer operational improvements. Typically, in this case the operating system and application would exist on one drive with the C: partition, with the other virtual disk being partitioned with the other functions, for example, D:, E:, etc. The table below will help to identify those different partitions.

The following table provides recommendations on how to allocate the different functional partitions to the available storage.

SKU	SKU Description	vCPU	RAM	HDD	Server Description	C:	D:	E:	F:	G:	H:	I:
						Operating System	Application, SQL	SQL Logs	Replication Cache	Replicated SQL Data	Replication Logs	SSRS Dataand Logs
54011245	CC-N App Web NVPServer 16Max Infra	8	12	190	All-in-One Small	60	130					
54011246	CC-N App Web NVPServer 30Max Infra	10	12	210	All-in-One Medium	60	150					
54011248	CC-N App Server 60Max Infra	4	8	120	DiallerSmall	60	60					
54011250	CC-N App Server 120Max Infra	6	8	150	Dialler Mid-Small	60	90					
54011471	CC-N App Server 240Max Infra	8	8	210	Dialler Medium	60	150					
54011477	CC-N App Server 500Max Infra	12	16	340	DiallerLarge	60	280					
54011249	CC-N NVP Server 60Max Infra	4	6	120	NVP Small	60	60					
54011469	CC-N NVP Server 120Max Infra	6	6	150	NVP Medium	60	90					
54011475	CC-N NVP Server 300Max Infra	8	8	230	NVP Large	60	170					
54011466	CC-N Web Server 60 Max Infra	4	8	160	Web Small	60	100					
54011468	CC-N Web Server 120 Max Infra	6	8	190	Web Medium	60	130					
54011473	CC-N Web Server 400 Max Infra	8	12	320	Web Large	60	260					
54011467	CC-N SQL Server 60 Max Infra	4	14	190	SQLSSRS	60	70	40				20
54011251	CC-N SQL Server 120 Max Infra	6	24	440	SQLPLus	60	100	60	40	100	60	20
54011472	CC-N SQL Server 240 Max Infra	6	30	640	SQLPlus	60	160	90	60	160	90	20
54011478	CC-N SQL Server 500 Max Infra	12	54	1060	SQLPlus	60	290	150	100	290	150	20

The SQL server "CC-N SQL Server 60 Max Infra" (SKU: 54011467) does not include SQL replication, asthe reporting service runs directly on the live SQL database. This server is designated as "SQLSSRS".

All other SQL servers include replication and are designated as "SQLPlus".

Although there are a number of logical partitions identified for the different servers, it may not always be efficient to allocate one logical partition to one particular disk drive, or disk volume, for example, disk volumes may not be available in these specific sizes. Where the server specification is primarily the operating system (drive C:) and the application (drive D), and maybe some smaller data partitions, the preference is to deploy a single disk, or volume, and divide this into the relevant partitions. However, where volumes can match the required storage size at equal cost, it is still valid technically to allocate one partition per logical drive.

For the larger SQL servers, a larger quantity of data needs to be handled, and performance is key. In this case, the preference is to separate this data onto a separate disk, or volume, separate from the application and operating system. Ideally this active data should be stored on a disk volume with higher throughput and IOPS, and SSD is recommended for this usage.

In many cases the disk volume may only be available in predetermined sizes. The values shown above are considered minimum values for these servers to operate correctly. Choose a disk volume that is at least larger than the sum of the partitions and spread out any additional capacity. The database and logs are especially storage hungry.

Deployment Configurations

As this deployment is not voice critical, it is only deployed as a single data centre deployment and is not resilient. The solution can be deployed as part of a Disaster Recovery, using Zerto. In this case the databases need to be replicated into the secondary data centre so that they can be recovered in the event of a disaster.

The deployment 'steps' and server definitions are provided below:

Single Data Centre

Non-Resilient Deployment, Single Data Centre						
Agents	**Server Description**	**Server Size**	**Quantity**	**vCPU**	**RAM**	**HDD**
16	All-in-One	Small	1	8	12	190
	SQLSSRS	Small	1	4	14	190
30	All-in-One	Medium	1	10	12	210
	SQLSSRS	Small	1	4	14	190
60	Dialler Application	Small	1	4	8	120
	NVP	Small	1	4	6	120
	Web	Small	1	4	8	160
	SQLSSRS	Small	1	4	14	190
120	Dialler Application	Mid-Small	1	6	8	150
	NVP	Medium	1	6	6	150
	Web	Medium	1	6	8	190
	SQLPlus	Mid-Small	1	6	24	440
	Dialler Application	Medium	1	6	8	210

Non-Resilient Deployment, Single Data Centre						
240	NVP	Large	1	6	8	230
	Web	Large	1	6	12	320
	SQLPlus	Medium	1	6	30	640
500	Dialler Application	Large	1	12	16	340
	NVP	Large	2	8	8	230
	Web	Large	2	8	12	320
	SQLPlus	Large	1	12	54	1060

Refer to the Server Definitions to cross reference the server descriptions and SKU.

Disaster Recovery (Zerto)

Non-Resilient Deployment, Single Data Centre						
Agents	**Server Description**	**Server Size**	**Quantity**	**vCPU**	**RAM**	**HDD**
16	All-in-One	Small	1	8	12	190
	SQLSSRS	Small	1	4	14	190
	Secondary Storage Replication (SQL)					190
30	All-in-One	Medium	1	10	12	210
	SQLSSRS	Small	1	4	14	190
	Secondary Storage Replication (SQL)					190
60	Dialler Application	Small	1	4	8	120
	NVP	Small	1	4	6	120
	Web	Small	1	4	8	160
	SQLSSRS	Small	1	4	14	190
	Secondary Storage Replication (SQL)					190
120	Dialler Application	Mid-Small	1	6	8	150
	NVP	Medium	1	6	6	150
	Web	Medium	1	6	8	190
	SQLPlus	Mid- Small	1	6	24	440
	Secondary Storage Replication (SQL)					440
240	Dialler Application	Medium	1	6	8	210
	NVP	Large	1	6	8	230

Non-Resilient Deployment, Single Data Centre						
Agents	Server Description	Server Size	Quantity	vCPU	RAM	HDD
	Web	Large	1	6	12	320
	SQLPlus	Medium	1	6	30	640
	SecondaryStorage Replication(SQL)					640
500	Dialler Application	Large	1	12	16	340
	NVP	Large	2	8	8	230
	Web	Large	2	8	12	320
	SQLPlus	Large	1	12	54	1060
	SecondaryStorage Replication(SQL)					1060

Refer to the Server Definitions to cross reference the server descriptions and SKU.

Basic Deployment Rules

- Keep SQL as a separate server (minimize cost) and include the SSRS function with the database.Microsoft SQL licenses are needed in addition to server licenses
- All-in-One includes Dialler, API/Web and NVP. All-in-one caters for up to, and including, 32 agents.Above this, NVP has to be broken out
- SQLSSRS = SQL + SSRS in a single server.
- SQLPlus = SQL Plus Replication + SSRS. Replication is needed to minimize DB access and blocking.Replication is on the SQL server to minimize SQL licenses and costs. SSRS is part of the SQL and Replication and therefore part of this server
- SQLSSRS caters for deployments up to, and including, 64 agents. Above this, SQLPlus is needed
- NVP scales up to, and including, 300 agents. Above this, multiple NVP servers are needed
- Web server scales up to, and including, 400 agents. Above this, multiple web servers are needed
- Anything more than 500 agents requires a customer deployment and approval. Please contact yourMitel Account Team.
- It is assumed that there are 1 supervisor per 25 agents, and supervisors also count as an agent froman application and operational view.

Additional Pre-Requisite Information

Additional information on the deployment and prerequisites for that deployment can also be found in the 'MiCC Outbound Prerequisites' document on Mitel Doc Center. Note that this document applies to all MiCC Outbound deployments and information may not be applicable in all deployments. Information that is relevant and different is highlighted in this section.

IP Ports

Most of the connection are local to the solution, although some connections are needed to web services, such as the web server for control, management, dashboard and agent notification. There is also more dedicated and functional client that might be used by supervisors or managers. Currently the recommendation is to only use this client on-net, or when connected to the private network via a VPN. It is not recommended for connections over the Internet. The web service is currently not available via the MBG, and this is also currently restricted to on-net services.

Internal IP Connections

These are IP Ports that are on the customer private network. The following ports need to be opened, in order to ensure correct operation of the solution to and from the customer network:

Port	Function
80	Web Access (Unsecured http://)
81	Web based applications
443	Web Access (Secured https://)
8900	Web Services Communication
8901	Web Services Communication

External IP Connections

Connection to the web server is also available via the MiVoice Border Gateway web-proxy. The following ports are required:

Port	TCP/UDP	Direction	Description
443	TCP	To webserver from remote worker	Access to web server via MiVoice Border Gateway web-proxy

Access is provided to the following functions:

- Live Monitor/Dashboard
- NVP Monitor
- Strategy (Route Manager) and Campaign Management

Access is not available for MiCC Outbound Interaction Studio, which still requires an on-net connection.

Training Material and Professional Services

The deployment and integration of the MiCC Outbound solution is quite complex. A good understanding of the product and essential components is a requirement to a successful deployment. It is strongly recommended to take training on the product and how to integrate this. Use of Professional Services for installation is a requirement. Training is available through the Mitel Learning Management System (LMS).

Alternative Deployments

Can't seem to get a deployment to meet your requirements? Maybe you've exceeded some of the boundaries? Not sure where to go to next? Please contact your Mitel Account team for further advice.

Mitel Revolution

Mitel Revolution is Mitel’s next generation mass notification solution that was built for today’s modern organization to improve network-wide communications, security, and emergency responsiveness across the organization.

With Mitel Revolution, organizations can facilitate enterprise-wide communication for virtually any communications need - including real-time and automated notification alerts for emergencies, large scale notification for routine (or non-critical) communications, facility-wide live overhead paging announcements, scheduling of bells/prerecorded announcements, and mobile-centric communications for mobile employees, students, or others who registered using its self-service portal.

Mitel Revolution provides a way to centrally manage creating and sending notifications. Creating notifications involve three main steps:

- Assign the triggers that send notifications.
- Create the content (image, audio, text) to be sent.
- Assign the endpoints and contacts to receive the notifications.
- Collect responses, if desired, and set up post notification activities including conferences and lockdowns

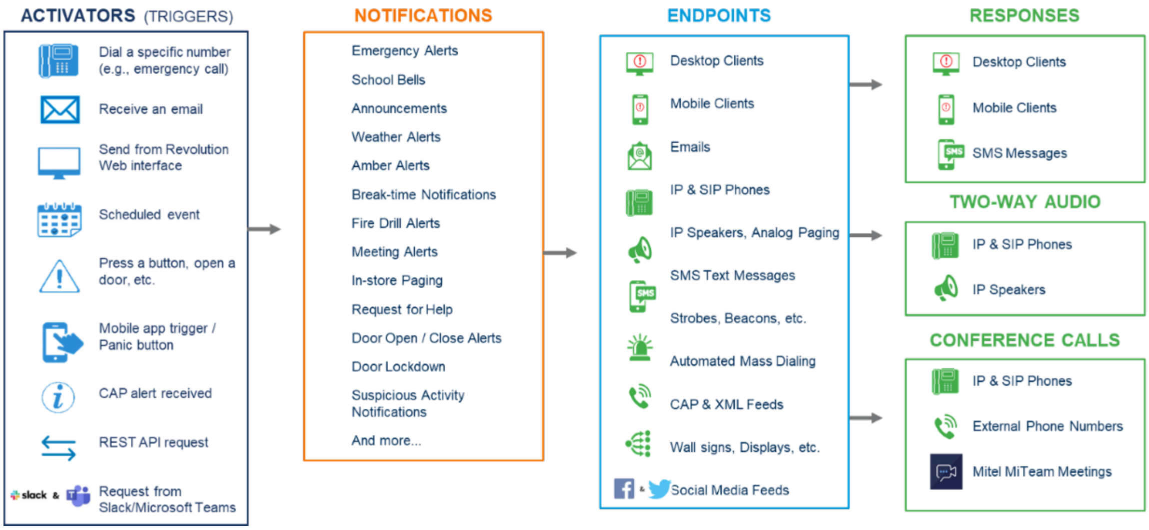

Mitel Revolution can be particularly important for emergency calling, where someone within the business has made an external emergency call. This allows the call to provide notifications, say to internal business emergency teams, or personnel, and can also provide notification for someone to guide emergency services that arrive at the front door to the location of the emergency.

This trigger is provided by the MiVoice Business in the MiCloud Flex solution to the Mitel Revolution server. The MiVoice Business and Revolution server must be on the same network, that is if they are physically in separate locations, they can still be reached through the private WAN connection. This is not an OTT service.

Further information on deployment, requirements and server resources can be found on Mitel Doc Center.

Powering connections