MiCloud Flex (Google Cloud)

Solution and Engineering Guidelines

February 2022

NOTICE

The information contained in this document is believed to be accurate in all respects but is not warranted by Mitel Networks™ Corporation (MITEL®). The information is subject to change without notice and should not be construed in any way as a commitment by Mitel or any of its affiliates or subsidiaries. Mitel and its affiliates and subsidiaries assume no responsibility for any errors or omissions in this document. Revisions of this document or new editions of it may be issued to incorporate such changes.

TRADEMARKS

Contents

Revision History

Document Version	Document Release Date	Description
1.0	May 15, 2020	Initial release
1.1	July 22, 2020	• Updated the following sections: o Customer MPLS Connections o Phone Support o Feature Availability Comparison between OTT and MPLS Connections o Mitel Interaction Recording Limits (Call and Screen Recording) o Inclusion of Domain for FQDNs o Certificate Management o User and Device limits o MiVoice Border Gateway (MiVBG) Limits o Mitel Speech Analytics Speech Transcription Limits o Mitel Speech Analytic Keyword Spotting Limits o Available Data Centres and Points of Presence • Updated tables under the following sections: o IP Port Listing MPLS Connection to Customer Site o IP Port Listing OTT Connection to Google Cloud o Resiliency and Operations o Feature Availability Comparison between OTT and MPLS Connections • Added the following sections: o Default Generation of FQDN o Custom/Customer or Partner Generation of FQDN
1.2	Oct 29, 2020	Made the following changes for bundle 1.2: • Added the following sections: o WebRTC number of supported clients o WebRTC Clients o Availability with a Single region and Dual Zones o Single Region, Dual Zone Deployment o MiCollab Client and DNS SRV records for MiCollab Client under DNS Support o Onsite Gateway Base Kit o Onsite Gateway • Updated the following: o Updated Why Google? section o Updated What is the Google Virtualisation Infrastructure? section o Updated Target Market Segment section o Updated Capabilities section

Document Version	Document Release Date	Description
		o Updated Scalability section o Updated Certificate Management section o Updated the 'Resource Definitions for the UC Bundles' table under Google Cloud Resources. o All graphs and wording under section User and Agent Mix Limits o Updated the OTT port tables and wording under section IP Port IP Port Listing OTT Connection to Google Cloud o Updated devices table under Feature Availability Comparison between OTT and MPLS Connections section. SIP DECT 112 is resilient but has limited characters for FQDN. o Updated diagram and table under Available Data Centres and Points of Presence section o Updated diagram under Dual Region Deployment o Updated Call Signaling and Media Routing section o Updated diagram under Resiliency and Operations o Updated Overlay of Applications section o Updated Capabilities and Changes from Previous Deployments section o Updated Phone Support section o Updated System and Application Availability section o Updated Applications of Flex on Google Cloud section o Updated Call and Screen Recording (WFO MIR) section o Updated MiVB Networking section o Updated table under Resiliency and Operations section o Updated MiVB Networking section o Updated DNS Support section o Updated Connections to SIP Trunk Providers section o Updated the "Resource Definitions for the Call Recording Bundles" for license bundles 54010625 and 54010626, under Call Recording Bundles section o Updated IP Port Listing MPLS Connection to Customer Site and IP Port Listing OTT Connection to Google Cloud sections. Additional ports (3999 and 3998) added to allow connection from 5300 phones to MiVB and MBG for teleworker.
1.3	Jan 20, 2021	o Updated the reference information link in the section, MiCollab Client and DNS SRV Records.
1.4	Feb 21, 2022	o Updated resource definitions for the UC bundles in section Google Cloud Resources.

Solution Overview

This section provides an overview of the MiCloud Flex hosted solution when deployed on Google Cloud public compute infrastructure. More detailed information on specifics of the deployment can be found under the Engineering Guidelines section of this document.

Why Google?

Google Cloud provides a number of deployment options and models, ranging from:

- Google Compute Engine (GCE - Virtual Machines)
- Google Kubernetes Engine (GKE - Virtual Applications deployed in Containers)
- Google Cloud Functions is a serverless Compute Platform for creating scalable micro-services

As such it provides a road-map capability from virtual machines, through containers to micro-services. All architectures can live and operate together, making this a good platform for product migration and evolution. The use of Kubernetes and micro-architectures also makes applications easier to scale, more portable between platforms and faster to develop. This enabled increased innovation and acceleration of time to market for new features.

Google runs a number of secure data centres with a global footprint. Google's backbone network uses advanced software defined networking (SDN) and edge caching services to deliver fast, consistent and scalable performance, providing for global access, or restricted to a single region. The sites are secured and monitored for operation 24/7, within a defined Service Level Objective (SLO). Google also includes certification for privacy and security from a number of recognised compliance standards and recommendations.

What is the Google Virtualisation Infrastructure?

The Google Cloud virtual infrastructure platform contains two fundamental functional engines. These are:

- Google Compute Engine(GCE)
- Google Kubernetes Engine (GKE)

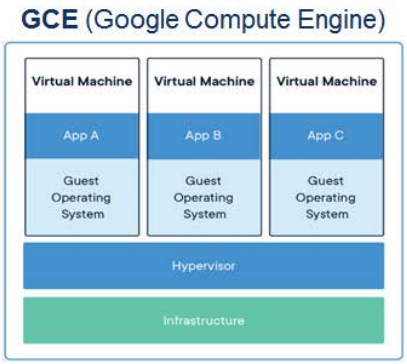

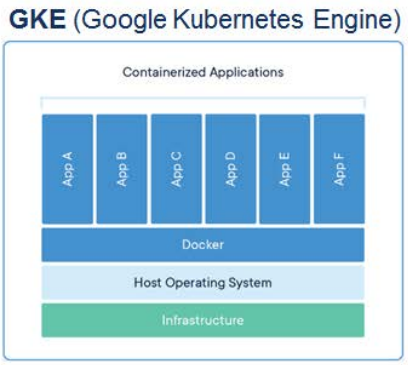

Google Compute Engine (GCE)

This is a virtualisation infrastructure using a common proprietary (Google) hypervisor to run a number of virtual machines and operating systems. The virtualisation is at the server level, or hardware level, and allows the applications, and operating systems, to function as though they were on a single server platform. This is similar in operation to many other virtualisation platforms such as VMWare and Hyper-V. Each application runs within a defined resource boundary for CPU, memory and disk storage.

Some key aspects of GCE include:

- Each VM runs in its own OS
- Hardware Level Virtualisation
- Scaling in distinct blocks
- Resource pre-allocation
- Ability to run existing applications on existing Operating System (OS)

The advantages of such an infrastructure are improved efficiency over a single server deployment. The applications appear to run as though they were on a dedicated server. It does allow multiple applications and operating systems, at different revision levels to exist on the same server. This is in line with other virtual machine technologies.

However, the requirement for a unique operating system with each application per virtual machine, increases the resource footprint when the operating system is repeated. Any outages typically involve recovery of the entire virtual machine, and usually involve a reasonable recovery time, because both operating system and applications need to be started together.

Google Kubernetes Engine (GKE) - Containers

Google GKE is an enterprise-grade platform for containerised application. All applications in the Kubernetes environment share this common operating system. This improves efficiencies by not duplicating code and resources. Leveraging GKE, it is possible to refactor applications into smaller functional blocks and have these run as individual decoupled components. Effectively these smaller components become functional containers. The architecture now allows for the dynamic scaling of these smaller functional blocks and gain performance improvements. Smaller applications, or containers, also contain less code, start faster and can be updated independently of other applications or containers. This provides for better dynamic scaling and maintaining up-time through product release levels.

Some key aspects of GKE include:

- Common host Operating System
- Rapid Startup time
- Application Virtualisation
- Dynamic Scaling
- Resource efficiency ($)
- Consistent behaviour through use of common Operating System
- More rapid and targeted updates
- Portability

- Enhanced Security through reduction of attack surface per component

The advantages of GKE environment include more rapid startup following an outage, more efficient use of resources, as well as consistent behaviour of the operating system between applications (common OS). Containers, like applications, can be started and stopped independent of the underlying operating system, and upgraded independently of the operating system. Having only one common operating system also allows a smaller resource footprint.

However, when this common operating system requires to be upgraded, all applications and containers on that node or virtual machine need to be stopped and started. This is the same as for any virtual machine deployment. Currently, production ready Kubernetes from Google supports Linux (Ubuntu based).

Target Market Segment

The target market for MiCloud Flex on Google Cloud is:

- Generally, customers of 100 users upwards; however, the solution will scale below this as well
- Customers that require a more bespoke deployment
- Customers that want to customise their deployments with a level of self service
- Customers that require optional capabilities including contact centres, and call/screen recording
- Customers that are highly mobile and may use multiple devices
- Deployment over dedicated network connections, such as MPLS, Open Internet (Over the Top (OTT) delivery) connections, as well as a hybrid mix of the two
- Deployment as a dedicated customer instance, and not shared with others. Some cloud services may be shared.

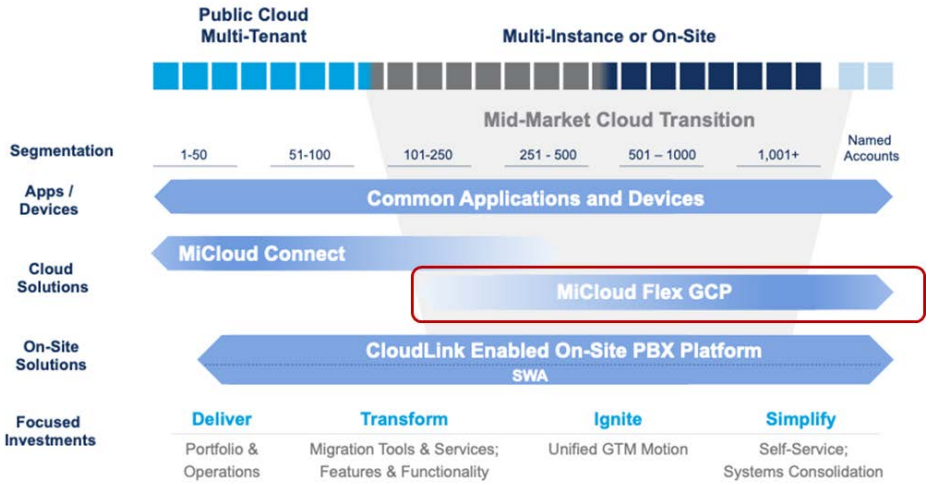

The current identified target markets for this release are:

- Belgium

- France
- The Netherlands
- UK
- USA

Capabilities and Changes from Previous Deployments

The following capabilities and changes are present in this deployment release:

Capabilities

- Rich Voice Services
- Desktop and Mobile UC
- Unified Messaging (Voicemail, Fwd to Email, Visual Voicemail)
- Audio/Web/Video (AWV) Conferencing
- Team Collaboration
- Site/Department/Team Auto Attendant Call Flow
- Mitel Workforce Optimization (WFO)
    - Mitel Interaction Recording (MIR) / Call and Screen Recording
        - MIR provides redundant Call Recording with a default storage duration of 1 year
        - MIR provides non-redundant Screen Recording with a default storage duration of 3 months
        - Additional storage capacity and user sessions can be purchased as optional add-ons
        - Mitel Quality Management
    - Mitel Speech Analytics
- Mitel Quality Management
- Mitel Teleopti Workforce Management (Contact Centre)
- Business Analytics
- MiCollab Unified Messaging (also known as NuPoint UM) functionality has been replaced with:
    - Voice Mail is provided via the Embedded Voice Mail on MiVoice Business (MiVB), and scaled up to 120 ports and 5000 mailboxes
    - Call Director like functionality is provided via "Call Flow" within MiVoice Business
    - Visual Voice Mail has been ported to MiVoice Business
    - Advanced UM is currently on the roadmap

Scalability

- Up to 2500 UC users (5000 devices) including both OTT and MPLS
- Up to 200 Contact Centre agents

- Solutions for Call Recording, either agent/user (200/2500) or trunk (500) based are included, depending on purchased licenses
- Larger systems may be available subject to Product Line Management and Sales Engineering approval

Operations and Management

- Partner Remote Provisioning Access
- Partner Fault & Performance Monitoring

Network Interconnect

- Customer Interconnect - MPLS via PoP*
- Customer Interconnect - OTT via PoP*
- SIP Carrier Services - OTT via PoP*

(PoP*: Point of Presence)

Redundancy

- Geo-Redundancy, 5 '9's Deployment Dual Data Centre, including Dual PoP connectivity
- Global network Interconnect
- Onsite Gateway

Business Process Automation and Support

- Mitel CPQ integration (Configuration, Pricing, Quoting)
- Automated Service Orchestration Tool
- Mitel Technical Support

Changes from earlier Release

- Introduction into Belgium and the Netherlands markets using existing PoP and Data Centres
- Introduction of resilient remote gateway to be deployed on customer premise
- Introduction of both/either agent or trunk recording
- Scaling improvements to allow simplified provisioning and mix of agents and users with MPLS and OTT
- Increased scaling of WebRTC connections and users

Mitel is leveraging the global presence presented through the use of Google data centres as well as use of existing Points of Presence (PoP). This allows MiCloud Flex on Google Cloud to be deployed on a global basis and to have data centres in different regions, or countries, linked through the Google highly provisioned and low latency private backbone network.

Connectivity

Mitel is leveraging the global presence presented through the use of Google data centres as well as use of existing Points of Presence (PoP). This allows MiCloud Flex on Google Cloud to be

deployed on a global basis and to have data centres in different regions, or countries, linked through the Google highly provisioned and low latency private backbone network.

Available Data Centres and Points of Presence

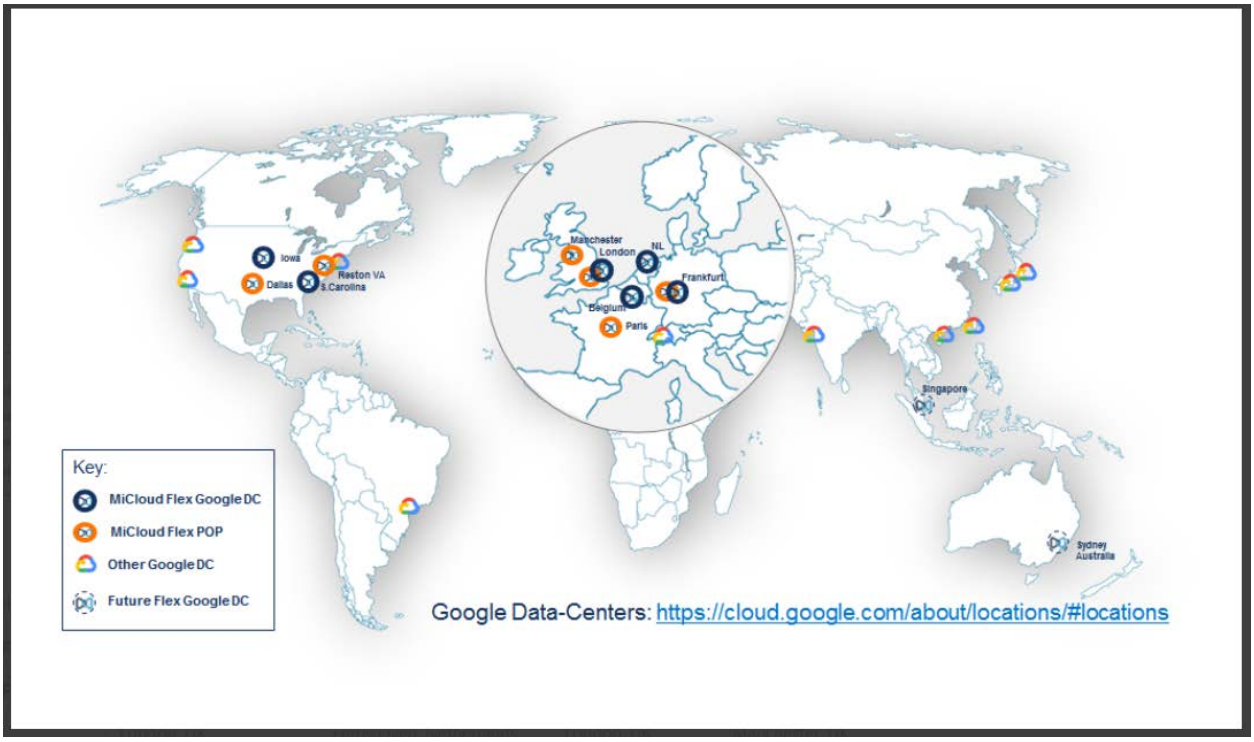

For the initial release the identified Data Centre locations and Points of Presence are:

MiCloud Flex Access	Google DC		Mitel PoP	
MiCloud Flex Region	**Primary Google Region**	**Secondary Google Region**	**Primary PoP**	**Secondary PoP**
UK	London, UK	Eemshaven, Netherlands	London, UK	Manchester, UK
France	St. Ghislain, Belgium	Eemshaven, Netherlands	Paris, France	Frankfurt, Germany
US East	South Carolina, US	Iowa, US	Reston, US	Dallas, US
US Central and West	Iowa, US	South Carolina, US	Dallas, US	Reston, US
Belgium	St. Ghislain, Belgium	Eemshaven, Netherlands	Paris, France	Frankfurt, Germany
Netherlands	Eemshaven, Netherlands	St. Ghislain, Belgium	Frankfurt, Germany	Paris, France

Customer Connectivity

Customers can connect to their MiCloud Flex on Google Cloud via two main methods:

- MPLS Private Network, via Mitel PoP
- Over the Top (OTT) over the public Internet, via Mitel PoP

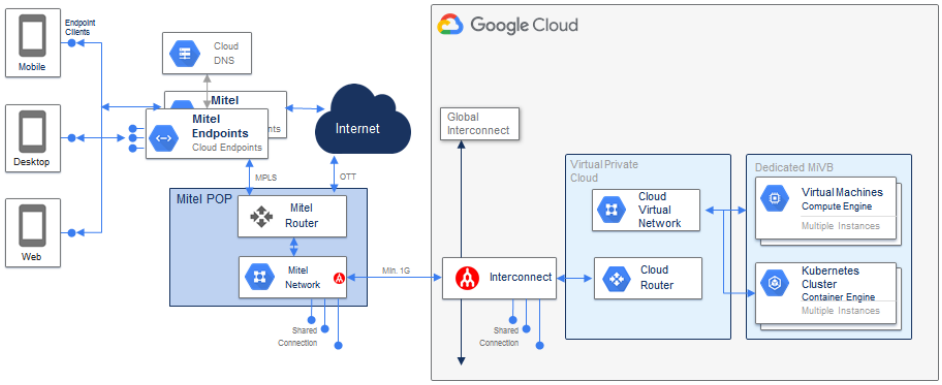

All access to the MiCloud Flex on Google Cloud is via the Mitel PoP, as indicated in the diagram above. This includes:

- User OTT access (teleworker and softphone connections)
- SIP Trunk access (channel partner provided) either OTT, or carrier provided NNI
- MPLS private network access (channel partner provided)

Deployment Architecture

Overlay of Applications

Within Google there are two main compute architectures. As highlighted earlier, these are the GCE and GKE compute environments. Both of these run on a common virtualisation platform, and are connected via a common Virtual Private Cloud (VPC) network. This allows an interconnection between the different compute environments (GCE and GKE), between the multiple PoPs and also between different Google regions.

The GKE is used to provide:

- MiVoice Business (Call Control)
- MiCollab (Collaboration Suite), which supports remote connections to CloudLink (hosted on Amazon Web Services) for Chat sessions, and optional MiTeam Meetings collaboration session. MiCollab AWV is also available for users not utilising MiTeam Meetings.
- MiVoice Border Gateway (Secure Session Border Controller for Teleworkers and SIP trunks) - The gateway may also be used as a Secure Recording Connector (SRC) when the Call Recording Solution is included. This records external user connections. Station side and trunk recording is provided.
- Optional and additional Secure Recording Connector (SRC). This is included in the solution when recording of users or trunks is required.
- Mitel Performance Analytics (MPA) probe, which functions as the secured management portal from the external cloud. This connection also provides performance and analytical data

to the cloud service in Amazon Web Services (AWS). Access by the partner is via the MPA Cloud service in AWS.

- Business Analytics probe, which collects Call Detail Records (CDR) for export to a cloud service (hosted on Microsoft Azure) providing call metrics, through a series of reports, dashboards and wallboards.

The GCE is a virtual environment and is used to provide hosting services for Microsoft Windows Servers for applications, and other applications that are not GKE compatible. These include:

- MiContact Center Business (MiCC-B) and Interactive Voice Response (IVR)
- Microsoft SQL Server (used with MiCC-B)
- Mitel Interaction Recorder: Call Recording Core Server, Computer Telephony Integration (CTI) connections and Recorder functions (voice and screen)
- Speech Analytics and Keyword spotting

Google also offers other functional services, which are used by a number of these applications in both GKE and GCE

- Google Cloud SQL Server is used by MiVoice Border Gateway within GKE
- Google Cloud SQL Server is used by Call Recorder Core Server, within GCE
- Cloud Storage is used by the Call Recording Core Server, within GCE
- Google Cloud Operations/Stackdriver is used by applications within GKE for real-time log management and analysis
- Google CloudDNS is used by applications within GKE, and must also be linked to the customer DNS to allow for customer network phone connectivity and Split-DNS operation
- Google Internal Load Balancers are used to terminate external Internet connections (OTT PoP) from the Mitel PoP for connection to MiVoice Border Gateway, within GKE

In addition to the core services running on the Google Cloud infrastructure, a number of other services, some mentioned above, are running on other cloud services, either Mitel provided or via Cloud partners, specifically:

- While the Mitel Performance Analytics (MPA) probe runs as a container in Google Cloud, the MPA Cloud Service is hosted in Amazon Web Services (AWS). The partner has an account in MPA on AWS which is used to gain access to the designated customer probes. The probe creates a secure tunnel back from Google Cloud to AWS.
- CloudLink Chat for Collaboration and Contact Center Messenger is provided by the Mitel CloudLink service, which is hosted on AWS
- MiTeam Meetings is an optional meeting (video and audio) and collaboration tool which is hosted on AWS. Initial connection is established from MiCollab on Google Cloud, but the main service is via AWS
- Business Analytics and Workforce Management tools are hosted by 3rd parties in the Microsoft Azure Cloud. Connections from the applications in Google Cloud connect to the services in Azure. The customer logs into an account on Azure to access the information.

The following diagram provides is a high-level overview of how some of the applications and functions are deployed:

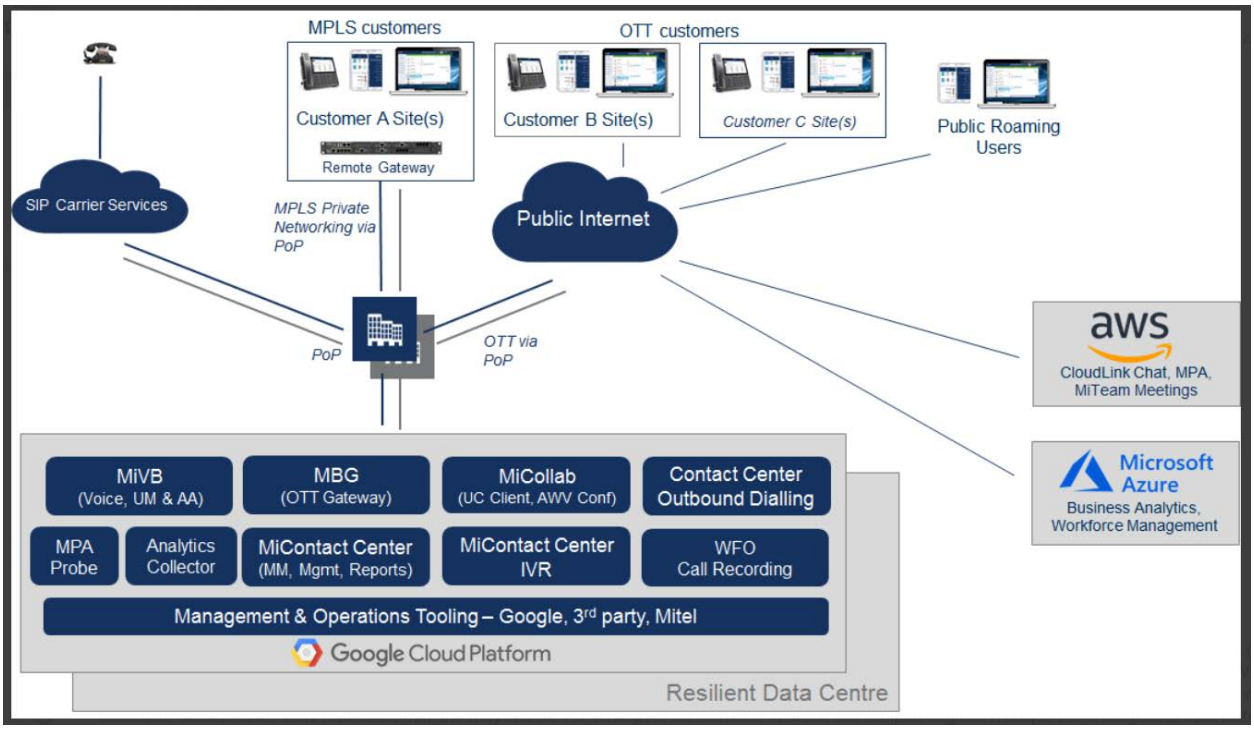

Customer Isolation and Routing via PoP

Individual Customers are deployed as separate Google Cloud projects for isolation from each other. Projects are used for organisation of cloud resources and providing granular Identity and Access Management (IAM). At the network and infrastructure level, customers are separated from each other through the use of different virtual compute engines, networked VLANs, VRF, Virtual Cross Connect (VxC) and Virtual Private Cloud (VPC) configurations. Each MPLS customer is assigned a separate "project" (shown in yellow in the diagram) which is a separate network from other customers. Isolation is therefore maintained from the entry into the PoP through to the individual networks and compute engines. A unique VxC is assigned to each customer and will carry both MPLS and OTT traffic for that specific customer.

For customers that are ONLY using the OTT PoP connections, an alternative solution exists. This allows for a single VxC interconnect connection to be shared by all OTT customers. Customers continue to be deployed in separate projects, using the shared VPC to terminate the OTT VxC connection. Within the shared project, customer network isolation is achieved through the use of the Public IP addresses and firewall rules. Note that some applications and features are only available with an MPLS connection and would not be suitable for this deployment. A view of which features are implicated can be found in the feature comparison list, in the section below.

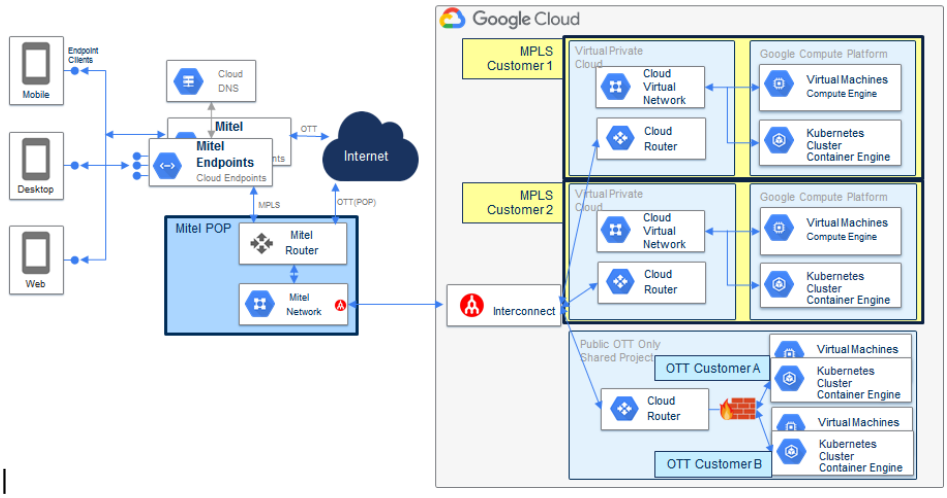

Call Signaling and Media Routing

Within each customer deployment (or project), is a mix of GKE and GCE application, and components. These are linked to a common Virtual Private Cloud (VPC) and this VPC is also peered to the secondary data centre over the Google backplane via the Google Cloud Router. The deployment provides connections to the customer over Private MPLS network as well as a Public connection to a SIP Trunk provider and also Public IP terminations for remote and teleworker users. All other application and components are on the "internal" network, even though this is provided in a Public Cloud infrastructure. See the diagram below for more information. For deployments that are purely OTT, i.e. without the MPLS connection, the deployment is the same, only missing the dedicated connection via the PoP. There may also be limits to certain features and applications. See the feature comparison list in the guidelines below.

The customer devices that are connected to the customer's internal network connect to the service via the MPLS service provider. This connection is terminated in the Mitel PoP, where a VRF and VLAN are assigned to maintain customer to customer and customer to Internet isolation. Phones will then terminate signalling on the MiVoice Business on an internal network address. If the phones are assigned for call recording, they will register through the internal Secure Recording Connector (shown as SRC Int), before the signalling is forwarded to the MiVB. The SRC then forwards the signalling to the MiVB, and any recording media to the Mitel Workforce Optimization (WFO) Mitel Interaction Recording (MIR) Call Recording server(s).

Applications, such as those for contact centre, or for WFO suite of applications, that are on the customer LAN will connect directly to these servers on the internal private IP address.

Phones that are remote, or connect as teleworker phones, connect to the Public IP address presented by the MiVoice Border Gateway (MiVBG). This application terminates the external connection on an external Public IP address, and forwards signalling and media to the internal applications using an internal IP address. If the external device is assigned for call recording, the device connection also has this feature invoked through the external SRC (shown as SRC ext). Connections are then forwarded internally; signalling to the MiVB and recording media to the WFO Call Recording server(s).

SIP Trunk connectivity is also terminated on the external IP address of the MiVBG. The connection to the SIP Service Provider is also terminated on a Public IP address, and is also routed via the Mitel PoP.

The use of components in GKE also means that the IP addresses may change, within defined ranges. In order to counter this, use of DNS is required, and this also needs to linked to the customer DNS. This ensures that phones located on the customer network receive the correct IP address for registration, such as on the Internal SRC or the MiVB.

The diagram illustrates the descriptions from above:

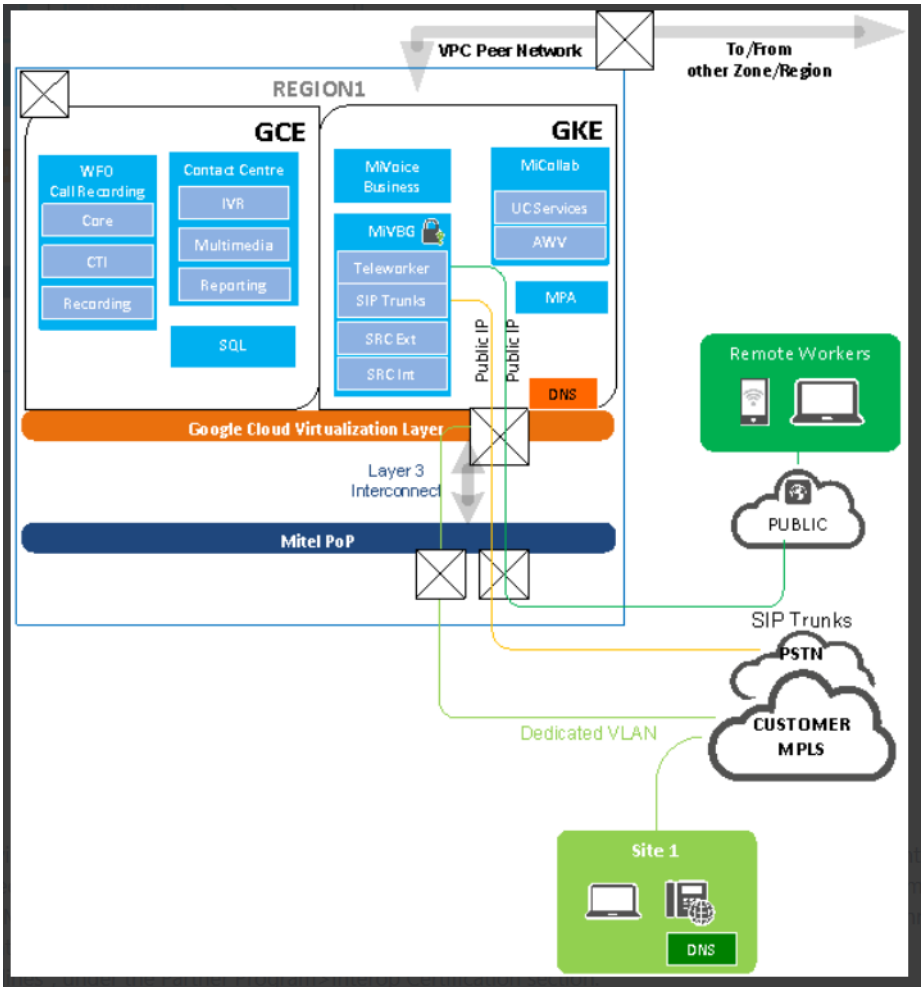

SIP Trunk Connectivity

SIP Trunk connectivity is provided via the Mitel PoP. The partner is required to provide SIP Trunk connectivity. The deployment requirement is that the carrier SIP Trunk SBC presents a Public IP Address. This may be provided either via OTT connection or via an agreed NNI connection. The SIP Trunk configuration information and successful integration for the named SIP Service Provider must also be included in the Mitel SIP Centre of Excellence (CoE) approved list. This list can be accessed under Mitel InfoChannel > Mitel Solutions Alliance >Certified Services pages. Reference should also be made to the "Mitel Interop Certification Validity & Extensibility Guidelines", under the Partner Program>Interop Certification section.

The SIP Trunk provider may be the same as the MPLS provider or may be a different provider. Further details on connection requirements can be found under Connections to SIP Trunk Providers.

System and Application Availability

Google has available a number of global data centres, which are grouped together within a region, or country. Within that region, data centres are also deployed as zones, typically one per zone. Regions therefore contain multiple data centres (zones), and typically there are at least 3 zones per region.

Loads that run within a particular zone may move with that zone within a region. This may occur when a particular data centre and specific servers are targeted for upgrade. This can occur on a regular basis, and typically once per month. Loads that are running will migrate within the zone and continue to operate. Since the zones are running effectively as a layer2 infrastructure, this provides a level of high-availability operation, and allows the IP addresses to be maintained.

However, there is a possibility that all zones within a region may be impacted by a common incident, for example weather conditions, floods, etc. Google can offer a Service Level Objective of 4 '9's (99.99%) when operating within a region. This means that any deployment will at best offer 3 '9's (99.9x%) availability. This reduction comes about because the service will need to be off-line for some period of time, for example software upgrades, and cannot be 100% available. As a result, it is not possible to be better than 4 '9's without changing the architecture. To counter this situation, deployments in multiple regions are used. Deployments that utilise a single region, but dual zones get the availability benefit of parallel operation in two locations, but overall will only achieve up to 4 '9's availability.

In order to improve the system availability, and to counter the threat of a large incident in one region impacting service, a secondary service is deployed in a different region, hence the requirement for multi-region deployments. In this way any outage in one region, can be handled by service from the secondary region. The chances of both regions being impacted at the same time, provides the 5 '9' capability for the solution.

The diagram above shows a number of regions highlighted within a single country, and the number of zones associated within that region.

Deployment across different regions offers the required 5 '9's (99.999%) availability. This availability is offered across the core voice services. It does not apply to all of the applications. For further details look under the section Resiliency and Operations.

Dual Region Deployment

In order to get the required level of resiliency, both a primary and hot-standby secondary system are deployed. An additional requirement is that the primary and secondary systems be deployed within different regions. This ensures that a system is always available, even if one of the regions becomes unavailable.

Google provides connectivity on the data backbone between data centres in different regions and locally between zones within a region. Connections within a region, between zones, can be considered as LAN connected, with little delay between zones. Connections between regions are considered as WAN with routing and may include an element of delay through the cloud router and backbone connection between locations.

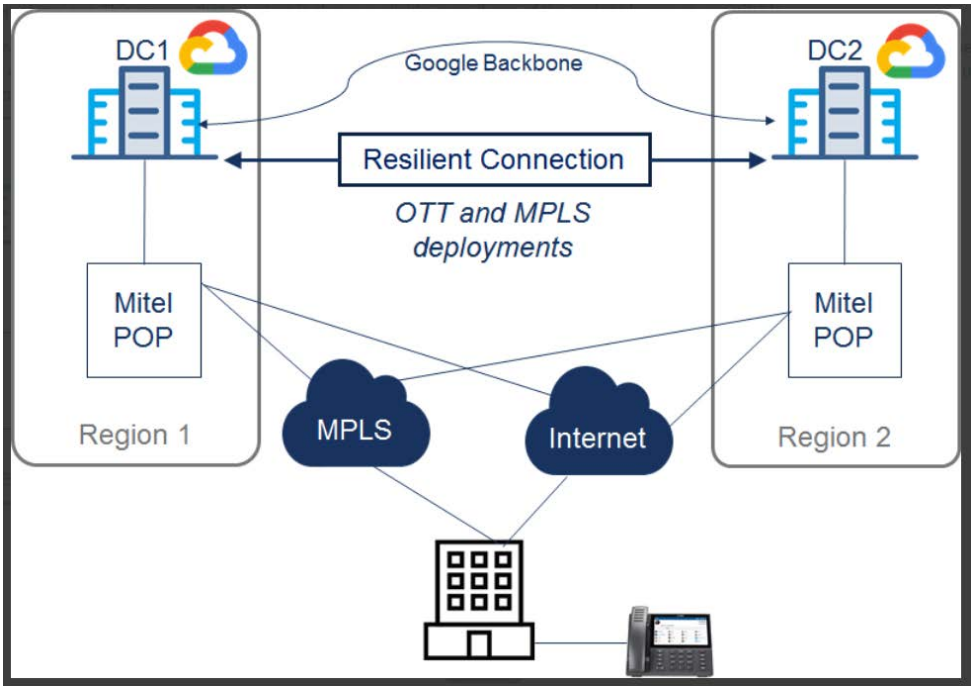

For further details on the resilient configuration, see under Resiliency and Operations.

Single Region, Dual Zone Deployment

In certain locations and geographies, it is not possible to deploy in dual regions. In such a case maximum availability is attained by deploying in two zones within the single available region. Although this will not provide the availability of a dual region deployment, it does improve the availability of the system over a single location deployment. Typically, such a deployment will achieve up to 4 '9's availability. Both a primary and hot-standby secondary system are deployed. This ensures that a system is always available, even if one of the zones within a region becomes unavailable.

Google provides connectivity on the data backbone between data centres in different regions and locally between zones within a region. Connections *within* a region, between zones, can be considered as LAN connected, with little delay between zones. Connections *between* regions are

considered as WAN with routing and may include an element of delay through the cloud router and backbone connection between locations.

Note that PoPs may still be deployed in different regions providing the high availability access connections, even if the internal cloud routing of those PoPs is to a common compute region.

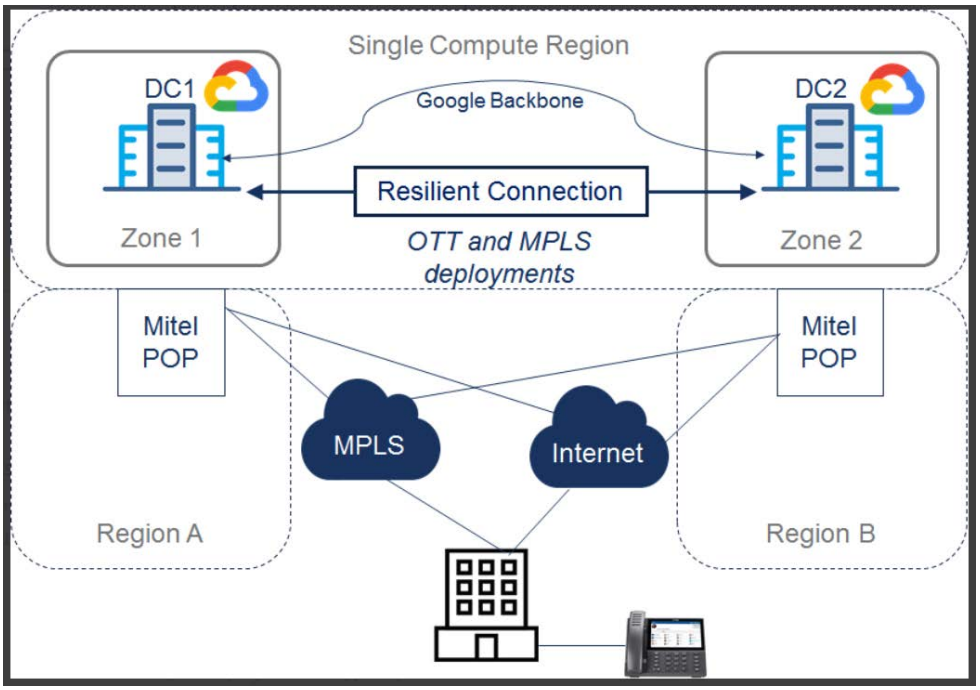

For further details on the resilient configuration, see under Resiliency and Operations.

Infrastructure License Bundles and Deployment Footprint

Infrastructure comes in three main functional blocks, which are further broken into sub functional deployment footprints:

- Base UC services
  - 250 UC Users (250 users/625 devices), also used for 50 agents
  - 500 UC Users (500 users/1250 devices), also used for 100 agents
  - 1000 UC Users (1000 users/2500 devices), also used for 200 agents
  - 2500 UC Users (2500 users/5000 devices), also used for 200 agents
  - 1000 UC users expansion
  - 2500 UC users expansion
- Contact Centre
  - 50 agents. Options include with and without SQL Express, with and without integrated IVR
  - 100 agents, with external Microsoft SQL. Options include with and without integrated IVR
  - 200 agents, with external Microsoft SQL. Options include with and without integrated IVR

- Mitel Workforce Optimization Interaction Recorder
  - 50 agents/250UC Call and Screen Recording
  - 100 agents/500UC Call and Screen Recording
  - 200 agents/1000UC Call and Screen Recording
  - 2500UC Call Recording
  - Additional Cloud Storage packages in blocks of 50 agents for 1, 3 and 5 year extensions of audio, and 50 agents, 3 month extensions for screen recording
  - Call Transcription and Keyword Base package (includes queuing server and transcription/decoder server)
  - Call Transcription and Keyword Server (additional transcription and keyword capacity to add to the base configuration)

Refer to the Engineering Guidelines section below for further details on how these sub-bundles and options are assigned to SKU numbers and resources, along with any associated deployment considerations. CPQ will define which bundles are selected based on customer and partner provided information. This will also look at any mix of users and agents on the base UC bundles to select the most appropriate bundle. See the section on User and Agent Mix Limits for further details.

System Configuration and Management Tools

There are a number of management tools that can be used on the deployment. For the partner, the main tools are CPQ in creating the order, and Mitel Performance Analytics (MPA) to then manage the deployment. Once access via MPA is achieved the system can then be configured through the Initial Configuration Wizard and then individually via the Solution Manager.

Once the CPQ is authorised, an automated Service Orchestration tool builds out the system, and provides a welcome e-mail to the partner to then gain management access to the solution via MPA. The Google cloud console is used by the Mitel operations team to manage the underlying infrastructure and is not accessible to the partner.

The tools and organisation flow for the partner are simplified in the following overview diagram:

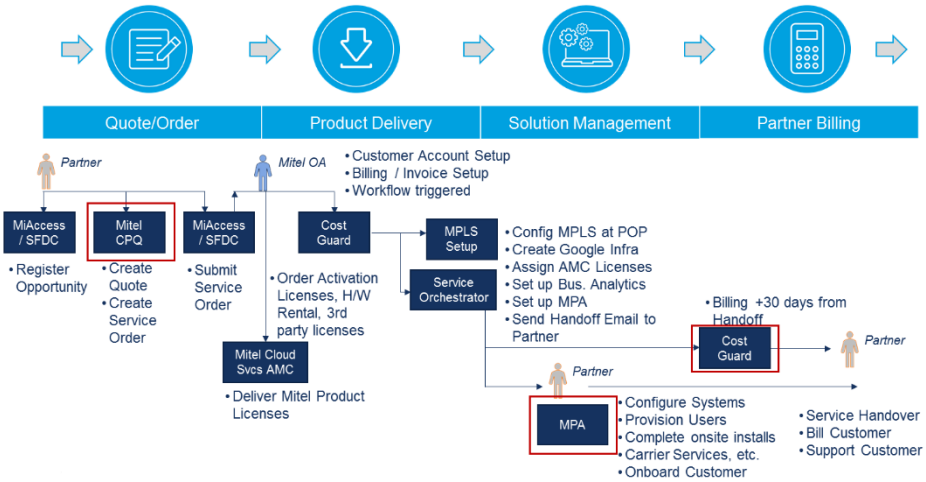

The management flow through to the main deployment is highlighted below:

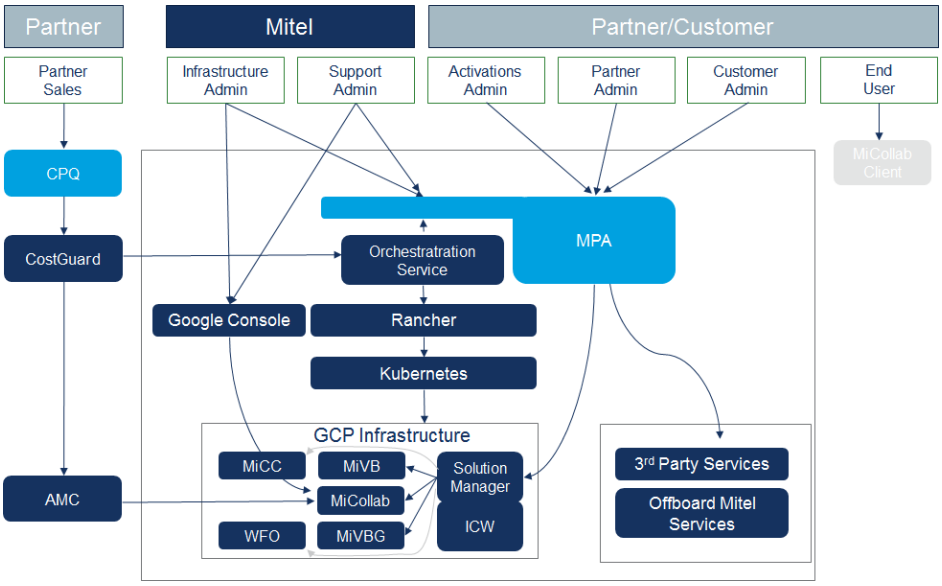

The services highlighted in the lighter blue boxes are those directly accessible by the partner. Access to the applications is through MPA and Solution Manager for the Partner and Customer (if provided).More details on who has access to which part are included in the Shared Responsibility section, below.

Configuration and Price Quoting Tool (CPQ)

CPQ, or Configuration and Price Quoting tool, is used to define the customer requirements and configuration. Initial access to Mitel CPQ is gained via the Mitel MiAccess Portal. Look for the CPQ entry in the table on the left. This tool is accessible by the Partner and is the primary configuration tool to build out the Flex on Google Cloud solution.

To build out a new configuration, enter a customer name and select the appropriate model, in this case 'MiCloud Flex'. Continue to fill out the tabs and finally produce the report. This information is provided as input to the Billing and Operations Support System (OSS) (shown as CostGuard in the diagram), which will be picked up by the automated Service Orchestration tool to deploy the underlying infrastructure and licensing. The automated Service Orchestration operates 'behind-the-scenes' and is not visible nor accessible by the partner.

Service Orchestration Tool

The Service Orchestration tool is an automated tool which uses a number of inputs to create a solution within the Google Cloud Platform, as well as linking the customer's unique application configuration to the Mitel PoP for MPLS connection and Public Internet OTT connections (OTT PoP). This tool is **not** accessible by the Partner. It is an automated infrastructure deployment tool.

The Service Orchestration Tool takes its primary input from the Billing and Operations Support System (shown as CostGuard), where the output from CPQ is delivered. This input includes the configuration information for the deployment, including the different license bundles, number of users and phones, etc. Other inputs include details such as customer name and location, selection of PoP, a specific customer FQDN (Fully Qualified Domain Name), etc. Once the solution is deployed, and e-mail is sent to the Partner, providing remote access details and how to connect to the solution via MPA in order to start with the initial configuration.

The Service Orchestration Tool greatly simplifies and speeds up the provisioning of the deployment. Partners need to be aware that this process takes minutes to complete, and that once the system is deployed it is chargeable and will be billed. Partners should ensure that they are ready for the Initial Configuration to commence before triggering this stage of the deployment.

Mitel Performance Analytics (MPA)

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 utilisation, 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 behaviour.

The Mitel Performance Analytics consists of two components:

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

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.

Solution Manager

The Solution Manager provides access to the Initial Configuration Wizard, as well as access to other application management interfaces. The initial configuration step requires use of the Initial Configuration Wizard. This is a requirement before other configuration and user information is included. Items that can be managed from Solution Manager include:

- Initial Configuration Wizard
- System User Administration
- Time Zone Configuration
- AMC Synchronisation
- Backup and restore
- SNMP Configuration
- MPA Configuration
- View system logs

Initial Configuration Wizard

To simplify the initial configuration of the solution, the Initial Configuration Wizard is made available. This will assist in provisioning users, programming numbers, connecting to SIP trunks across all of the units of the solution. This tool will provide the bulk of the initial configuration to simplify deployment. However, it may not provide all possible configuration combinations. In which case the partner will have access to the Solution Manager and other unit management interfaces.

Some items covered by the Initial Configuration Wizard include:

- Configure Resiliency
- Configure administration services including e-mail
- Configure languages
- Configure numbering plan
- Configure SIP trunking
- Configure incoming and outgoing calls
- Configure optional services including MiCollab AWV, Hot-Desk, Music-on-Hold, etc.

Google Cloud Console

The Google Cloud Console provides operational access to the underlying infrastructure that is running the applications. It allows the Mitel operations team to monitor the status of the services in terms of resource usage and any operational logs. It is not accessible by the Partners, nor does it allow access to the user content on the applications

Shared Responsibility Model

Responsibilities

There are a number of deliverables needed for a customer to obtain service and ongoing updates and maintenance. The deliverables include those provided as part of the MiCloud Flex on Google Cloud (Wholesale) and those provided by the Partner and/or Customer.

The following diagram identifies the split in these deliverables and responsibilities:

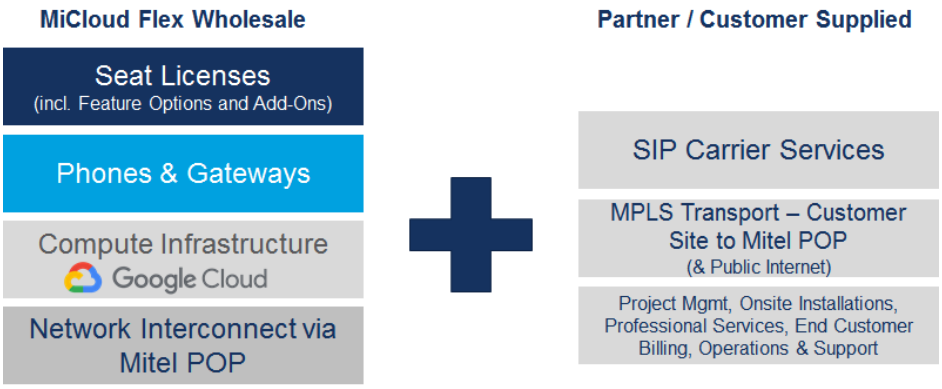

The MiCloud Flex solution includes the underlying infrastructure, the compute resources and allocation, the application licenses, the interconnect to the Mitel PoP, to allow Internet access and connections to MPLS, and any end user devices. The Partner and/or Customer need to provide connections to a SIP Trunk service provider, connections to the MPLS circuit (where used) and connection to the public Internet. The Partner is also responsible for the first levels of support, project management and customer billing. The Partner is also responsible for any professional services, which may include engaging with Mitel Professional Services, where needed.

Solution Definition

In a cloud deployment model, there are a number of key functional areas that are linked together to eventually provide the service to the end customer. There is a level of shared responsibility in making sure that the solution works and is correctly maintained. With the MiCloud Flex solution on Google Cloud there are four main players in the overall solution delivery:

- End Customer and User
- Partner
- Mitel
- Google Cloud

The Solution can be defined in 4 layers, which maps to the 4 different parties involved in the deployment, specifically:

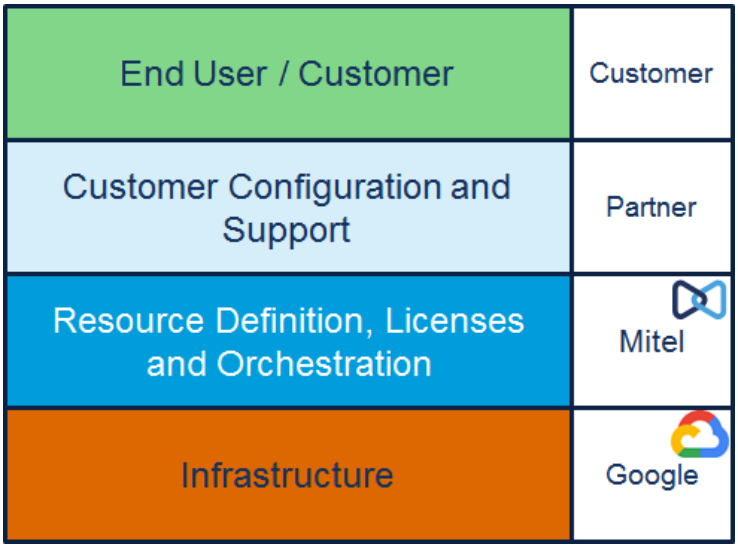

Infrastructure

The underlying infrastructure and compute platforms are provided by the Google Cloud platform. Google provides a number of global data centres in a number of geographic regions. These regions can be interconnected through the Google 'backbone', which allows the applications to be deployed across multiple regions and provide the capability for high availability.

In addition to the Google compute infrastructure, the MiCloud Flex on Google Cloud solution also includes a number of Mitel provided Points-of-Presence (PoP). These PoPs provide connections to the public Internet as well as the option to connect to Partner networks and providers for SIP trunks and MPLS connections.

Google provides the physical security of the data centres as well as the underlying network and compute infrastructure. Google also provides security updates and patches to this infrastructure as well as to the host server operating systems, including Kubernetes.

Resource Definition, Licenses and Orchestration (Mitel Cloud Operations)

The applications, resources and their compute definitions are pre-defined as part of the MiCloud Flex on Google Cloud deployment. These building blocks are predefined by Mitel. Output from CPQ, based on inputs from the partner, define which of these blocks to deploy. The solution deployment uses orchestration tools to automatically provision the resources and configuration of the Google Cloud infrastructure. The orchestration tools will also push the license requirements and configuration onto the Application Management Center (AMC) license server, used by the application to determine features, functionality and number of users and devices on the system. This includes automatic creation of ARIDs and Designated License Manager/Unified License Manager (DLM/ULM) as needed.

Access is provided to the Partner to provide initial configuration and on-going maintenance of the applications for the customer. Mitel will provide application updates and patches as necessary, including security updates. Partners are encouraged to carry out a regular and scheduled application updates in order to roll in these updates. Updates at least once every 6 months is encouraged.

The underlying infrastructure will require regular and scheduled updates, including security updates. From time to time it will be necessary for Mitel to take systems out of service to achieve this infrastructure upgrade. This will be coordinated with the partner and will take place during a defined maintenance window. Primary and Secondary locations will be upgraded during separate window periods to ensure that service can still be provided through the secondary data-centre and gateway.

Customer Configuration and Support (Partner)

Configuration of the applications is provided through MPA and is the primary and secured access portal for the Partner. This allows the applications to be provisioned and configured to meet the end customer requirements. The Partner also provides the connections to the SIP Trunk provider and MPLS provider, who will nominally connect into the Mitel PoP. The partner in consultation with the end customer is responsible for maintaining the E911 SIP Trunk provider location database, where applicable.

The Partner is responsible for the deliver any on-site equipment that is needed and will provide training to the end-customer users. The partner should ensure that Mitel application software updates and patches are applied in a timely and scheduled basis. The initial levels of customer support are provided by the Partner, including Moves/Adds/Changes as well as handling customer license allocation and future option growth.

End Customer

The end customer is the user of the service. The end customer works with the partner to outline the business and system functional requirements. This is normally carried out prior to deployment, and is the main driver into CPQ, which determined the licenses and resources to be used as part of the automated deployment and orchestration. The partner will continue to support the customer as well as managing any option and license growth.

The customer is responsible for meeting and obtaining any solution compliance certifications. Mitel cannot be responsible for the customer network and compliance processes. The infrastructure and applications can be used in specific deployments, it is up to the customer to

ensure that the overall solution, including any customer specific applications and network controls will meet their compliance requirements.

MPLS Provider

The Partner (and/or Customer) will provide the connection to the MPLS provider (where used) and this will terminate any circuits in the Mitel PoP.

SIP Trunk Provider

Nominally the Partner will provide the connection to the SIP Trunk provider. This may be the same partner as the MPLS provider, for example, using another circuit for SIP SBC connections. The SIP Trunk provider must be included in Mitel's SIP CoE list of approved suppliers. These suppliers have been found to inter-work with Mitel products and the configuration settings to be used in the deployment are well understood.

Information on the SIP Trunk Provider interoperability reports can be found under Certified Services (SIP CoE) of the Mitel Solutions Alliance on Infochannel. If you do not have access to this portal, please contact MSAInfo@mitel.com for further information.

Maintenance and Support

Front line support is provided by the partner, and the partner will handle most day-to-day operations or changes. Additional support to the partner can be provided through normal Mitel support channels.

Security

Security is an area where all parties need to play their part. The solution covers many vendor boundaries, and failure of one could lead to security compromise. The Google infrastructure, and the Mitel PoPs, are designed to provide customer to customer isolation as well as Internet isolation. The workloads are defined to run in multiple geographic regions, so that should one region fail, the secondary region can take over operations and provide continued availability. The infrastructure includes additional checks, monitors and internal firewalls to identify and isolate unintended traffic between operational boundaries. Partners and customers are provided with a secured management portal that allows access to that customer, or group of customers under that partner.

External connections via the MPLS network and SIP Trunk provider also need to provide appropriate levels of security. The MiVoice Border Gateway provides the capability for TLS signalling and SRTP encryption. See Connections to SIP Trunk Providers.

Further information can also be found in the "*MiCloud Flex on Google Cloud Security Whitepaper"* available in the Doc Center under MiAccess, under the Security>Technical Papers.

Engineering Guidelines and System Boundaries

The following sections highlight some of the Engineering boundaries that apply to MiCloud Flex on Google Cloud deployment. The boundaries come from two main sources:

- Engineering boundaries related to the deployment within the Google Cloud Platform (GCE and GKE)
- Product and configuration limitations

The guidelines in this section are focused on the limitations associated with the Flex on Google Cloud environment. These guidelines should be used in conjunction with the existing product engineering guidelines. Where there are specific changes to product limitations, these will also be highlighted in the following sections. A number of services, such as CloudLink Chat and MiTeam Meetings are provided on Amazon Web Services, which utilises a redundant and scalable architecture. Where limitations apply to these services, they are included in the sections below.

Additional product Engineering Guidelines can be found at the following documentation location.

- MiVoice Business Engineering Guidelines
- MiCollab Engineering Guidelines
- MiVoice Border Gateway Engineering Guidelines
- MiContact Center Engineering Guidelines
- Mitel Interaction Recorder

Applications of Flex on Google Cloud

The following products are deployed within Flex on Google Cloud:

Product	GCE/GKE	Function
MVoice Business	GKE	Call control, including Voice Mail
MiCollab	GKE	Collaboration suite, plus proxy link to CloudLink Applications on AWS
MiVoice Border Gateway	GKE	Gateway for Teleworker phones and SIP Trunks
Secure Recording Connector	GKE	Secure streaming tap for user and trunk call recording, to the call recording equipment
Mitel Performance Analytics	GKE	Performance monitor and reporter. Also provides secure management tunnel to customer Google Cloud environment
Mitel Business Analytics	GKE	Collection Client for SMDR/CDR collection and transfer to cloud service
Solution Manager	GKE	Management portal to applications
Initial Configuration Wizard	GKE	Used for initial configuration of system, and linking in to MPA management connection

Product	GCE/GKE	Function
MiContact Center	GCE	Contact centre core server and reporting
Teleopti Workforce Management	GCE	Combined with MiContact Center application
Workforce Management	GCE	Traffic analysis probe (included with MiCC) for forwarding data to analysis tool in Azure cloud
Interactive Voice Response	GCE	Front end contact centre call handling, and call direction handler
WFO Interaction Recorder	GCE	Call and Screen Recording. Consists of three prime units: Core server with playback and storage handling, Computer Telephony Integration service, and call and screen recording
WFO Quality Management	GCE	Additional license-able application running on Core Call and Screen Recorder server
WFO Speech Analytics	GCE	Conversion of speech calls to key text words, or full transcription. Requires connection via Quality Management application (Includes Transcription or Keyword Spotting)

Google Cloud Resources

The following identifies the License bundles used in the Flex on Google Cloud deployment. Details of the required resources are also included, although much of this information is integral to the operation of the Service Orchestration Tool. There is no need for the user/partner to adjust these settings. Further details on the bundles are also provided in the sections below.

UC Bundle

The UC bundle comes in four varying license bundles. The underlying infrastructure comes in two main options and there are variations in the infrastructure for primary and secondary data-centre deployments. The resource bundles are grouped to cater for the 250/500 user and 1000/2500 user deployments.

The following table identifies the core UC bundles along with description and nominal user and device capability

Bundle SKU	SKU Name	Description
54010602	cUC Small Infrastructure	250 UC Users (250 users/625 devices), or up to 50 agents*
54010603.	cUC MidSmall Infrastructure	500 UC Users (500 users/1250 devices), or up to 100 agents*
54010604	cUC Medium Infrastructure	1000 UC Users (1000 users/2500 devices), or up to 200 agents*

Bundle SKU	SKU Name	Description
54010605	cUC Large Infrastructure	2500 UC Users (2500 users/5000 devices), or up to 200 agents*
54010609	cUC Medium Infrastructure Expansion	1000 UC users expansion (*subject to PLM approval*)
54010610	cUC Large Infrastructure Expansion	2500 UC users expansion (*subject to PLM approval*)

Note*: A mix of UC Users and Agents is possible for a hybrid system. There is also variations depending on the level of deployment of users and agents as an OTT or MPLS, or hybrid, deployment. Further details can be found under the User and Agent Mix section, below.

Resource Definitions for the UC Bundles:

	Current SKU Definition						Alternative Small SRC[Note2]				Alternative Large SRC[Note2]			
Bundle SKU	**SKU Name**	**Description**	**Location**	**vCPU**	**RAM**	**HDD**	**Location**	**vCPU**	**RAM**	**HDD**	**Location**	**vCPU**	**RAM**	**HDD**
54010602	cUC Small Infrastructure	UC Bundle	Primary	6	26	150	Primary	8	30	190				
		UC Bundle	Secondary	6	6	40	Secondary	8	10	80				
		Postgress DB	Primary	1	3.75	10	Primary	1	3.75	10				
		Postgress DB	Secondary	1	3.75	10	Secondary	1	3.75	10				
54010603	cUC MidSmall Infrastructure	UC Bundle	Primary	6	26	150	Primary	8	30	190				
		UC Bundle	Secondary	6	6	40	Secondary	8	10	80				
		Postgress DB	Primary	1	3.75	10	Primary	1	3.75	10				
		Postgress DB	Secondary	1	3.75	10	Secondary	1	3.75	10				
54010604	cUC Medium Infrastructure	UC Bundle	Primary	8	30	270	Primary	12	34	310				
		UC Bundle	Secondary	8	7.2	140	Secondary	12	11	180				
		Postgress DB	Primary	1	3.75	10	Primary	1	3.75	10				
		Postgress DB	Secondary	1	3.75	10	Secondary	1	3.75	10				

	Current SKU Definition						Alternative Small SRC[Note2]				Alternative Large SRC[Note2]			
54010605	cUC Large Infrastructure	UC Bundle	Primary	12	34	270	Primary	12	34	310	Primary	16	38	310
		UC Bundle	Secondary	12	11	140	Secondary	12	11	180	Secondary	16	14.4	180
		Postgress DB	Primary	1	3.75	10	Primary	1	3.75	10	Primary	1	3.75	10
		Postgress DB	Secondary	1	3.75	10	Secondary	1	3.75	10	Secondary	1	3.75	10

Note: These deployments define the change, or addition, of resources to the cUC Infrastructure bundles when an internal SRC for call recording is required. The cUC large provides two options depending on the number of users that require the call recording service, with a split at around 300 streaming channels, or around 1500 users.

The UC Bundles come with Current and Alternative definitions. The Current bundle is the standard and default bundle that will be used with the solution deployment. The alternative definitions arise because of the requirement to include additional resources for the MPLS connected user SRC function within the existing GKE cluster. This results in a larger resource footprint than the default deployments. Most deployments include one alternative to add SRC resources. The cUC Large bundle includes two alternative bundles, that are driven by the number of UC users that require call recording. The first (small) alternative caters for up to 1500UC users being recorded and the other (large) alternative caters for up to 2500UC users being recorded.

The Alternative Small SRC bundle should replace the existing Current bundle when the following recording SKUs are in use:

- 54010611: cWFO CR Audio Small Infrastructure
- 54010612: cWFO CR Audio MidSmall Infrastructure
- 54010613: cWFO CR Audio Medium Infrastructure
- 54010624: cWFO CR Audio + Screen Small Infrastructure
- 54010625: cWFO CR Audio + Screen MidSmall Infrastructure
- 54010626: cWFO CR Audio + Screen Medium Infrastructure

The Alternative Large SRC bundle should replace the existing Current bundle when the following recording SKUs are in use:

- 54010614: cWFO CR Audio Large Infrastructure

Contact Centre Bundles

The following table identifies the Contact Centre bundles along with description and nominal user and device capability.

Bundle SKU	SKU Name	Description
54010623	cMiCC Business Single IVR Infrastructure	Secondary IVR for MiCC that includes integrated IVR in the primary unit
54010685	cMICC Business Dual IVR Infrastructure	Remote Resilient IVR for MiCC
54010622	cMiCC Business Reporter Infrastructure	MiCC Business Reporter with remote SQL server
54010686	cMiCC Business Small Infrastructure SQL-Express	MiCC Core controller with integrated SQL-Express (limited to 50 agents)
54010687	cMiCC Business Small Infrastructure SQL-Express Multi-Media	MiCC Core controller with integrated SQL-Express (limited to 50 agents)
54010618	cMiCC Business Small Infrastructure	MiCC Core controller with remote SQL server (limited to 50 agents)
54010619	cMiCC Business Small Infrastructure Multi-Media	MiCC Core controller with remote SQL server (limited to 50 agents)
54010621	cMiCC Business Medium Infrastructure	MiCC Core controller with remote SQL server (limited to 200 agents)
54010620	cMiCC Business Medium Infrastructure Multi-Media	MiCC Core controller with remote SQL server (limited to 200 agents)

Resource definitions for the Contact Centre Bundles:

Bundle SKU	SKU Name	Description	Location	Quantity	vCPU	RAM	HDD	Snapshot
54010623	cMiCC Business Single IVR Infrastructure	IVR	Secondary	1	4	8	80G	
54010685	cMICC Business Dual IVR Infrastructure	IVR	Primary	1	4	8	80G	
		IVR	Secondary	1	4	8	80G	
54010622	cMiCC Business Reporter Infrastructure	MiCC	Primary	1	4	16	200G	
		SQL DB	Primary	1	2	7.5	50G	120G
54010686	cMiCC Business Small Infrastructure SQL-Express	MiCC	Primary	1	4	16	200G	

Bundle SKU	SKU Name	Description	Location	Quantity	vCPU	RAM	HDD	Snapshot
54010687	cMiCC Business Small Infrastructure SQL-Express Multi-Media	MiCC	Primary	1	8	32	200G	
54010618	cMiCC Business Small Infrastructure	MiCC	Primary	1	4	16	200G	
		SQL DB	Primary	1	2	7.5	50G	120G
54010619	cMiCC Business Small Infrastructure Multi-Media	MiCC	Primary	1	8	32	200G	
		SQL DB	Primary	1	2	7.5	50G	120G
54010621	cMiCC Business Medium Infrastructure	MiCC	Primary	1	4	16	200G	
		SQL DB	Primary	1	2	7.5	160G	480G
54010620	cMiCC Business Medium Infrastructure Multi-Media	MiCC	Primary	1	8	32	200G	
		SQL DB	Primary	1	2	7.5	160G	480G

Call Recording Bundles

The following table identifies the Call and Screen Recording bundles under the Workforce Optimization (WFO) application suite, along with description and nominal user and device capability.

Bundle SKU	SKU Name	Description
54010611	cWFO CR Audio Small Infrastructure	WFO Call Recording, audio only (limited to 50 agents, 250UC users), resilient audio recording with 1 year audio storage
54010612	cWFO CR Audio MidSmall Infrastructure	WFO Call Recording, audio only (limited to 100 agents, 500UC users), resilient audio recording with 1 year audio storage
54010613	cWFO CR Audio Medium Infrastructure	WFO Call Recording, audio only (limited to 200 agents, 1000UC users), resilient audio recording with 1 year audio storage
54010614	cWFO CR Audio Large Infrastructure	WFO Call Recording, audio only (limited to 2500UC users), resilient audio recording with 1 year audio storage
54010624	cWFO CR Audio + Screen Small Infrastructure	WFO Call Recording, audio and screen (limited to 50 agents), resilient audio recording with 1 year audio storage, non-resilient screen recording with 3 months of storage
54010625	cWFO CR Audio + Screen MidSmall Infrastructure	WFO Call Recording, audio and screen (limited to 100 agents), resilient audio recording with 1 year audio storage, non-resilient screen recording with 3 months of storage
54010626	cWFO CR Audio +	WFO Call Recording, audio and screen (limited to 200

Bundle SKU	SKU Name	Description
	Screen Medium Infrastructure	agents), resilient audio recording with 1 year audio storage, non-resilient screen recording with 3 months of storage
54010616	cWFO 2Yr Audio Storage Expansion	Provides an additional 2 years of audio storage for 50 agents (250UC users)
54010617	cWFO 3Mos ScrRec Storage Expansion	Provides an additional 3 months of screen recording storage for 50 agents

Resource Definitions for the Call Recording Bundles:

Bundle SKU	SKU Name	Description	Location	Quantity	vCPU	RAM	HDD	Snapshot	Notes
54010611	cWFO CR Audio Small Infrastructure	Server H Small (Core Server, DB and Replay)	Primary	1	4	8	210G		
		Server D Small (Audio Recorder and CTI)	Primary	1	4	4	160G		
		Server D Small (Audio Recorder and CTI)	Secondary	1	4	4	160G		
		Secure Recording Connector (SRC)	Primary	1	4	3.6	40G		This unit is deployed as part of the cUC Bundles. See 'Alternative Small SRC'
		Secure Recording Connector (SRC)	Secondary	1	4	3.6	40G		This unit is deployed as part of the cUC Bundles. See 'Alternative Small SRC'
		SQL DB	Primary	1	2	7.5	50G	120G	
54010612	cWFO CR Audio MidSmall Infrastructure	Server H Small (Core Server, DB and Replay)	Primary	1	4	8	210G		
		Server D Small (Audio Recorder	Primary	1	4	4	160G		

Bundle SKU	SKU Name	Description	Location	Quantity	vCPU	RAM	HDD	Snapshot	Notes
		and CTI)							
		Server D Small (Audio Recorder and CTI)	Secondary	1	4	4	160G		
		Secure Recording Connector (SRC)	Primary	1	4	3.6	40G		This unit is deployed as part of the cUC Bundles. See 'Alternative Small SRC'
		Secure Recording Connector (SRC)	Secondary	1	4	3.6	40G		This unit is deployed as part of the cUC Bundles. See 'Alternative Small SRC'
		SQL DB	Primary	1	2	7.5	80G	240G	
54010613	cWFO CR Audio Medium Infrastructure	Server H Small (Core Server, DB and Replay)	Primary	1	4	8	210G		
		Server D Small (Audio Recorder and CTI)	Primary	1	4	4	160G		
		Server D Small (Audio Recorder and CTI)	Secondary	1	4	4	160G		
		Secure Recording Connector (SRC)	Primary	1	4	3.6	40G		This unit is deployed as part of the cUC Bundles. See 'Alternative Small SRC'
		Secure Recording Connector (SRC)	Secondary	1	4	3.6	40G		This unit is deployed as part of the cUC Bundles. See 'Alternative

Bundle SKU	SKU Name	Description	Location	Quantity	vCPU	RAM	HDD	Snapshot	Notes
									Small SRC'
		SQL DB	Primary	1	2	7.5	160G	480G	
54010614	cWFO CR Audio Large Infrastructure	Server H Medium (Core Server, DB and Replay)	Primary	1	6	16	210G		
		Server D Large (Audio Recorder and CTI)	Primary	1	8	8	160G		
		Server D Large (Audio Recorder and CTI)	Secondary	1	8	8	160G		
		Secure Recording Connector (SRC)	Primary	1	8	7.2	40G		This unit is deployed as part of the cUC Bundles. See 'Alternative Large SRC'
		Secure Recording Connector (SRC)	Secondary	1	8	7.2	40G		This unit is deployed as part of the cUC Bundles. See 'Alternative Large SRC'
		SQL DB	Primary	1	2	7.5	160G	480G	
54010624	cWFO CR Audio + Screen Small Infrastructure	Server H Small (Core Server, DB and Replay)	Primary	1	4	8	210G		
		Server D Medium (Audio and Screen Recorder, and CTI)	Primary	1	6	6	160G		
		Server D Medium (Audio and Screen Recorder, and CTI)	Secondary	1	6	6	160G		

Bundle SKU	SKU Name	Description	Location	Quantity	vCPU	RAM	HDD	Snapshot	Notes
		Secure Recording Connector (SRC)	Primary	1	4	3.6	40G		This unit is deployed as part of the cUC Bundles. See 'Alternative Small SRC'
		Secure Recording Connector (SRC)	Secondary	1	4	3.6	40G		This unit is deployed as part of the cUC Bundles. See 'Alternative Small SRC'
		SQL DB	Primary	1	2	7.5	50G	120G	
54010625	cWFO CR Audio + Screen MidSmall Infrastructure	Server H Small (Core Server, DB and Replay)	Primary	1	4	8	210G		
		Server B Small (Audio Recorder)	Primary	1	4	4	160G		
		Server B Small (Audio Recorder)	Secondary	1	4	4	160G		
		Server B Large (Screen Recorder)	Primary	1	8	8	160G		Non-resilient deployment. Only primary site.
		Server C Small (CTI)	Primary	1	4	4	160G		
		Server C Small (CTI)	Secondary	1	4	4	160G		
		Secure Recording Connector (SRC)	Primary	1	4	3.6	40G		This unit is deployed as part of the cUC Bundles. See 'Alternative Small SRC'
		Secure Recording Connector	Secondary	1	4	3.6	40G		This unit is deployed as part of the cUC

Bundle SKU	SKU Name	Description	Location	Quantity	vCPU	RAM	HDD	Snapshot	Notes
		(SRC)							Bundles. See 'Alternative Small SRC'
		SQL DB	Primary	1	2	7.5	80G	240G	
54010626	cWFO CR Audio + Screen Medium Infrastructure	Server H Medium (Core Server, DB and Replay)	Primary	1	6	16	210G		
		Server B Small (Audio Recorder)	Primary	1	4	4	160G		
		Server B Small (Audio Recorder)	Secondary	1	4	4	160G		
		Server B Large (Screen Recorder)	Primary	2	8	8	160G		Non-resilient deployment. Only primary site.
		Server C Small (CTI)	Primary	1	4	4	160G		
		Server C Small (CTI)	Secondary	1	4	4	160G		
		Secure Recording Connector (SRC)	Primary	1	4	3.6	40G		This unit is deployed as part of the cUC Bundles. See 'Alternative Small SRC'
		Secure Recording Connector (SRC)	Secondary	1	4	3.6	40G		This unit is deployed as part of the cUC Bundles. See 'Alternative Small SRC'
		SQL DB	Primary	1	2	7.5	160G	480G	

Mitel Speech Analytics - Keyword Spotting and Transcription Services

Give a description of the resources needed and any channel capacity limits.

Bundle SKU	SKU Name	Description
54010607	cWFO Speech Analytics Base Infrastructure	Speech Analytics Queuing server (3 languages) and one Transcription/Keyword decoder server (3 channels)
54010608	cWFO Speech Analytics Add-on Infrastructure	Speech Analytics Transcription/Keyword decoder server (3 channels)

Resource Definitions for the Mitel Speech Analytics - Keyword Spotting and Transcription Services:

Bundle SKU	SKU Name	Description	Location	Quantity	vCPU	RAM	HDD	Snapshot
54010607	cWFO Speech Analytics Base Infrastructure	Keyword and Transcription Core and Queuing Server (3 Languages)	Primary	1	2	8	120G	
		Keyword and Transcription Decoder Server (3 channels)	Primary	1	8	52	120G	
54010608	cWFO Speech Analytics Add-on Infrastructure	Keyword and Transcription Decoder Server (3 channels)	Primary	1	8	52	120G	

Onsite Gateway Base Kit

The following table identifies the onsite gateway base kit used in the Flex on Google Cloud deployment. It is required even if existing hardware is being re-purposed as an onsite gateway.

Onsite Gateway License

Base Kit SKU	SKU Name	Description
54009825	SP Subscript - MiVB Cloud Gateway SW PKG	One MiVB Cloud Gateway software package

Phone Support

The following phones are supported on Flex on Google Cloud:

- 5300e series and 5304, 5312, 5324, 5320 *[1]
- 6905
- 6910
- 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, and
  - ACD Hot Desking SIP Softphone

Note 1: 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. FQDN support is included in Phone Firmware load 6.5.0.128 and higher. MiVB 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.

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.

Phones that are used with Flex on Google Cloud are required to support FQDNs. Already deployed 5300 phones may not fully support FQDN and will require to be upgraded for Phone Firmware version 6.5.0.128, or higher. The phone loads provided in MiVB Release 9.1 SP1 include this firmware version. Earlier 5300 phones, including 5302, 5330, 5340 (non-'e' versions) cannot be upgraded for FQDN support.

The SIP DECT 112 phone has limited capacity for FQDN length, such that the SIP user name 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.

Browser Support

The solution management and applications support browsers that are based on the Chromium platform. This includes latest versions of Microsoft Edge, and Google Chrome. Firefox is also supported.

Other browsers may work, but these have not been verified to work in all situations, for example Safari is known to work for a number of the applications but may not work on all.

Microsoft Internet Explorer is not supported.

DHCP and Customer VLAN Support

DHCP support for customer on-premise is not provided within the MiCloud Flex for Google Cloud deployment. DHCP must be provided locally with the devices using local DHCP servers, for example using an existing Windows server. For deployments that require the use of VLANs (recommended) for voice segregation, use of DHCP options 125 is recommended. In addition, the DHCP options must also support use of FQDN instead of IP address. Any FQDN entries must also be backed by a customer based DNS service which is configured with a forward helper to the Google Cloud DNS service. See DNS Support, in the section below. This local DNS will be synchronised with information from Google Cloud into the customer network.

For 6900 phones that connect as Teleworker phones, the Redirection and Configuration Service (RCS) can be used to direct the phones to the correct MiVoice Border Gateway Portal. However, the RCS cannot provide specifics associated with the local network such as DSCP and VLAN numbering. The MiCollab Softphones that register as Teleworker phones, will receive their registration details via the application specific Client Deployment Service provided through MiCollab.

Teleworker phones can also be statically programmed, instead of obtaining information via RCS or local DHCP. In this case the FQDN of the call server, or registration, will need to be entered manually from the phone keyboard.

Further details on the content for DHCP Option 125 can be found in the Mitel IP-Sets Engineering Guidelines (under General IP-Phone). Information on how to configure Option 125 on a Windows Server is also available within the MiCloud Flex on Google Cloud Deployment Guidelines. Take care to note which release version of the phone is deployed in order to have FQDN functionality.

FQDNs to use within DHCP will be provided in the configuration confirmation e-mail.

DNS Support

Kubernetes on Google Cloud uses dynamically assigned IP address within the Container environment. This means that each time a pod is started, the IP address assigned to a service may be different from the previous restart. Because IP addresses for services are dynamic, extensive use of FQDNs is now a requirement of all applications and end devices, in order to inter-operate in this environment. Certain IP addresses are static, such as those used by the Internet facing gateways (MiVoice Border Gateway), but internally, IP addresses may change, within subnet bounds.

In order that end-devices connect to the correct service, the IP addresses need to be published, and be able to be dynamically updated. For this reason, it is not recommended to manually enter IP addresses into end devices. The mechanism used to ensure that the correct IP address is provided to the end device, is through the DNS. The DNS will translate the provided FQDN (including those provided to devices through DHCP options) to a current IP address that the end-device can then use to make the correct connection.

In order to ensure that all DNS services are kept up to date with IP address changes, these all need to be synchronised together. There are four main services that need to be updated:

- Customer DNS
- Private Google Cloud DNS
- Public Internet Google Cloud DNS
- Kube-DNS

Changes within the GKE (container) environment, are reflected locally within the Kube-DNS service. The External DNS service takes this information and pushes updates to the internal Private Cloud DNS service, which also holds DNS entries for the more static GCE deployed

applications. Static-IP entries, for the GCE services are pushed from the Service Orchestration Tool to the Private Cloud DNS. The Service Orchestration Tool also pushes Static IP addresses for externally facing services to external Public Google Cloud DNS. These external DNS entries will point to the gateway where the applications can be accessed, whereas the same entry in the Private Cloud DNS will point directly to the application. This is split-DNS operation where the same FQDN resolves to different IP-addresses depending on the domain of reference. To complete the operation, a connection is now needed between the customer DNS service and the Private Cloud DNS service. See the diagram below for a pictorial view.

The customer DNS peering is also needed for customer network phone connectivity on the internal network, where connections are direct to the services and applications hosted in the Google Cloud Platform.

Clients will need to have two DNS servers programmed:

- Primary DNS points to the Private Cloud DNS
- Secondary DNS points to the Customer DNS

This is to ensure that the applications can access services hosted by the customer, where applicable, as well as accessing services within the Cloud.

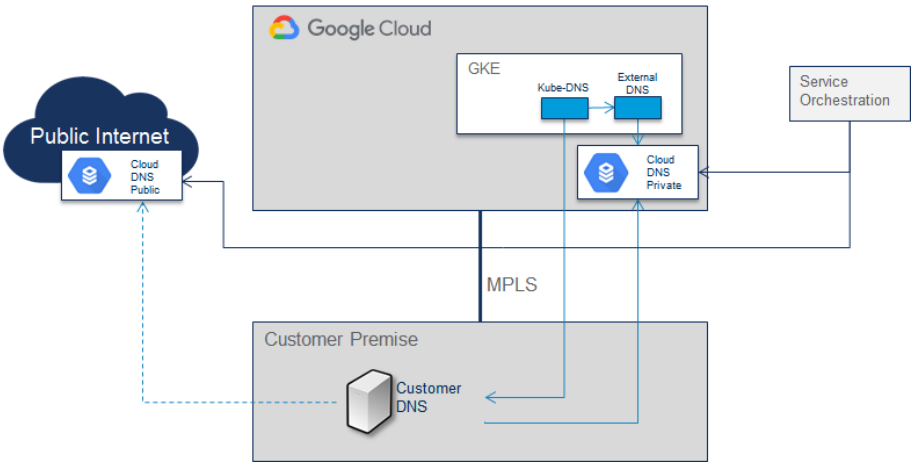

As part of the deployment, some additional information needs to be known and configured. This is provided to the customer as part of the configuration confirmation e-mail and is also an input into the Service Orchestration Tool.

There are two slightly different deployment processes, depending on whether the customer is supplying the FQDN, or the default is in use:

- Customer uses default FQDN
  - Split DNS rules are posted to the Public Cloud DNS and Private Cloud DNS
  - A DNS forwarding rule is configured in the Customer DNS to point to the Private Cloud DNS (IP provided in e-mail)
  - IP address and domain paring of the Customer DNS is provided to the Service Orchestration Tool
  - Kube-DNS now has link to Customer DNS
- Customer uses their own FQDN
  - Customer provides their own Public Cloud DNS with all relevant FQDN to point to the gateways (IP provided in e-mail)
  - Customer FQDN provided to the Service Orchestration Tool prior to deployment

- A DNS forwarding rule is configured in the Customer DNS to point to Private Cloud DNS (IP provided in e-mail)
- IP address and domain paring of the Customer DNS is provided to the Service Orchestration Tool
- Kube-DNS now has link to Customer DNS

MiCollab Client and DNS SRV records

When using MiCollab Client some additional requirements apply when deploying with a remote on-premise gateway. In this case the programming of the pairing profile of primary and secondary MIVBs needs to be changed. In a hosted only solution, the pairing profile identifies the primary and secondary MiVB controller hosted within Google Cloud. However, when a remote gateway is deployed and this is used as the secondary controller, a new pairing profile needs to be created identifying the primary controller in Google Cloud, and the secondary as the remote gateway.

Each pairing profile is associated with a DNS-SRV record so that the phones can register to the correct pair of controllers. When a new pairing profile is created, in order to include the remote gateway, a new DNS-SRV record needs to be generated and included into the customer DNS server. The MiCollab Clients also require a unique MiCollab Client Profile entry under the "PBX SIP Host' field. Essentially a new set of DNS-SRV records (3 needed to include TCP (port 5060), UDP (port 5060) and TLS (port 5061)) needs to be created for each primary and secondary MiVB controller pair and assigned to the appropriate MiCollab Client devices. Further details can be found in the MiCollab documentation, within the Document Center, under MiCollab Client within the *MiCollab Client Resiliency Guide* document.

Customer/Partner delivered FQDN

The customer or partner can provide a specific domain for FQDNs to permit their own advertisement on the Internet DNS system. For example, a customer may wish to include their business name, rather than the default setting, or the partner may wish to advertise their own domain in FQDNs.

Associated with the use of customer or partner specific provided domains in FQDNs is the requirement to generate and manage the associated authentication certificates.

The following sections identify the process to allow a customer or partner to use a specific FQDN and a valid certificate to the generated.

Inclusion of Domain for FQDNs

At the point the customer configuration is entered into CPQ an option is provided to include a customer specific domain to use for FQDNs. Behind the scenes during the infrastructure deployment process, the Service Orchestration Tool will take the provided domain information and use this to generate appropriate FQDNs. For external Internet connections, each FQDN will be assigned a public IP address from the Mitel pool as part of the MiCloud Flex service.

In assigning FQDNs to the customer the expectation is that the MiCloud Flex UC service will be provided a sub-domain to the main business domain, using an additional domain identifier. For example, if the Partner, "Joe's Telecom", normally identifies as "joestelecom.com", it would be expected to include an additional domain as well as the customer name, e.g. "customername.uc.joestelecom.com".

The generation of FQDN can take two forms:

- Default provided by Mitel Flex on Google Cloud using the customer name provided in CPQ
- Domain name provided by Customer or partner

Default Generation of FQDN

In the default configuration, the customer name, provided in CPQ, will be appended to the flex.gc.mitel.io domain. The FQDN will be generated in the form of:

- "<application>.<customer-name>.flex.gc.mitel.io" for the primary deployment, and
- "<application>.<customer-name>-b.flec.gc.mitel.io, for the secondary deployment.

The "<application>" field will be provided by the Service Orchestration Tool and will describe the underlying application that is being referenced, e.g. mivb for call control; micc for contact centre. The "<customer-name>" field will be derived from CPQ using predefined rules for removal of capitals and certain text characters that are not used in DNS. The additional "-b" is appended to the customer name to identify the secondary deployment from the primary deployment. The following table describes the 'application' field entries:

Application field	IP Address	Description
ucs	main address, e.g. .36	Unified Communication Service, including MiVB, MiCollab and MiVBG
awv	main address + 1, e.g. .37	Audio Web and Video conferencing of MiCollab. Uses dedicated IP address to provide service
micc	main address, e.g. 36	MiContact Centre
mbg	main address, e.g. .36	Mitel Border Gateway
mivb	main address, e.g. .36	MIVoice Business Call Control
micollab	main address, e.g. .36	MICollaboration Services, including presence information
wfo	main address, e.g. .36	Workforce Optimization, including call recording, quality management and transcription services

Future applications may be added to this list. The 'main address' can be the primary Point of Presence access location, or the secondary location, e.g. Primary could be XX.XX.33.36 and secondary YY.YY.36.36

For example, if the customer name is 'fredsgarage', then the following will result:

- Primary Location: "mivb.fredsgarage.flex.gc.mitel.io" - this would be the call control in the primary deployment location
- Secondary Location: "mivb.fredsgarage-b.flex.gc.mitel.io" - this would be the call control in the secondary deployment location

The FQDN generated using the default mechanism will be pushed to both the internal network DNS and also the public DNS service. This only applies to names that are subdomains of 'flex.gc.mitel.io'

Custom/Customer or Partner Generation of FQDN

For a custom, or partner/customer delivered FQDN, the following additional items are required:

- A valid domain name, e.g. mybusiness.com, or mybusiness.mypartner.com (the customer may provide the valid domain name, or if this is provided via a partner then the partner name may also be part of that valid domain name)

- Two different subdomains to be associated with the valid domain name. These are used to distinguish between the primary and the secondary locations, e.g. 'a' and 'b', or 'pri' and 'sec', or 'cloud1' and 'cloud2', etc.

The format of the FQDN will take the following format: "<application>.<subdomain>.<customer provided domain name>, where the "<application>" name will be provided by the Service Orchestration Tool, and the other entries are those provided by the customer or partner. The use of shorter entries is recommended where possible, remembering that the FQDNs may need to be manually entered into the phones from the dialpad.

As part of the deployment, the domain name, the subdomain names and applications will be used to create the necessary FQDN. This would lead to the following, if using 'fredsgarage.com' as the main domain name, and where 'pri' and 'sec' are used to identify the primary and secondary deployment locations, respectively.

- Primary Location example: 'mivb.pri.fredsgarage.com' - this would be the call control in the primary deployment location
- Secondary Location example: 'mivb.sec.fredsgarage.com' - this would be the call control in the secondary location

The partner name could equally be included with the customer name, such that 'fredsgarage.com' could be replaced with 'fredsgarage.mypartner.com'.

The customer/partner will be provided an e-mail with the relevant FQDNs, which will need to be registered with DNS. Mitel will not register these names, only those created using the default settings, as described in the section above. Included in this e-mail will also be some FQDN records related to access to the MiVoice Border Gateway. These will need to be provisioned as SRV records. An example of suitable records and information is provided below:

FQDN	TTL	Priority	Weight	Port	Target	Description
_sip._tcp.mbg.pri.mycompany.com	300	1	100	5060	mbg.pri.mycompany.com	Targets unsecured TCP signalling to the MiVBG in the primary data centre (Primary Subdomain used in FQDN)
_sip._tcp.mbg.pri.mycompany.com	300	2	100	5060	mbg.sec.mycompany.com	Targets unsecured TCP signalling to the MiVBG in the secondary data centre (Primary Subdomain used in FQDN)
_sip._udp.mbg.pri.mycompany.com	300	1	100	5060	mbg.pri.mycompany.com	Targets unsecured UDP signalling to the MiVBG in the primary data centre (Primary Subdomain used in FQDN)
_sip._udp.mbg.pri.mycompany.com	300	2	100	5060	mbg.sec.mycompany.com	Targets unsecured UDP signalling to the MiVBG in the secondary data centre (Primary Subdomain used in FQDN)
_sips._tcp.mbg.pri.mycompany.com	300	1	100	5061	mbg.pri.mycompany.com	Targets TLS secured TCP signalling to the MiVBG in the primary data centre (Primary Subdomain used in FQDN)

FQDN	TTL	Priority	Weight	Port	Target	Description
_sips._tcp.mbg.pri.mycompany.com	300	2	100	5061	mbg.sec.mycompany.com	Targets TLS secured TCP signalling to the MiVBG in the secondary data centre (Primary Subdomain used in FQDN)

Further examples of FQDN information provided in the welcome e-mail can also be found in the Appendix of the Flex Deployment Guide, available under 'Doc Center>Business Phone Systems > MiCloud Flex'.

Certificate Management

The front-end access into the MiCloud Flex on Google Cloud solution requires that a valid certificate be generated and provided so that customers connecting to the system can correctly verify the connection. An automated certificate management service, included with the deployment, will obtain a valid Domain Validated (DV) certificate using Let's Encrypt. Mitel is also the handling periodic certificate refresh updates. Customer specific Extended Validation (EV) level certificates provided by other certificate authorities are not currently supported.

The generated certificates include a validity period. The validity period is generally short (weeks) in duration and will need to be updated a number of times during a customer contract period. The automated certificate management service will handle this update and renew the certificate in advance of the expiry time, ensuring continuously available connections.

Where customers provide their own domain name, a welcome e-mail will be provided highlighting a number of FQDNs for the deployment. This may include FQDN for product that is not currently provisioned for the customer, but may be at a later time. When generating the certificate, all of the provided FQDN should be used. Failure to do so may result in future connections failing and a requirement to generate and deploy a new certificate.

Customer MPLS Connections

The partner needs to ensure delivery of customer MPLS terminations in both PoP in different regions in order to provide full resiliency functionality. The deployment of MiCloud Flex on Google Cloud will use default IP addresses for the deployment. The partner should ensure that this will not cause a conflict with any existing customer IP addressing. If there is a conflict the default address range can be changed at deployment time. The default IP address range that will be advertised is based on a /24 address subnet.

The default address ranges advertised to the end customer are:

- 10.13.220.0/24
- 10.20.220.0/24

If these addresses cause conflict with the existing customer deployment, these can be changed to other /24 subnets within the ranges 10.0.0.0/8 and 192.168.0.0/16. Google infrastructure uses 172.16.0.0/12 addresses internally and it is not recommended to use selections within this range. Mitel also uses subnets within this range for networking, specifically 172.20.0.0/16 and 172.21.0.0/16.

Voice Quality

Voice quality is paramount to customers. This can be impacted by a number of factors, with the primary factor being packet loss. Packet loss can occur for a number of reasons, but generally it is due to congestion, either in the network devices, or in the capacity of the

connection. Although network configurations can be applied, and also honoured by different network equipment, the real determination of voice quality is at the end devices. Mitel end devices and gateways monitor the performance of the voice terminations, including packet jitter and packet loss. From this a determination can be made as to the voice quality of the connection. Mitel IP-Phones also report the connection statistics back to the call controller during and at the end of a call.

The Mitel Performance Analytics tools collects the statistics for the calls and can provide an almost immediate view of the call quality. It can also provide alarm indication and notification should any unexpected network condition create unexpected poor voice quality. Although the connections via MPLS are generally expected to be good, the Mitel Performance Analytics probe can also be deployed at a customer site, and a predetermination made for suitability of Voice over IP for the network. This can be of great assistance when used prior to deploying Teleworker phones at a location where network quality may be unknown.

Additional information can also be found in the following documents: *IP-Set Engineering Guidelines*, *Networking for IP Telephony* and *VQ Troubleshooting Guide*, available in the Doc Center.

Speeds and Feeds

Bandwidth for Services

The main services that consume bandwidth are:

- Voice Streaming
- Web Streaming and Collaboration
- Video Streaming
- Call Setup and Signalling
- E-mail and chat

In terms of peak bandwidth, the first three items (Voice, Web Collaboration and Video) in the list are the prime consumers of bandwidth capacity and determine the required bandwidth of any connection. Typically, these services use UDP connections and are latency and jitter sensitive. The last two services are less critical services and use significantly less peak bandwidth. Typically, the last two items use TCP as their transport medium, for guaranteed delivery, where latency and jitter are of much less concern.

Where overall consumption is considered, all these items continue to consume capacity over the longer period, including periods where there is no call traffic. These consumption figures, and associated costs, are included within the deployment bundle licenses.

As some simple rules of thumb, the following values can be used to estimate the peak bandwidth required of a connection:

Media Connection	Rule of Thumb Peak Bandwidth	Notes
Audio	100kbits/s	This is based on G.711 base CODEC. Other CODECs typically consume less, often providing better audio quality. Some CODECs do consume more.
Web Collaboration	50kbits/s per user	This is a variable rate and may even increase to 250kbits/s. However, this is a good starting value for standard resolutions
Video Streaming	2.5Mbits/s	This is a variable rate and may increase up to 10Mbits/s in some cases. This is resolution dependent, but for conference sessions, this is a reasonable starting value.

Bandwidth Consumption

Although a number of services require peak bandwidths, there are other services that consume background levels of bandwidth., or bandwidth consumption. Often these services include background tasks, such as:

- Signaling
- Logs
- E-mails
- Downloads
- Dashboard updates

Based on nominal traffic patterns, the service consumption is included in the license and deployment bundles.

Trunks

SIP Trunks consume audio bandwidth and are utilised close to 100% during the peak hours for office workers (UC Users), typically for 2 hours each business day. Contact centre deployments consume 100% bandwidth for a complete shift, typically 8 hours for each working day. As a minimum requirement, SIP Trunk providers must offer G.711 to maintain compatibility with the Public Switched Telephone Network (PSTN). Some offer additional compression, such as G.729. The SIP gateways (a.k.a. Session Border Controllers) use ITU-T based CODECs, for network to network compatibility. The ITU-T G.711 CODEC usually consumes the most audio bandwidth and offers a simple rule of thumb for bandwidth calculation.

The signaling is SIP based, providing compatibility between different call control platforms and SIP service providers.

As a simple rule of thumb, the Peak Bandwidth of the G.711 CODEC can be used.

Media Connection	Rule of Thumb Peak Bandwidth	Notes
Audio	100kbits/s	This is based on G.711 base CODEC

When calculating bandwidth consumption, the level of traffic and type of installation needs to be considered. Based on the following assumptions:

- A working office day (Standard Office) consists of 8 hours, with typically 2 peak hours per day
- A contact centre work shift consists of 8 hours, with all hours running at peak levels. There are no off-hours for contact centres!
- Assume that the trunks are utilised 100% during peak hours and 30% during off hours for office traffic

The expected bandwidth consumption is highlighted in the table below:

Office Deployment	Per day/shift consumption per trunk	Notes
Standard Office	180M Bytes	This is 2 hours of Busy Hour Traffic and 6 hours of nominal background traffic
Contact Centre	360M Bytes	This is 8 hours of Busy Hour Traffic

UC User, ACD and SIP Trunk Traffic

Call traffic is generally described in Calls per Hour and also how long a call is in hold for. Typical call rates and hold time are:

Connection	Call Rate and Hold Time	Usage time
UC User	6 Calls per Hour at 100 seconds hold time. Two busy hours per day	~16%
ACD Agent	27 Calls per Hour at 100 seconds hold time. Eight busy hours per day/shift	~75% to 100%
Trunks	27 to 36 Calls per Hour at 100 seconds hold time. 2 to 8 hours busy per day	~75% to 100%

The ratio of user to trunks is governed through Erlang traffic calculations and can become quite complex. These also include an overhead to cater for expected peak traffic. However this can be simplified to the following rules for the traffic ratios shown:

User	User to Trunk Ratio
UC User	5:1 (5 users to 1 trunk)
ACD Agent	1:1 (1 agent to 1 trunk)

Interactive Voice Response (IVR)

Interactive Voice Response services are used to supplement agents in a contact centre environment. They can can be used to provide automated functions, effectively replacing agents for all but the most complex call, or they can be used to process a call prior to the call being forwarded to a particular holding queue. Because an IVR is often the front interface within a contact centre environment, these units are generally as busy as the agents, if not more.

A typical ratio of IVR units to agents is:

Service	IVR to Agent Ratio	Trunks per IVR channel
IVR	1:2 (1 IVR per 2 agents, i.e. 50%)	1

When deploying a contact centre the number of active agents will also drive the number of IVR ports and the combined number of trunks. An example contact centre might deploy:

- 100 active agents
- 50 IVR channels
- 150 trunks (100 for agents, and 50 for IVR)

The quantity of available IVR channels is defined in the MiContact Center (MiCC) Guidelines which can be found in the Document Center here. The information is repeated below:

- Remote IVR provides up to 120 channels
- Remote IVR should be used when there is a requirement for 40 or more IVR channels. Below this number, a local IVR can be deployed on the MiCC server, for the local primary IVR. The secondary IVR will still need to be a standalone remote server.

Quantity of IP Ports needed for Media Connections

When making connections through MiVoice Border Gateway (MiVBG) a number of IP ports into the gateway need to be reserved for different services and also to handle potential

changes in service type during a call. The number of ports being reserved for use, at any one time, is dynamic and driven by the number of calls in progress. The number of ports needed per connection are:

Connection Type	IP Ports Reserved	Notes
IP User	4 ports	Line 1 and Line 2 operations (most phones handle 2 lines) plus RTCP
SIP Trunks	8 ports	Multiple ports and possibly different services require different M-Lines, such as CODEC change, that requires 4 RTP ports and associated 4 RTCP ports
WebRTC (audio)	4 ports	Two ports for RTP and media changes and associated 2 RTCP ports
WebRTC (audio and video)	8 ports	Two ports for audio RTP plus 2 ports for associated video RTP, and 4 associated RTCP ports

IP Port Listing MPLS Connection to Customer Site

The IP Port listing considers the IP ports that need to be opened between the customer and a deployment within Google Cloud. It also considers the possibility that the customer may have existing equipment on-premise or may wish to deploy in the future that needs to be cross linked between Google Cloud and on-premise.

Ports	TCP/UDP	Direction	Function
25	TCP	Bidirectional Customer	SMTP
67	UDP	From Customer	DHCP
68	UDP	To Customer	DHCP
69	UDP	From Customer	TFTP
80	TCP	From Customer	HTTP - Web Browser
143	TCP	To Customer	E-Mail service
161	UDP	From Customer	SNMP
162	UDP	To Customer	SNMP-Trap
389	TCP	To Customer	LDAP
443	TCP	From Customer	HTTPS - Web Browser
587	TCP	To Customer	SMTP-TLS
636	TCP	To Customer	LDAPS
1023	TCP	To Customer	FTP Upgrade
1067	TCP	Bidirectional Customer	IP-Trunks - TLS

Ports	TCP/UDP	Direction	Function
1433	TCP	Bidirectional Customer	SQL Server
1752	TCP	To Customer	SMDR
1801	TCP	From Customer	MiCC
2601	TCP	Bidirectional Customer	MiCC/WFO MIR
3268	TCP	From Customer	MiCC
3269	TCP	From Customer	MiCC
3999	TCP	From Customer	SAC protocol for 5300 phone devices
4000	TCP	From Customer	Powerplay Pro Client (Call Recording)
4003	TCP	From Customer	Powerplay Pro Client (Call Recording)
4040	TCP	From Customer	Web Playback (Call Recording)
4443	TCP	From Customer	MiCollab AWV
4498	TCP	From Customer	Screen Recording
4499	TCP	From Customer	Screen Recording
4711	TCP	From Customer	Client Access (Call Recording)
5024	TCP	Bidirectional Customer	MiCC Supervisor Advisor
5025	TCP	Bidirectional Customer	MiCC Supervisor Auditor
5026	TCP	Bidirectional Customer	MiCC MiTai Proxy
5030	TCP	Bidirectional Customer	MiCC Call Recording
5060	TCP	From Customer	SIP
5061	UDP	From Customer	SIP-TLS
5090	TCP	From Customer	MiCC
5091	TCP	From Customer	MiCC
5152	TCP	From Customer	MiCC
5320	TCP	From Customer	MiTai Driver
5400	TCP	From Customer	MiCC
5432	TCP	From Customer	SQL

Ports	TCP/UDP	Direction	Function
6800	TCP	Bidirectional Customer	MiNET
6801	TCP	From Customer	MiNET - TLS
6802	TCP	From Customer	MiNET - TLS
6803	TCP	From Customer	MiNET - TLS
6809	TCP	Bidirectional Customer	MiVBG peering
6810	TCP	Bidirectional Customer	SRC Registration
6815	TCP	Bidirectional Customer	SRC Redirect
7000	TCP	From Customer	MiCC
7001	TCP	From Customer	MiCC Configuration
7003	TCP	From Customer	MiCC Messaging
7011	TCP	Bidirectional Customer	Data Services
7050	TCP	Bidirectional Customer	SDS
8080	TCP	From Customer	MiCollab Redirection
8085	TCP	From Customer	MiCollab AWV
8293	TCP	From Customer	MiCollab Web Service
8492	TCP	From Customer	MiCollab Web Service
9200	TCP	From Customer	MiCC
9300	TCP	From Customer	MiCC
10001	TCP	From Customer	MICC
10118	TCP	From Customer	MiCC
10443	TCP	From Customer	MiCollab Web Service
10991	TCP	Bidirectional Customer	RMI - TLS
15373	TCP	Bidirectional Customer	ACD Real-Time Events
15374	TCP	From Customer	IP-PMS
18100	TCP	Bidirectional	MiCollab SIP peering

Ports	TCP/UDP	Direction	Function
		Customer	
20001	UDP	From Customer	TFTP
36008	TCP	Bidirectional Customer	MiCollab Client client
61616	TCP	From Customer	Transcription Service (Speech Analytics)
12000 - 13999	UDP	Bidirectional Customer	AWV Audio
14000 - 14611	UDP	Bidirectional Customer	AWV Video
20002-29999	UDP	Bidirectional Customer	MiVoice Border Gateway Audio
33000-33999	UDP	Bidirectional Customer	MiVoice Border Gateway WebRTC
35000 - 36999	UDP	Bidirectional Customer	SRC Audio
50000 to 51399	UDP	Bidirectional Customer	Voice Media SRTP

IP Port Listing OTT Connection to Google Cloud

The following is a list of ports that are opened from the Internet to be terminated on the Public IP of the MiVoice Border Gateway in the Flex on Google Cloud solution. This list is a subset of ports defined in the MiVoice Border Gateway Engineering Guidelines.

Port	TCP/UDP	Function
80	TCP	HTTP
443	TCP	HTTPS
3998	TCP	SAC protocol for teleworker 5300 phone devices
4000	TCP	Call Recording and Playback Client
4498-4499	TCP	Call Recording and Playback Client
4711	TCP	Call Recording and Playback Client
5024-5026	TCP	MiContact Center
5060	TCP	SIP
5060	UDP	SIP
5061	TCP	SIP-TLS
5063	TCP	WebRTC TLS (WAN)

Port	TCP/UDP	Function
6801	TCP	Secure MiNet (SSL)
6802	TCP	Secure MiNet (AES)
6809	TCP	Inter MiVBG Communication
6881	TCP	Avatars Teleworker Mode
7001	TCP	MiContact Center
7003	TCP	MiContact Center
8030	TCP	MiContact Center
20000	UDP	Teleworker Analyzer Termination
20001	UDP	TFTP
20002-29999*	UDP	Voice SRTP
32000-32499*	UDP	WebRTC
36008	TCP	MiCollab Client client

Note*: These are gateway default values. See below for deployment defaults.

For the MiCloud Flex on Google Cloud, the full external (OTT) port range of the MiVoice Border Gateway is not required. The following port ranges will be applied to the specific cUC configurations and supersede the default settings.

Infrastructure Deployment	Voice SRTP Port Range	WebRTC Port Range
cUC Small Infrastructure	20002-22001 (250 trunks)	32000-32239 (59 WebRTC/Audio connections)
cUC MidSmall Infrastructure	20002-22001 (250 trunks)	32000-32239 (59 WebRTC/Audio connections)
cUC Medium Infrastructure	20002-24001 (500 trunks)	32000-32767 (191 WebRTC/Audio connections)
cUC Large Infrastructure	20002-26001 (500 trunks+500 user channels)	31767-32767 (250 WebRTC/Audio connections)

The limits on trunks and WebRTC connections and devices (see section WebRTC number of supported clients) are derived from the information in this section and highlighted in the sections of Deployment Boundaries, below. When video is used in conjunction with audio, the number of ports will double, or the number of available channels and users will halve. Audio and Audio/Video connections can be mixed and on demand.

Deployment Boundaries

The information in this section is supplemental to that provided in the product guidelines, highlighted above. The deployment boundaries highlighted here are specific to the Flex deployment on Google Cloud.

MiVoice Business Limits (Call Control)

User and Device limits

There are two main areas that govern the capacity limits for the different cUC (Containerised Unified Communication) bundles. These are:

- Software defined limits
- License defined limits

In most cases, the limits are defined by the license type and the cUC bundle.

Infrastructure Deployment	Number of Users	Number of Devices	Nominal Device to User ratio	Notes
cUC Small	250	625	2.5	Device limit is license bound. Higher device count is possible.
cUC MidSmall	500	1250	2.5	Device limit is license bound. Higher device count is possible.
uUC Medium	1000	2500	2.5	Device limit is license bound. Higher device count is possible.
cUC Large	2500	5000	2	Device limit is software constrained.

Solution user limits are highlighted under the User and Agent Mix Limits section below, of which MiVoice Business is one component. Values shown here relate to this product function capacity.

Call Handling Capacity

The call handling capacity is defined in conjunction with other services. Both media channels and call capacity need to balance each other. The call capacity can peak higher than the nominal amount, but with an impact on the number of media channels available. The system is designed for a nominal call rate of:

- UC User: Typical Office traffic of 6CPH at 100 seconds hold time during the busy hour periods (2 hours per day)
- ACD User: Typical ACD traffic of 27CPH at 100 seconds hold time for a complete shift of 8 hours.

Traffic rates and hold times may vary, depending on the business type.

The systems have a nominal traffic capacity of:

Infrastructure Deployment	Number of Users	Nominal Call capacity
cUC Small	250	6000CPH at 100 seconds hold time
cUC MidSmall	500	6000CPH at 100 seconds hold time
cUC Medium	1000	12000CPH at 100 seconds hold time
cUC Large	2500	12000CPH at 100 seconds hold time

The nominal call rates are sufficient to handle a mix of UC users and ACD agents in line with the ACD and User Mix Limits, below.

Embedded Voice Mail

Embedded Voice Mail provides for storage of around 8GBytes of data for around a total of 250 hours of storage. However, it should be noted that backups for Voice Mail are limited to around 30% of the storage, so up to 80 hours. If there is a requirement to back up all voice mail messages, then steps must be taken to limit the usage to this amount.

The storage capacity applies to all the deployment configurations and provides the following average storage limits, based on the 250 hours of storage:

User Deployment	Average Storage Time* per User	Average Number of Messages per User**
2500	6 minutes	12
1000	15 minutes	30
500	30 minutes	60
250	60 minutes	120

*Divide by 3 to fit within the 80 hours backup limit. **Based on a 30 second message

The maximum number of concurrent calls for Embedded Voice Mail is 120 channels. Note that this number of channels is part of the total available pool of Media Channels, see below. The limits from Media Channels is overriding. It is not necessary that all 120 voice mail channels are active for the different deployments. The default setting is 120. The following table provides recommendations for the maximum number of EMEM channels needed for the different system deployments:

System Deployment	Number of Users	Recommended maximum EMEM ports
uUC Small	250	30
cUC MidSmall	500	30
cUC Medium	1000	60
cUC Large	2500	60

Embedded voice mail includes additional capabilities, previously accessible within MiCollab NuPoint:

- Voice Mail Call Flow (replacement for Call Director)
- Visual Voice Mail

Advance UM is not available in this release. Options are under consideration for future releases.

Where EMEM connections to e-mail servers are required, some level of authentication on the connection may be required. Should this be required, an e-mail forwarder will be required. This should be configured to provide the level of authentication to the e-mail service, as well as limiting input connection from the EMEM service. Note that since EMEM is part of MiVoice Business the IP address is considered dynamic. Use of FQDN is therefore recommended. Include both Primary and Secondary EMEM/MiVoice Business in the access controls. Note that use of "open" e-mail forwarders is not recommended for cloud-based e-mail services. Further details can be found in KBA HO2569.

Media Channels

The number of media channels provided is limited to the following values in the table below

Infrastructure Deployment	Number of Users	Number of Media Channels
cUC Small	250	64
cUC MidSmall	500	64
cUC Medium	1000	128
cUC Large	2500	128

WebRTC number of supported clients

Based on the number of available WebRTC channels and based on standard office traffic the following limits on number of WebRTC clients is recommended:

Infrastructure Deployment	Number of WebRTC Clients
cUC Small	300
cUC MidSmall	300
cUC Medium	1000
cUC Large	1250

The number of available channels is identified in section IP Port Listing OTT Connection to Google Cloud

MiVB Networking

The default networking configuration is for the primary and secondary MiVB to be hosted within the Google infrastructure. With the introduction of the remote on-premise gateway(s), it is now possible to introduce additional secondary controllers into the configuration. Due to the unknown number of remote gateways and unknown configurations to be encountered, networking of the MiVB controllers will need to be manually configured.

Onsite MiVoice Business gateways directly networking to the Cloud MiVoice Business instances will require use of FQDNs and must therefore be running a minimum of MiVoice Business 9.1. Where 5300 phones are used in the deployment, then MiVB Release 9.1 SP1, or higher, is required.

MiCollab Limits

The majority of limits that apply to MiCollab are those described for the Multi-application, single deployment configuration. Information can be found in the MiCollab Engineering Guidelines. Some overriding limits for Flex on Google Cloud are included in the following sections.

User Limits

Flex on Google Cloud is limited by default to 2500 users. Larger deployments are subject to PLM approval.

Web Collaboration

Web Collaboration meetings are different from normal phone traffic, as they are often:

- Long duration calls, typically 1 hours
- Involve a number of users at the same time
- Often includes external users who are not part of the system, i.e. guests
- May include simple audio connections, may include web collaboration, may include video.

Although collaboration sessions can vary from customer to customer, the following table provides guidelines for typical numbers of concurrent web collaboration users and sessions or conferences:

Configuration	Users	Lower Collaboration Users	Upper Collaboration Users	Number of conferences
cUC Small	250	6	15	2
cUC MidSmall	500	12	30	3
cUC Medium	1000	25	60	6
cUC Large	2500	40	100	10

General use suggests conference usage with the lower number of collaboration users. However, at certain times the numbers may peak, and the upper limit should always be considered with respect to bandwidth and network requirements.

Web Collaboration and video conference calls may be handled by one of two services:

- MiCollab AWV, on MiCollab
- CloudLink MiTeam Meetings

When using MiCollab AWV, the collaboration session is hosted on the MiCollab Application. Users will connect over the MPLS network, for internal users, or via the web proxy for external users and Teleworker users. Typically for AWV Web Collaboration, the ratio of calls used for Web collaboration is 90%, compared to 10% for video. For nominal bandwidth usage see under Bandwidth for Service.

For users with CloudLink MiTeam Meetings, the service is hosted directly as a cloud service, from AWS. MiCollab provides the connection and setup proxy. All users are connected over the public network to this AWS cloud service. Internal users will therefore consume Internet bandwidth, and this should be considered when sizing the Internet connection to and from the customer's Internet gateway. With MiTeam Meetings, the level of video bandwidth will increase. Nominally there will be a single video stream to each user, with each conference participant, up to 16 concurrent/most recent active participants, being presented as a subset of that bandwidth. Audio will be provided as a single mixed stream from the bridge, derived from all participants. Customers using MiTeam Meetings should consider the number of concurrent active users when calculating the Internet connected bandwidth.

Chat

Chat is now provided exclusively by CloudLink Chat in AWS. MiCollab is used as a proxy connection. All chat capabilities are now provided by this Cloud service. Nominal peak bandwidth for chat sessions are low, and the information is less time sensitive.

MiCollab Client

MiCollab Client continues to be provided by MiCollab and will provide status updates on user presence. This may also be linked to the customer Directory Services for authentication. Some of the capabilities within MiCollab Client, such as chat, are proxied to the CloudLink Chat Cloud services with the client application making direct connection to this service.

MiVoice Border Gateway (MiVBG) Limits

The MiVoice Border Gateway provides a number of key functions within the deployment, specifically:

- External Gateway for Teleworker devices
- External Gateway for softphones, and "Toaster Popup" notifications
- External Connection to SIP Trunk provider
- Secure Recording Connector (SRC) for Externally connected devices.

In addition, a separate MiVoice Border Gateway will be deployed internally to the cUC configurations to handle Secure Recording Connector (SRC) functions for MPLS connected internal devices.

The nominal capacity of the different deployments are:

Infrastructure Deployment	Number of Streams	Number of Users	Number of Devices	Notes
cUC Small	250	1250	2500	Assumes multiple devices per user
cUC MidSmall	250	1250	2500	Assumes multiple devices per user
cUC Medium	500	2500	5000	Assumes multiple devices per user
cUC Large	500	2500	5000	Assumes multiple devices per user
Plus internal SRC Small	333	2500	2500	Additional Resource added to cUC bundles. This provides streaming for internally connected users and devices
Plus internal SRC Large	667	2500	5000	Additional Resource added to cUC bundles. This provides streaming for internally connected users and devices

Solution user limits are highlighted under the User and Agent Mix Limits section below, of which MiVoice Border Gateway is one component. Values shown here relate to this product function capacity and may show available capacity, which may be different from the overall solution capacity.

When the cUC bundles have SRC enabled, additional streams are needed to connect to the Call Recording solution. This results in an increased streaming requirement, which results in the overall through streaming capacity being reduced. The additional call-recording streams counts as an additional 50% overhead, or the number of through channels is reduced to 67% of the nominal value. For example, in a non-SRC deployment, the through capacity of 500 streams would be reduced to 333 streams when SRC is deployed.

Note that the MiVoice Border Gateway is used for multiple functions, including teleworker users and also trunk connections. Where a user is calling a trunk, a streaming channel will be needed for the user AND also for the trunk. In the case where a user is calling a user, behind the same public address, local streaming can be enabled, and thereby consuming no channels of the MiVoice Border Gateway. However, if a user calls another user, both with different public addresses, then two streams will be consumed as the two users are cross connected via the MiVBG.

Customer to customer connections are not provided directly between individual customer MiVoice Border Gateways. Calls between customers transition through the SIP trunk provider and are billed accordingly.

Where call recording is enabled, the MiVoice Border Gateway also provides a secure 'tap' into the connection. Audio streaming from the two parties is provided to the Call Recording Equipment as two half streams (one way), or effectively a stereo connection with each participant in each channel. Call recording is provided for station side connections. This means that the tap is associated with a named user. The MiVoice Border Gateway provides an SRC function associated with both Teleworker users (External SRC) and also for internal LAN users (Internal SRC). For users that are teleworker connected making a trunk call, the audio is streamed via the MiVoice Border Gateway anyway, and so this has little impact on the audio streaming and routing. However, for internal phone calls, if call recording is enabled for internal calls, then the streaming between the two devices, must be anchored, and hair-pinned via the MiVoice Border Gateway, i.e. where local streaming would be the normal path, this now requires that both end devices connect via the Internal SRC function.

MiCC Limits

it is recommended that the guidelines for MiCC are consulted for the deployment. However, a couple of areas that are relevant to the MiCloud Flex on Google Cloud are highlighted here. The main areas concern the options available in CPQ for selecting the different license bundles. These include options for:

- Use of Embedded Interactive Voice Response (IVR)
- Use of SQL-Express (SQL-X)

MiCC Collector Node

It is recommended to include the MiContact Center (MiCC) Collector Node on the secondary IVR. This collects any contact centre voice interaction data when the primary MiCC Business server is not available. When the primary MiCC Business server becomes available, this call data will be synchronised. This ensures that any historical data collected on the secondary IVR, while disconnected, will not be lost.

Embedded IVR Limits

IVR (Interactive Voice Response) ports are typically deployed along with agents to offset load and re-direct calls to specific queues and sets of agents. Another consideration for separating IVR and the main MiCC Business server is the operation under resilient conditions. In the case where the MiCC Business server is unavailable, the IVRs will still be active and directing calls at primary and secondary sites. For smaller deployments, or where call traffic is lower than normal, there may not be a great demand for IVR ports. In this case, the deployment of an integral IVR with the MiCC Business server may offer cost and management advantages.

One of the options within CPQ is to allow the use of an integral IVR service with MiCC Business server. This selection will trigger a single remote IVR for the resilient secondary site. This option is possible where the number of required IVR ports is low and where the call traffic is also low. The details of when to integrate and IVR function with the MiCC Business server are included in the MiCC Engineering Guidelines under the section "RECOMMENDATIONS FOR COLLOCATING SERVER APPLICATIONS". Typically, this option is available when the number of IVR ports is less than 40.

SQL-Express

The SQL database is used to store configuration data and historical call records of contact centre agents. For many business services that use contact centres, there is often a legal requirement to store this information for a minimum period of one year. For a reasonable sized contact centre with expected levels of traffic, the number of records to be recorded requires substantial storage capacity, and exceeds that available to a SQL-Express deployment. However, there may be cases where it is possible to deploy an SQL-Express server.

An SQL-Express Deployment can handle the following traffic with the following storage duration:

Agents	Call Rate	Shift Duration	SQL-Express Storage Duration
50	27CPH	8 hours (5 days per week)	3 months

If there is no data retention requirement, then the SQL-Express server can possibly be used. However, in most cases the requirement is for a minimum of one year.

It is possible to extend the duration by adjusting some of the initial parameters. For instance, if the number of agents is reduced to 12, the storage duration extends to 12 months. Or, if the call rate reduces to 8-9CPH, the storage duration will extend to 12 months. Or, maybe the contact centre is only active for 2 hours per day?

There are ways to extend the duration to allow the SQL-Express to be deployed. This is useful for the smaller contact centre, but often less so for the larger deployments. CPQ will offer the option to select SQL-Express for smaller contact centre agent quantities.

The partner selecting the SQL-Express option, in CPQ, must be aware of any legal requirements with respect to minimum storage duration, and any traffic considerations to allow SQL-Express to meet this requirement.

Mitel Workforce Optimization Suite

The Mitel Workforce Optimization suite provides a number of advanced customer and user interaction capabilities. The main applications under this category description:

- Mitel Interaction Recording (Call and Screen Recording)
- Mitel Quality Management (Quality management)
- Mitel Speech Analytics (Transcription and Keyword spotting)

The Mitel Interaction Recording solution is the base to offer Mitel Quality Management, Mitel Speech Analytics and Mitel Coaching and Learning. Of specific interest for Engineering Guidelines are Mitel Interaction Recording and Mitel Speech Analytics. These are described in the sections below.

Mitel Interaction Recording Limits (Call and Screen Recording)

The Mitel Interaction Recording (Call and Screen Recording) solution provides capabilities to record both audio calls and screen displays. Some of the key functions of the deployment include:

- Core Server, Playback and Storage
- CTI Connection
- Recording (Audio and Screen)

For smaller deployments it is possible to combine some of these functions onto a common server. In other cases, based on the requirements, these services need to be provided on separate servers. For resilient operation it is also necessary to split some of these functions.

The Core Server is associated with the primary data-centre. This is a non-resilient component as it does not deal with time sensitive audio and screen recording. One of the functions this Core server provides is de-duplication of multiple recordings that may arise from deploying a resilient recording solution. In this way if a call is recorded with multiple recorders, this services can identify the duplication and remove the additional recording.

The CTI Connection service connects to the relevant call control server and controls when the recording servers kick into operations. Since the overall solution includes both a primary and secondary call control, a CTI Connection server is associated with each of these. Thus, if a call is active on the primary server, the primary recorders will record the call. If the call is active on the secondary controller, the secondary recorders will record the call.

The Recording servers have capacity for both audio and screen recording, although screen recording is much more compute intensive. The primary service for the call recording is the audio connection, with screen recording taking secondary priority. For this reason, the Audio recording servers are deployed in a primary/secondary resilient configuration, and the screen recording is only deployed in a primary non-resilient configuration.

The recording servers will store the information from the call until the connection is complete. At this time the file and information is forwarded to the Core server. However, if the Core server is not available the file will not be forwarded and will remain locally on the recording server. Once the Core server becomes available, the files and information will be forwarded for processing. In this way, if the primary site becomes unavailable, the recording servers will continue to operate and provide resilient operation from both the primary and secondary deployment regions.

The storage requirements are also quite extensive. Storage of an audio call involves recording of both sides of the conversation, with around 16kBytes/second being generated per call. Calls are recorded as stereo audio files, one channel for each party. By comparison a screen recording will generate around 300kBytes/second, requiring around 20 times the storage capacity.

The default storage duration associated with the different license bundles is 1 year for audio recording and 3 months for screen recording. Additional recording storage can be purchased. The default values for the license bundles are based on sufficient storage capacity for fair usage and typical call handling and levels of screen recording. Mitel reserves the right to police the usage of storage and either charge more, or to limit the storage capacity and retention period should the boundaries of fair usage be exceeded.

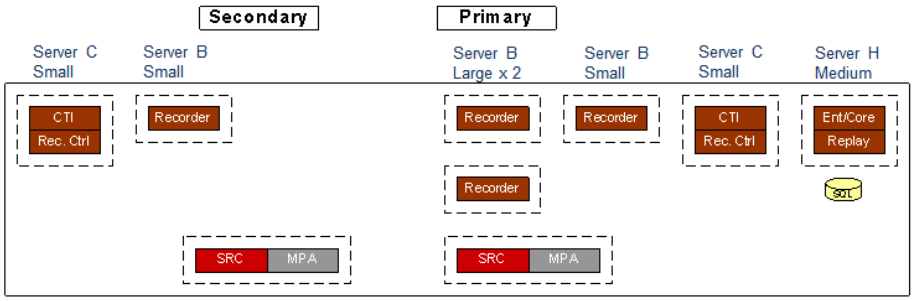

The diagram above shows the different servers associated with the larger recording license bundle. The components of the bundle include:

- Google Cloud SQL Server for call record information
- Google Cloud Storage (not shown)
- Server H Medium - Core Server and Playback
- Server C Small - CTI connection to primary and secondary call control
- Server B Small - Audio recorder for primary and secondary regions
- Server B Large (x 2) - Screen recorder at primary region only
- SRC - This is the Secure Recording Connector resource needed to record the audio from the agents and users on the MPLS network. The MiVoice Border Gateway includes this function for externally connected Teleworker agents and users. This resource is added to the existing GKE deployment, when this license bundle is selected

In terms of Engineering Rules, the deployments and resources have been defined within the specific license bundle. The main requirement is to provide the information into CPQ and allow it to size the deployment with the relevant license bundle. The applicable licenses for WFO MIR are managed in the following manner:

License and Feature	Licensing Model
WFO MIR Voice Rec Subscription	per concurrent channel
WFO MIR Screen Rec Subscription	per named user
WFO MIR Quality Monitoring Subscription	per named user

Per 'concurrent channel' is based on resource definitions, and can be used by any agent, as long as there is resource capacity. Per 'named user' means that only information for that particular user can be processed. For example, suppose the system is designed for 50 active agents, and there are three work shifts, with a total of 150 agents. With a 'concurrent channel' model then only 50 agents are active at any time and 50 licenses are needed. With a 'named user' model, 150 licenses are needed for each agent to access the resource.

Mitel Speech Analytics Speech Transcription Limits

Additional services can be attached to the WFO MIR call recording service. The Mitel Speech Analytics Speech Transcription service transcribes recorded calls into full language text. This information is stored with the file as additional meta-data, and enables rapid data searching for specific file information. This is especially useful if files need to be search for a particular call at a particular, but non-specific time. Full transcription provides a complete view of any interaction, enabling the full conversation to be included within other media, such as documents and e-mail.

The Transcription service runs in the background against existing files. The deployment consists of two main components:

- A queuing server
- A decoder or transcription server

The queuing server will fetch files that need to be transcribed and append the metadata once transcription is completed. It is also the storage point for the language dictionaries used for the transcription and keyword search service. For MiCloud Flex on Google Cloud the queuing server has been sized to handle up to three language files.

The decoder, or transcription server, does the actual work of transcribing the file. This is a compute intensive activity. The server provides a number of parallel channels and can handle a limited number of hours of recordings within 24 hours. Multiple transcription servers may be deployed against a single queuing server. For MiCloud Flex on Google Cloud the transcription server has been scaled to provide 3 parallel channels (Note that the server may refer to 6 channels. However, these are 6 *mono* channels. The recorded audio calls include both sides of the conversation and are recorded as a *stereo* file. Therefore, the real throughput is 3 "stereo" channels). The Speech Analytics engine works on G.711 encoded files. Files must therefore be recorded in G.711, or a suitable transcoder used to convert the files, or Analytics performed prior to storage.

In terms of license bundles two options are provided:

- Base: This includes one queuing server and one decoder, or transcription, server (3 languages and 3 channels)
- Additional Transcription Server: One decoder, or transcription, server (3 channels)

Note that in order to enable Transcription, at least one Quality Monitoring license is required.

The applicable licenses for Speech Analytics Speech Transcription services are managed in the following manner:

License and Feature	Licensing Model
WFO MIR Transcription Subscription	per named user

Transcription Capacity

The following capacity limits apply:

Activity	Transcription
Hours of Recording per channel processed in 24 hours	24
Number of channels per server	3
Nominal number of agents *	12

Note* One decoder channel can Transcribe 24 hours of recordings within 24 hours. Given a typical agent shift of 8 hours and 75% recordings (6 hours), then one channel will effectively handle recordings for 4 agents. A server can therefore nominally handle transcription for 12 agents. Note that different a shift duration or recordings per agent will adjust this figure.

Mitel Speech Analytics Keyword Spotting Limits

Additional services can be attached to the WFO MIR call recording service. The Mitel Speech Analytics Keyword Spotting service transcribes recorded calls for specific key words into text. This is faster than full text transcription, since the search dictionary is reduced. This identified keyword information is stored with the file as additional meta-data, and enables rapid data searching of the file information. This is especially useful where files only need to be identified by certain key words, and rapid search is important.

The Keyword Spotting service runs in the background against existing files. The deployment consists of two main components:

- A queuing server
- A decoder or transcription server

The queuing server will fetch files that need to be keyword searched and append the metadata once transcription is completed. It is also the storage point for the language dictionaries used for the transcription and keyword search service. For MiCloud Flex on Google Cloud the queuing server has been sized to handle up to three language files.

The decoder, or transcription server, does the actual work of transcribing the file. This is a compute intensive activity. The server provides a number of parallel channels and can handle a limited number of hours of recordings within 24 hours. Multiple transcription servers may be deployed against a single queuing server. For MiCloud Flex on Google Cloud the transcription server has been scaled to provide 3 parallel channels (Note that the server may refer to 6 channels. However, these are 6 *mono* channels. The recorded audio calls include both sides of the conversation and are recorded as a *stereo* file. Therefore, the real throughput is 3 "stereo" channels) Keyword searching is effectively a limited language search library, which can result in faster processing with minimal search impact. However, keyword searching does not transcribe every word. The Keyword Spotting engine works on G.711 encoded files. Files must therefore be recorded in G.711, or a suitable transcoder used to convert the files, or Keyword Spotting performed prior to storage.

In terms of license bundles two options are provided:

- Base: This includes one queuing server and one decoder, or transcription, server (3 languages and 3 channels)
- Additional Transcription Server: One decoder, or transcription, server (3 channels)

Note that in order to enable Transcription, at least one Quality Monitoring license is required.

The applicable licenses for Speech Analytics Keyword Spotting services are managed in the following manner:

License and Feature	Licensing Model
WFO MIR Keyword Spotting Subscription	per named user

Keyword Spotting Capacity

The following capacity limits apply:

Activity	Keyword Spotting
Hours of Recording per channel processed in 24 hours	40
Number of channels per server	3
Nominal number of agents *	20

Note* One decoder channel can Keyword spot 40 hours of recordings within 24 hours. Given a typical agent shift of 8 hours and 75% recordings (6 hours), then one channel will effectively handle recordings for 6.7 agents. A server can therefore nominally handle keyword spotting for 20 agents. Note that different a shift duration or recordings per agent will adjust this figure.

Connections to SIP Trunk Providers

The MiCloud Flex on Google Cloud will connect to a SIP Service Provider (SP), as supplied by the partner. The network requirement is that the SIP Session Border Controller (SBC) is accessible using a Public IP address. The SIP SP will either provide this connection over the Internet, or will provide a Network to Network Interface (NNI) to be terminated within the Mitel PoP. The SIP SP may be the same provider as the MPLS provider, or may be independent.

Where a connection is provided over NNI, the SIP SBC public IP address will be seen within the Mitel PoP. The BGP protocol is used to determine the most appropriate route for the SIP trunk connection, either over the public Internet or via the NNI connection. This requires that agreement be reached on which public IP addresses will be advertised from the Mitel PoP to the Carrier and vice-versa. For example if the carrier only provides SIP Trunk services, then only those addresses might be advertised. However, where the carrier provides Internet services to the end customer as well as SIP Trunk services, they may wish to advertise the customer addresses as well. This will enable Teleworker phone and softphone connections across a dedicated network, rather than over the public network. If this is not available, teleworker phones and softphones can still connect, taking the longer route via the MPLS service provider Internet Gateway.

Use of the BGP to advertise the routes provides the option to take the best route to the SIP SBC. Ideally, this would be via the NNI connection. If, for whatever reason, this connection becomes unavailable, the BGP protocol will recognise that the SIP SBC is still accessible over the public Internet connection, assuming that the SIP SP advertises these addresses to the public network. Although a longer router, this mechanism provides improved resiliency to the SIP Trunk connection. Service providers that only offer static routing may not be able to take advantage of the BGP capabilities.

An advantage of routing the data over a dedicated MPLS connection, or via NNI, is the ability to honour Quality of Service (QoS) settings within the private network and provide improved and consistent voice quality. It also provides a level of isolation and added security on the connections. Internet connections are considered best effort.

The SIP Service provider must provide an SBC (Session Border Controller) with a Public IP address (non-RFC1918). The SIP SP must be be included in the SIP Interoperability list under the Mitel Solutions Alliance (MSA). This list can be found under InfoChannel, under Certified Services.

The diagram below shows a deployment where the partner may bring both MPLS and SIP SBC, or where these services are provided by two different carriers. The terminations to the Mitel PoP, may be logically or physically independent.

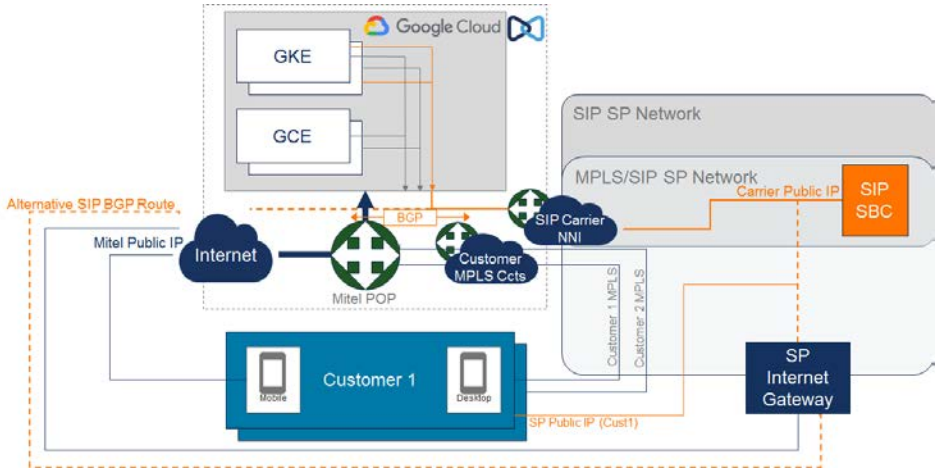

SIP Trunk can use the following authentication mechanisms:

- SIP Username/password registration
- Trusted IP

The MiVoice Border Gateway provides encrypted connections for the following:

- Signalling: SIP-TLS
- Media: SRTP to AES_CM_128_HMAC_SHA1_32 and AES_CM_128_HMAC_SHA1_80, as per RFC3711. The MAC-32 bit is the preferred and default option. MAC-80 bits is selectable.

E911 Emergency Services

E911 location information is associated with the SIP Trunk Provider. Each of the users can be assigned a unique DID to attach this location information to. Typically, a general business number will also have location information associated with it. It is up to the Partner to ensure that this location information is maintained for the customer. Additional Customer Emergency Service Identifier (CESID) numbers may be required to be purchased and configured for the solution to meet the customer's E911 needs. Emergency call routing and notification is as per standard MiVoice Business practices. Note that the channel partner must program both 911 and 9911 access in the USA to meet Kari's Law requirements, i.e. an untrained user must be able to dial 911 from the phone without understanding of access codes. Be aware that for North America, external long distance often starts with '91'.

It should also be noted that because of the location independence of IP connectivity, there is an additional challenge on location information to provide to the E911 location database. Teleworker phones, for example, may be located at a remote location (where location can be determined), or may be more mobile. Likewise, the softphone is also a mobile IP device. This information should also be conveyed to the users of these devices and also to the E911 provider. The MiCollab softphone recognizes the '911' call and will route the call through the native phone device allowing for the possibility of local triangulation, see MiCollab Administration Guidelines for further details.

FAX

FAX termination and generation is not supported directly within the solution. However, FAX (and T.38) connections can be established between a supported SIP ATA device and the SIP Trunk provider, provided that the SIP Trunk provider SBC can support T.38 transcoding. EMEM does not support termination of T.38 FAX calls.

If T.38 is not available, a SIP ATA may still provide FAX pass-through in the audio channel. This circuit will provide best effort operation.

User and Agent Mix Limits

The UC deployments are sized to allow for a mix of agents and users. There are a number of limiting factors that impact the ratio of this mix of users and agents. These include:

- Device limits within MiVB
- User Limits within MiCollab
- Expected user and agent call rates
- Trunk streaming and registration limits within MiVBG
- Ratio of OTT to MPLS connections
- Call Recording and whether this is all users, including agents, OR only agents and not UC users.

The following sections provide an overview of the available mix of users and agents on a system based on standard call traffic rates, as described earlier in this document. The sections also consider that call recording is predominantly agent based. The results are provided for both 100% teleworker deployment and 100% MPLS deployment, with and without call recording. A realistic deployment is probably somewhere in between. All these considerations are included within CPQ analysis, and options are provided at the configuration input to select the teleworker/MPLS ratio. If for any reason a deployment does not fall within the normal settings, it is recommended to contact Professional Services for a more detailed analysis and recommendation.

cUC Small Infrastructure

The deployment limits for this configuration are relatively simple. For Teleworker and MPLS configurations as well as with and without call recording, the results are the same. The system can cater for up to 250 users, or 50 agents, or a simple ratio in between.

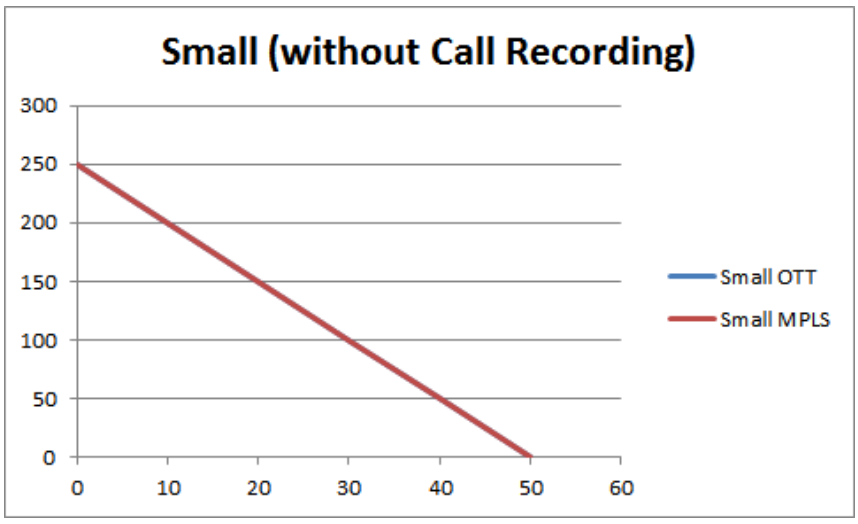

For the cUC Small deployment, without call recording, the results for the number of users to agents is the same for both OTT and MPLS deployments. The OTT results are overwritten by the MPLS results, and are the same.

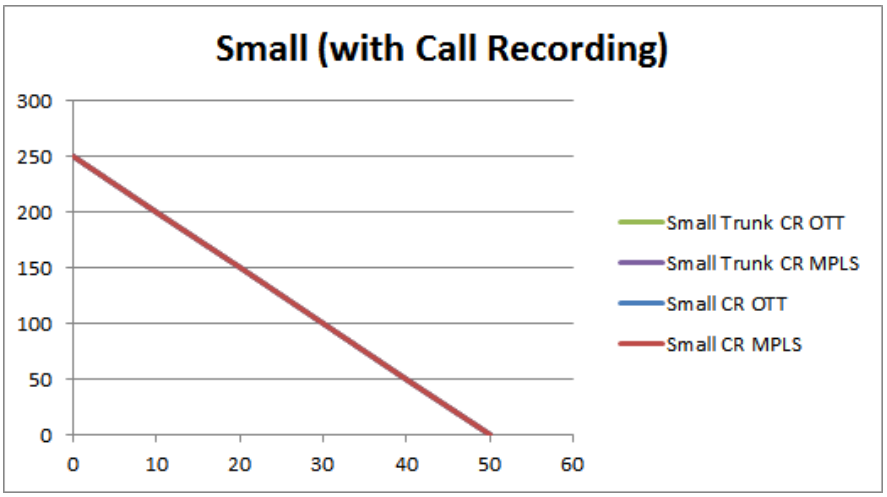

In both graphs, all results line up and only one line appears, even though multiple lines are identified in the key, i.e. for Small Trunk CR OTT and Small CR MPLS are the same.

cUC MidSmall Infrastructure

For the cUC MidSmall deployment the results with call recording for OTT deployments results in a decreased number of agents, compared to the MPLS deployment.

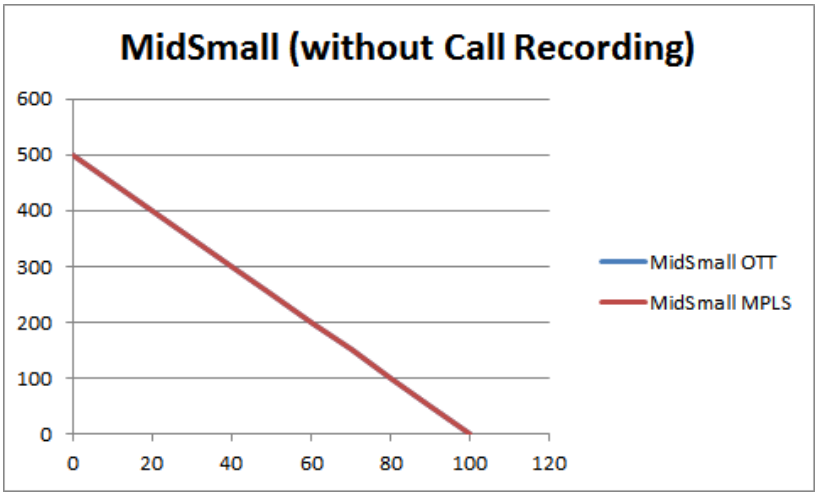

For the cUC MidSmall deployment, without call recording, the results for the number of users to agents is the same for both OTT and MPLS deployments. The OTT results are overwritten by the MPLS results and are the same.

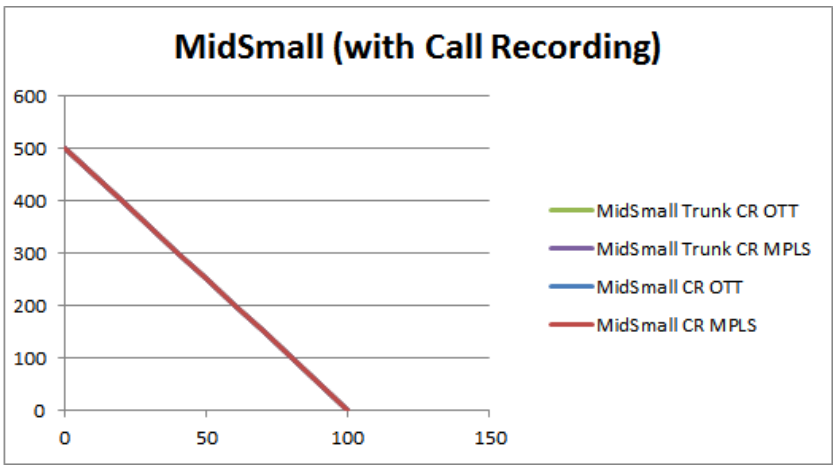

In both graphs, all results line up and only one line appears, even though multiple lines are identified in the key, i.e. for MidSmall Trunk CR OTT and MidSmall CR MPLS are the same.

cUC Medium Infrastructure

For the cUC Medium deployment, there is a reduction in the number of supported agent, with call recording, compared to the MPLS deployment:

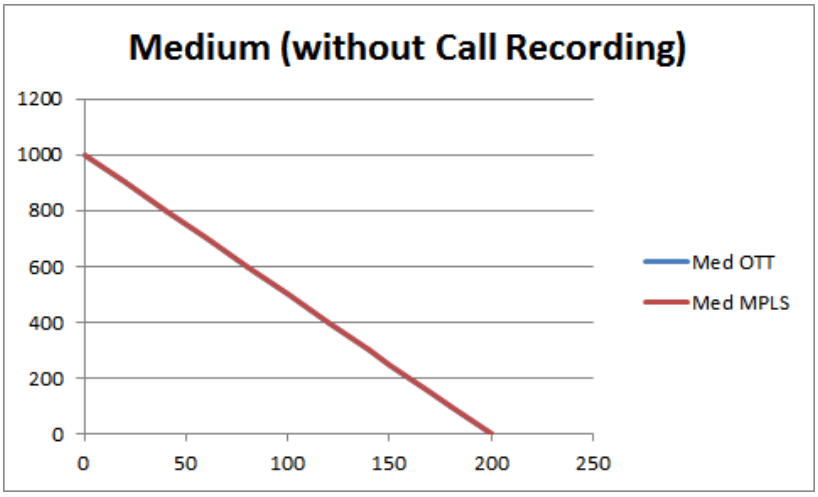

For the cUC Medium deployment, without call recording, the number of agents to users is the same for OTT and MPLS deployments. The OTT results are the overwritten by the MPLS results and are the same.

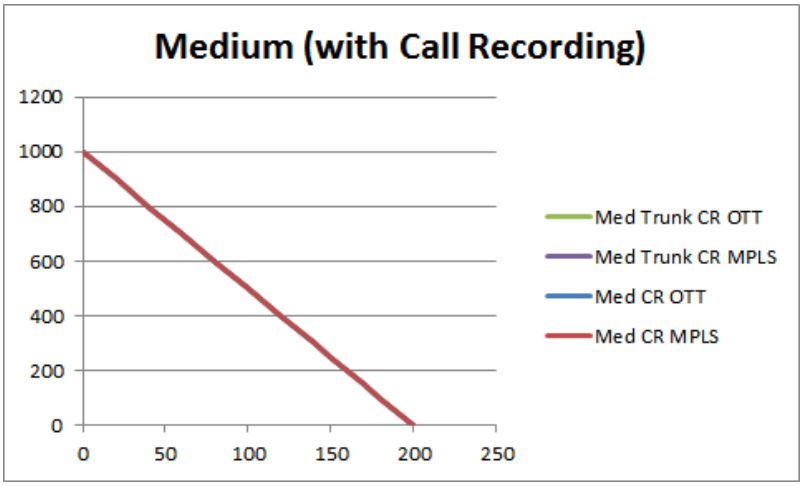

In both graphs, all results line up and only one line appears, even though multiple lines are identified in the key, i.e. for Medium Trunk CR OTT and Medium CR MPLS are the same.

cUC Large Infrastructure

The Large infrastructure includes a number of limits specifically around the number of users that can be supported in an OTT deployment, as well as differences depending on whether call recording is enabled, or not.

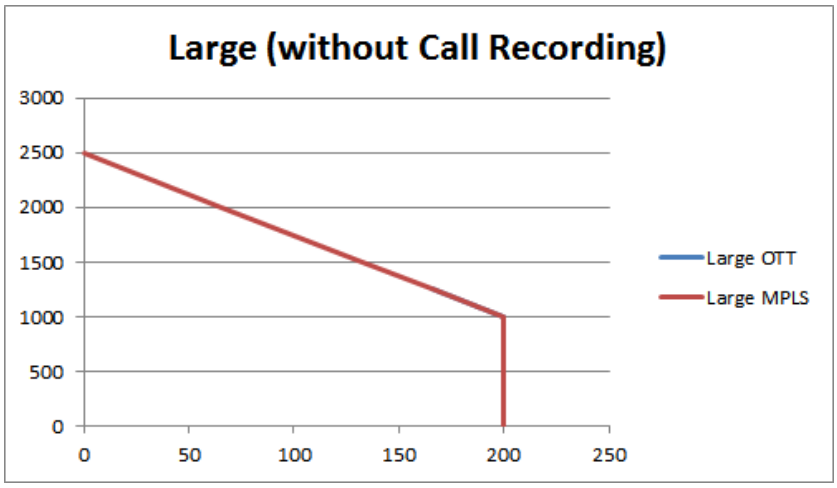

For this configuration, without call recording, the upper number of OTT users is 2500. Agents are limited to 200. The number of users to agents is scaled accordingly, with the possibility of 1000 users *and* the maximum of 200 agents

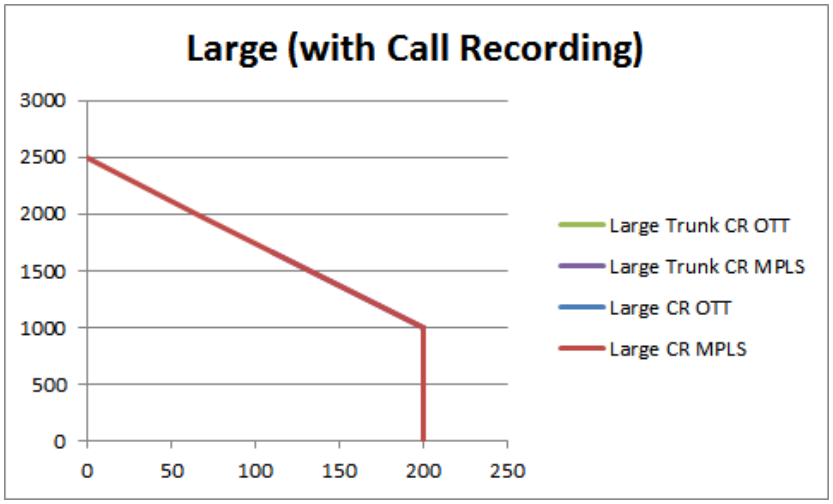

For this configuration, with call recording, the upper number of OTT users is 2500. Agents are limited to 200. The number of users to agents is scaled accordingly, with the possibility of 1000 users *and* the maximum of 200 agents.

In both graphs, all results line up and only one line appears, even though multiple lines are identified in the key, i.e. for Large Trunk CR OTT and Large CR MPLS are the same.

Feature Availability Comparison between OTT and MPLS Connections

This tables below will help quickly identify if a feature the customers requires will work Over-The-Top (OTT - Internet connected) or whether MPLS is required to deliver the functionality required.

Contact Center

Flex on Google Cloud Feature	OTT	MPLS	Notes
Email Queue as part of Multimedia	•*†	•*	* 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

Flex on Google Cloud Feature	OTT	MPLS	Notes
Chat Queue as part of Multimedia	•	•	* Chat is CloudLink based and uses OTT connections, and requires Internet access
SMS Queue as part of Multimedia	•*	•*	* The customer will need to sign up to Twilio for the SMS service. The implementation requires Mitel Professional Services
Contact Center Client	•	•	
Ignite Web Client	•	•	
IVR Customer Onsite Database Lookup	–	•	A secure on-net path to the customer DB is required. This does not affect SFDC.
IVR Customer Web/ Integration Onsite Server	–	•	If the Web/ Integration server is onsite there must be path for OTT services to reach the server

Devices

Flex on Google Cloud Feature	OTT	MPLS	Notes
Survivable Gateway	–	•	On-premise remote and survivable (secondary controller) gateways is available with MiCloud Flex on Google Cloud solution
IP Phones (53xx Series)	•	•	5300e series and 5304, 5312, 5324, 5320. Earlier 5302, 5330 and 5340 (non-'e') are not supported.
IP Phones (69xx Series)	•	•	
MiVB Console	•	•	
MiCollab Softphone	•	•	Desktop client (SIP) or Web Client (WebRTC)
112 DECT and RFP12	•	•	Restrictions on number of FQDN characters may limit deployments (limited to 63 characters)
IP DECT	-	•	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 premise.
SIP DECT 600	-	•	Only Onsite OMM is supported. It is not possible to host the physical base-station in the Data Centre. This needs to be deployed on the customer premise.
WebRTC Clients	•	•	Note that WebRTC clients are not resilient. Nor can they be recorded directly with call recording/WFO MIR.
ATA TA71xx	•	•	ATA device sourced through Partner

Features

Flex on Google Cloud Feature	OTT	MPLS	Notes
OnSite CRM Integration	–	•	A secure on-net path to the customer CRM is required
Active Directory Integration / Support	–	•	Mitel does not currently support Active Directory Federation Services (ADFS). Office365 AD is only supported through use of an onsite AD server from the customer.
Email Integration	–	•	EMEM does not currently support SMTP authentication. Limited to onsite deployments only. Use of an inline e-mail forwarder, to provide authentication, may be required.

Interconnection

Flex on Google Cloud Feature	OTT	MPLS	Notes
Hybrid Mode – Multiple sites, each with their own system. Ability to connect cloud to premise systems	–	•	Not currently available with MiCloud Flex on Google Cloud
Syslog	–	–	Not available

Call and Screen Recording (WFO MIR)

Flex on Google Cloud Feature	OTT	MPLS	Notes
Call Recording IP Stations Internal	•	•	And/Or associated Trunk side call recording
Call Recording IP Stations External	•	•	And/Or associated Trunk side call recording
Screen Recording	–	•*	* 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.
PCI Stop / Start Recording	•*	•	Call Recording agent available OTT and MPLS. Screen Recording only available MPLS.
Playback (Audio) and scoring	•	•	Scoring is provided via a license option: "WFO MIR Quality Monitoring"
Playback (Video)	–	•	Only available for MPLS
Call Recording Archive	•	•	Google Cloud is used for storage. Archive server is not available.

Resiliency and Operations

At the heart of the resilient operation is the deployment of the applications and network access via the PoP, between two regional data centres. Resiliency provides for continued operation of the core voice services and customer network access via OTT and

MPLS. Applications that provide additional functions may not be resilient. Typically, these will be reporting and management applications, some of which are described below.

An overview of which functions and features that are resilient is provided under System and Application Availability.

The diagram below provides an overview of the network connectivity for the resilient configuration.

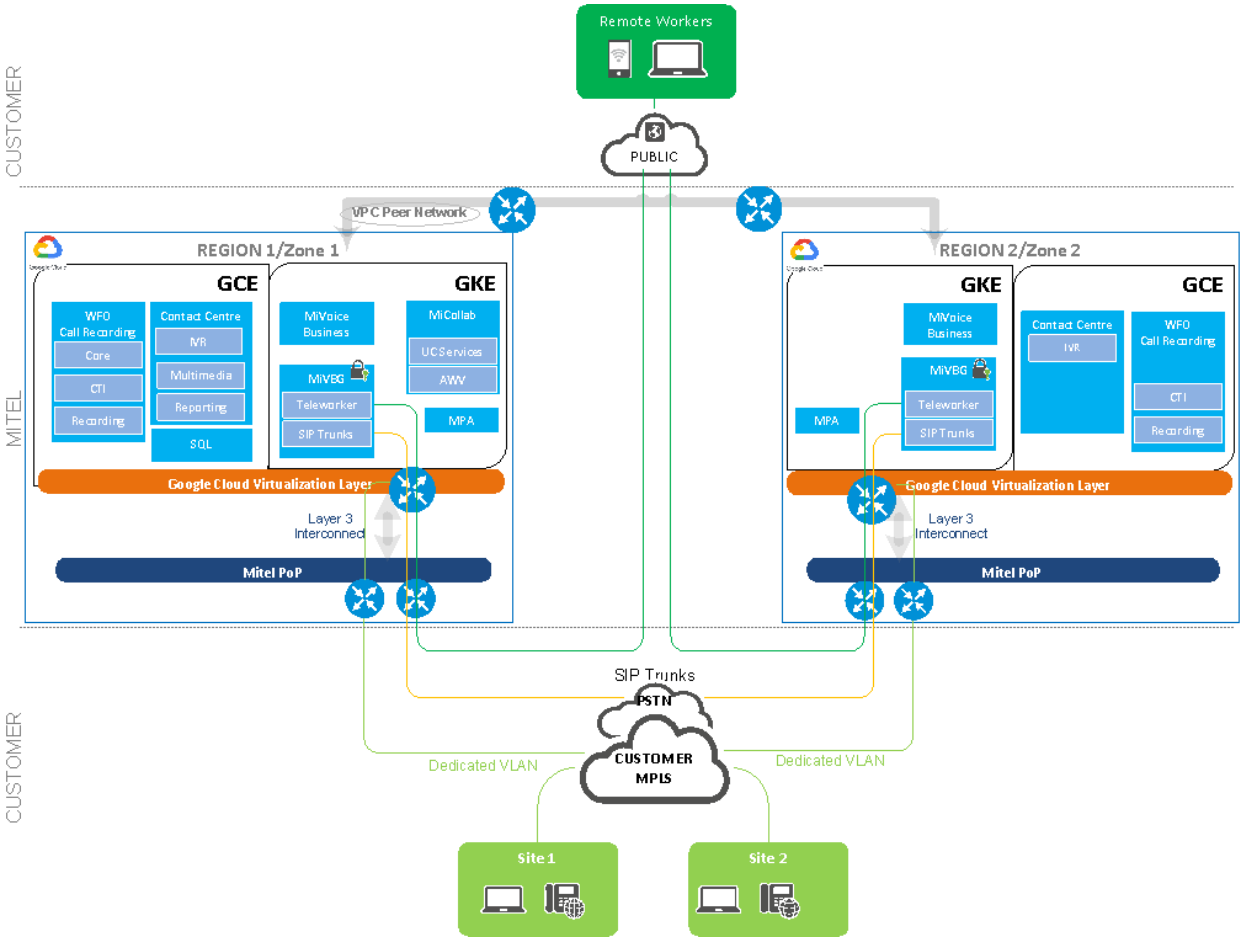

All access to the solution is duplicated. The infrastructure within the PoPs is duplicated as well as duplication of access to the applications and services via multiple PoP connections. Google can provide multiple data centres (zones) within a region and multiple regions within a geographic area. However, there is a limited availability within a region. Only by deploying across multiple regions can the target of 99.999% availability become achievable. Note that not all countries offer multiple data centre regions, and within a geographic area, multiple regions may be located within different countries. Some deployments may only be able to exist within multiple zones within a region. In this case the expected availability target will be 99.99%.

The two regional deployments are connected to the same Virtual Private Cloud (VPC), which extends between the two regions via the Google backbone. The regions exist in different subnets, so any connection is at Layer 3 and routable, the same as an on-premise deployment across a WAN link.

The endpoints run heartbeat protocols to determine access to the underlying applications. Should this heartbeat fail, a set of protocols triggers the device to identify the alternative secondary connection point. It is possible that during this transition, there may be minor loss of service, or trunk disconnection, depending on the level of the outage. Streaming connections between phones, or between phones and SIP SP will continue, if the network path exists. The devices will then re-home once the call is cleared. Recovery back to the primary application and services is managed to ensure transition occurs during an idle

period. Maintenance and updates uses a similar failover and recovery mechanism, however, in this case the failover is managed to ensure this occurs during an idle period.

Within a region, Google may move workloads between different servers and also between zones within a region. A region is effectively a layer 2 infrastructure and any such moves can be achieved very rapidly. Any interruption to service is minimal.

It is up to the partner to ensure that for MPLS connections, a connection is available at each of the PoP locations. In the event that only a single connection is achieved, there is a possibility of increased outage probability if the one connected region fails. Although this might be considered minimal risk, Google publishes Service Level Objectives (SLO) for regional and data centre outages, i.e. outages are expected. In such a condition the applications are subject to the same outage probability. In this case it is expected that service will only be possible up to 99.9% availability.

The following table describe some of the operation of the different applications during a resilient configuration:

Application	Primary	Secondary	Resilient	Primary Service outage outcomes
MiVoice Business Call Control	Yes	Yes	Yes	Voice Service switches to secondary controller (including remote survivable on-premise gateway, when provisioned)
MiVoice Border Gateway	Yes	Yes	Yes	Voice Service switches to secondary gateway
Mitel Performance Analytics and Management	Yes	Yes	Yes	Dual deployment allows access to both primary and secondary regions
Business Analytics	Yes	Yes	Yes	Deployed with MiVB in GKE. Connected to service hosted in Microsoft Azure Cloud
MiContact Center (MiCC)	Yes			Available on recovery of primary service. Deployment of a Collector Node at the secondary data centre (located on the IVR server) is recommended
MiCC Interactive Voice Response (IVR)	Yes	Yes	Yes	Voice Services and call routing continue in service. The SIP SP must also be configured to route to both primary and secondary gateways. Manual override of the IVR hunt group is also recommended (i.e. call rerouting)
MiCC SQL	Yes			SQL server is associated with the MiCC Enterprise server.
Workforce Management	Yes			Deployed with MiContact Center Business Enterprise Server. Connects to service hosted in Microsoft Azure Cloud
Call Recording Playback	Yes			Call playback and access to the database will not be available. Available on recovery of primary service
Call Recorders	Yes	Yes	Yes	Call recording of voice, and CTI connections (call status), continue in service. Recordings are cached until the core storage and playback server returns to service, then data is uploaded for future storage and playback access

Application	Primary	Secondary	Resilient	Primary Service outage outcomes
Call Recording User Client	Yes			Requires connection to WFO MIR core server for playback and other control. Call recording will continue otherwise. Ability to stop/start Call Recording may be impacted.
Screen Recorders	Yes			Screen recordings are available on recovery of primary service
MiCollab Softphone Call	Yes	Yes	Yes	Softphone will re-home to primary or secondary region. Calls can be made and received, provided they were previously logged in.
MiCollab Softphone Directory Calling	Yes			Directory lookup may be impaired for new searches. Cached information may still be used.
MiCollab Chat	Yes		Yes	MiCollab Chat redirects to use the CloudLink Chat on AWS. Once initial registration is established, connection is direct to CloudLink chat.
MiCollab Presence	Yes			MiCollab Presence information will become available on recovery of the primary service.
MiCollab AWV Collaboration	Yes			MiCollab Collaboration will become available on recovery of the primary service. MiCollab Collaboration using MiTeam and CloudLink will continue to be available
MiTeam Meetings	Yes		Yes	MiTeam meetings are provided via CloudLink MiTeam Meetings application. MiCollab is used primarily for management and initial configuration.
MiNET Phones	Yes	Yes	Yes	Service moves to secondary gateway or controller
WebRTC clients	Yes			Client would need to register with secondary gateway. Dependent on client.
3rd party SIP phones	Yes	Yes	Yes	For phones that have been tested for interoperability and that support DNS-SRV, service switches to the secondary gateway or controller

Further product details can be found in the sections below.

MiVoice Business

The MiVoice Business (MiVB) Call control operates in both the primary and secondary data centres. These are linked by a management protocol that maintains synchronisation of the database. Synchronisation of the call state is not provided. An end-device is provisioned on both controllers, so failover between the two controllers will occur with minimal user impact.

EMEM (Embedded Voice Mail)

Voice mail is provided by the MiVoice Business Call control. Voice mail services are configured and provided for the user in the two regional locations, with two different storage locations. The two locations are linked together through the visual voice mail application.

This application takes the status information from both locations and combines these to a common view. Dual SIP Trunk connections are provided, so in the event one region is inaccessible, such as the primary location, calls will be directed to the secondary location. In this case the secondary EMEM service will handle the call and storage.

During a failover condition the primary region may not be accessible by the user. During this period, access to the recorded calls on EMEM, at the primary location, may not be availabel, nor visible in the application. On recovery of service, access and visibility of the voice mails at both locations will be returned.

Auto Attendant functionality should be duplicated in the MiVoice Business programming in each region.

UCC licensing includes a voicemail license for primary and secondary controllers.

MiVoice Border Gateway

The MiVoice Border Gateway (MiVBG) components are clustered and share database information. There is both an internal network and external network connection. In the event one application fails, the other is available to take over registrations and service. The MiVBGs are scaled such that either primary or secondary locations can handle the expected call traffic load.

Teleworker phones and end devices will use information provided to them, or they will use DNS-SRV to identify the two gateways. In the event that one unit cannot be reached, a connection will be established via the secondary unit.

MiCollab

The MiCollab itself is not resilient. Certain functions may not be available if access to the primary location is not possible. These would include MiCollab Client presence information and access to Audio-Web-Video (AWV) Collaboration tools. However, user of Chat will be available, as will access to MiTeam Meetings on Amazon Web Services (AWS). Once access to the primary site is available, the MiCollab services will resume.

Softphones that linked to MiCollab will loose MiCollab Client presence information. However, the softphone, and the capability to make calls will still be available, as these will link directly into the MiVB Call Control.

MiCollab Unified Messaging feature set is not included with the MiCloud Flex on Google Cloud solution.

Mitel Performance Analytics

A Mitel Performance Analytics (MPA) management probe is provided 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.

Business Analytics

A Business Analytics probe is provided in both primary and secondary locations. In the event the primary site is inaccessible, yet somehow calls are made to or from this site, the probe will store the information locally. Once the primary site recovers, the probe will reconnect to the hosted cloud service and forward the cached information. In this way the data is not lost, it just may not be immediately available.

MiCC Contact Centre

The MiCC Business application is not resilient. In the case of an outage, access to the unit may be limited. However, the remote Interactive Voice Response (IVR) servers will continue to handle incoming calls, with some reduced capability. Core call handling and voice streaming will continue, albeit with some reduced capability. The remote IVR must be configured with cross linked hunt groups so that call calls can be directed to either remote IVR to handle possibility that one of the IVR units is also impacted by the outage.

Interaction Recording and Speech Analysis

Call recording is configured for resiliency. However, screen recording is only configured for the primary location. In the event that the primary region is unavailable, screen recording will also be unavailable. The Core server is also only located at the primary location. The call recording servers are duplicated in each regional location and matched to duplicate MiVoice Border Gateway (MiVBG) and Secure Recording Connector (SRC) functions in each regional location.

In the event that the call recording server cannot reach the Core server, it will continue to record calls and store these locally. Each of the MiVoice Business (MiVB) controller also has an associated Computer Telephony Integration (CTI) server. This server determines when calls to be recorded are being created and will direct the appropriate recording server when to record the calls, or the screen. Once the connection to the Core server is recovered, the stored recordings will be forwarded for any analysis and local storage. The Core server includes a de-duplication service which is applied to the recordings. In the event that a call is recorded by multiple recorders, this service will remove duplicates to ensure efficient use of storage and search resources.

Playback client access is required to the Core server. Playback services will not be available if access is not achievable to the Core server. However, services will resume when connections recover.

Transcription and Keyword spotting are services that run on servers located at the primary location. They connect to the Core server to access recording files, process these and re-store them. In the event that connection to the primary location is lost, it is possible that these services will continue to operate. However, user access will not be available. Recovery of the connection will recover client access to these services.

MPLS to OTT failover

Although this is touted as a possible alternative to dual MPLS connections, in reality it negates any of the benefit of any existing MPLS circuit. This is because in order to take advantage of this service the end-devices need to be considered, and configured, as Teleworker devices. In order to gain access via the MiVoice Border Gateway(s), the phones need to be included in the registration database. Unidentified devices are not granted access. Therefore, the end-devices are Teleworker.

Since the end-devices are programmed as Teleworker devices, they will register to the public IP address(es) of the MiVBG units. Thus, any connection will go via the customer Internet Gateway, and not necessarily via the MPLS connection. The MPLS connection is intended to securely handle internal private addresses. However, a connection from an end device that registers with a Public IP will appear to the MiVBG as an external device and will be identified by the customer public IP address, the phone internal address being NATed by the customer's internet gateway.

Currently this configuration is not supported.

WebRTC Clients

WebRTC clients are not resilient.

Availability with a Single region and Dual Zones

Google indicates that a single region will offer 4 '9's of availability, or 99.99%. Based on this it is determined that a single zone within a region offers around 99% availability. It also infers that although it is possible to get close to 4 '9's availability, it is not quite possible to achieve. Therefore, a deployment with two zones within a single region can offer 3 '9's availability, albeit a high 3 '9's. This equates to an unexpected outage period of around 8 hours, or less, per year. For applications that are running in a single zone, the availability is only 2 '9's. When the solution is deployed within a single region, it will use two zones to improve availability for the core voice applications. Users and installers should be aware of the potentially longer outage period that may result from use of a single region, compared to the standard dual region deployments.

Onsite Gateway

MiCloud Flex on Google Cloud customers who connect to MiCloud Flex through MPLS private networking via Point-of-Presence (PoP) may choose to deploy an onsite gateway. This gateway can support local analog extensions as well as local functions such as paging adapter support or music on hold. It can also support local PSTN connections via PRI E1, PRI T1, or analog trunks.

IP phones located at this site, registered to MiCloud Flex, can optionally fail over to this gateway instead of to a controller in a secondary data center. This would enable local voice extension calling at the gateway site should the primary controller become unavailable. Resilient voicemail service can also be provided by the onsite gateway using the embedded voicemail capabilities of the gateway. Configuration of resiliency is on a per-set and per-user basis. Some IP phones and users at the site can be configured to failover to the onsite gateway while others can be configured to fail over to the secondary data center.

The onsite gateway is a Mitel MiVoice Business controller running Release 9.1 or higher. This minimum release is required in order to support Fully Qualified Domain Name (FQDN) addresses.

The onsite gateway can be monitored from the MiCloud Flex MPA.

The hosted Mitel Business Analytics service can be used for call analysis, using a connection from the collector client at the hosted site to the remote gateway. If the network connection to the hosted site is lost, the CDR data will be cached and sent to the client once the connection is re-established.

MiCloud Flex users at this site are registered to the MBGs or controller in the MiCloud Flex deployment They receive their primary voice services from MiCloud Flex. In a failover scenario. If the MiCloud Flex users are programmed to fail back to the onsite gateway, they can then receive their PSTN services from the onsite gateway.

Analog users connected to the onsite gateway, as well as IP set users who are not MiCloud Flex users, may receive their PSTN services from the onsite gateway or via the centralized trunk services in the cloud depending upon configuration and licensing.

Onsite Gateway Scenarios

The following diagrams show the components within the gateway. There are three possible scenarios, depending on the customer.

No Existing Controller

In this first scenario, the customer does not own any Mitel MiVoice Business controllers prior to becoming a MiCloud Flex with Onsite Gateway customer.

In this case, the customer may purchase or lease a Mitel MiVoice Business EX or AX controller.

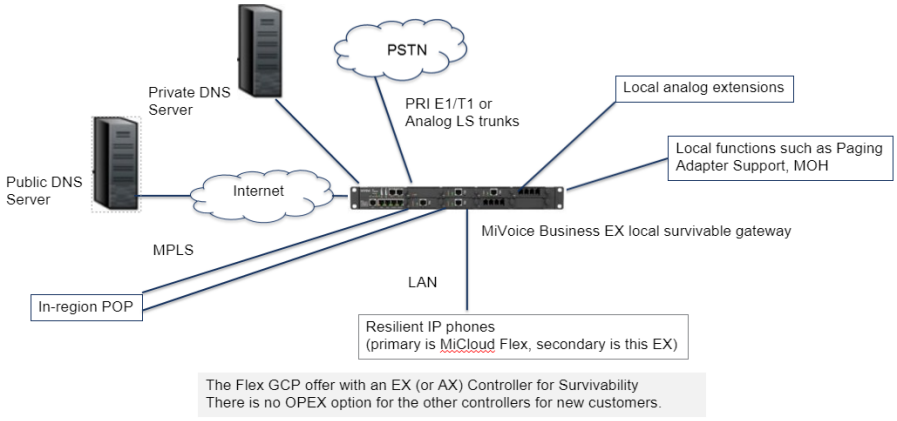

Existing MiVoice Business Controller Re-Purposed as a Gateway

In this second scenario, the customer already owns a Mitel MiVoice Business controller that can be upgraded to release 9.1 and above.

In this case, after re-licensing, the customer may re-use their existing controller as an onsite gateway.

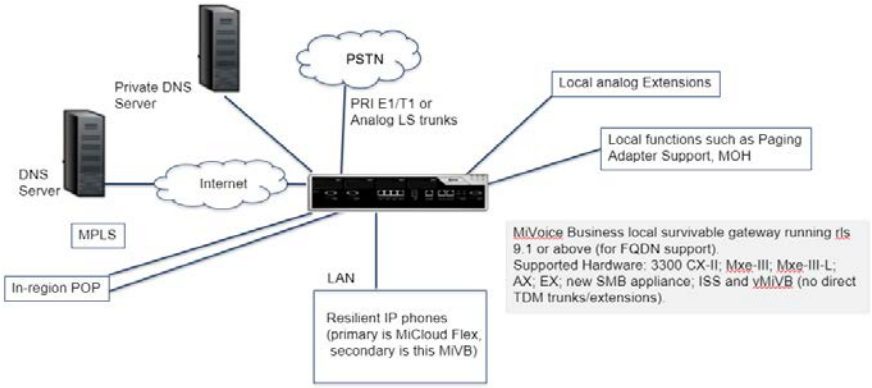

Front Ending Existing Controller that Cannot Be Upgraded to MiVoice Business 9.1

In this third scenario, the customer already owns a Mitel MiVoice Business controller that cannot be upgraded to release 9.1 and above.

In this case, the customer will need to buy or lease a controller that will act as the gateway between the existing hardware that cannot be upgraded and the MiCloud Flex instance. They may choose to retain or remove their existing controller. In many cases it may make sense to migrate the users to the new controller in a time period determined by the customer.

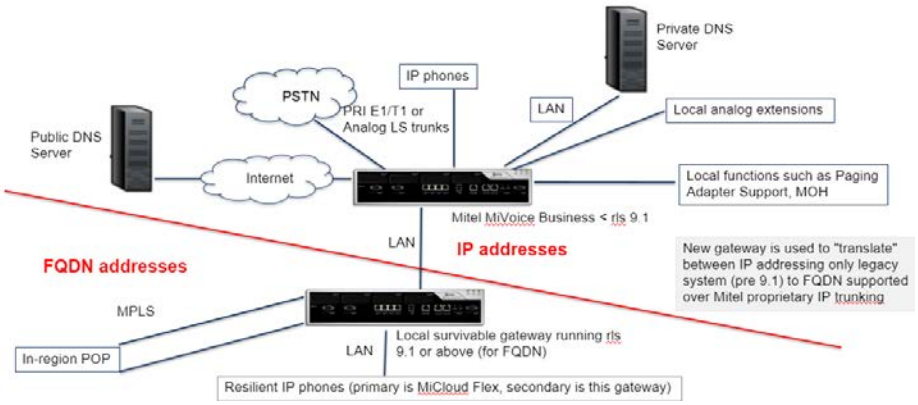

Phone Type and Location	Outage type: Teleworker MBG unreachable	Outage type: SRC MBG unreachable	Outage type: cMiVB unreachable	Outage type: MiCollab unreachable	Outage type: Private Network unavailable (complete outage)
Desk phones on network - registered to cMiVB	Not applicable	Not applicable	Failover to onsite gateway for those programmed this way.	Not applicable	Failover to onsite gateway assuming path is available
Teleworker desk phones - registered to cTeleworker MBG	Up - goes to secondary MBG in the data center	Not applicable	Up - the cTeleworker connects to on-premise MiVB. Failover to onsite gateway for those programmed this way.	Not applicable	Not applicable
Any phone that is registered to or via cSRC MBG	Not applicable	Up - goes to secondary SRG MBG in DC	Failover to onsite gateway for those programmed this way.	Not applicable	TW up; On network - down.
MiCollab SIP softphone on network - cMiVB	Not applicable	Up - goes to secondary SRG MBG in DC	Failover to onsite gateway for those programmed this way.	Up - the user loses presence and AWV. Telephony and chat still work.	Up - the user loses presence and AWV. Telephony and chat still work.
MiCollab SIP softphone off network - registered to cTeleworker MBG	Up - goes to secondary MBG in DC	Up - goes to secondary SRG MBG in DC	Failover to onsite gateway for those programmed this way.	Up - the user loses presence and AWV. Telephony and chat still work.	Up - the user loses presence and AWV. Telephony and chat still work.

WebRTC softphone always registered to cTeleworker MBG	Down - WebRTC is only supported on a single MBG	Not applicable	Down	Down	Up

Onsite Gateway Resiliency Scenarios

The following table shows the cases where the Onsite Survivable Gateway can be used to provide voice survivability.

Onsite Gateway Limitations and Best Practices

- Controller Hardware: New customers will purchase or lease an EX or AX to be the onsite gateway. Existing customers with controllers that can be upgraded to MiVoice Business 9.1 and above may use them as MiCloud Flex Onsite Survivable Gateways. Please see the list of supported controllers in the MiVoice Business Migration Guidelines document.
- IP Set Hardware: The introduction of Onsite Survivable Gateway does not alter the list of supported IP phones that can register to MiCloud Flex. The onsite gateway will still support the same phones as an on-premise MiVoice Business of the same release. (Some limitations apply to older 5300 devices in this configuration)
- Software: The onsite gateway software should be running release 9.1 or newer because FQDN support is required for this solution. It is recommended to upgrade the onsite gateway software to the same MiVoice Business release as the MiCloud Flex on Google Cloud offering. (MiVB Release 9.1 SP1 is required to support 53oo series phones - some limitations may apply to earlier versions of this series)
- PSTN Connectivity: The onsite Gateway supports PSTN connectivity via local analog and E1/T1 links. It does not currently support PSTN connectivity via SIP trunks.
- Connection to MiCloud Flex: The connection between the onsite gateway and MiCloud Flex on Google Cloud is via Mitel proprietary IP trunks. These trunks need to be enabled at firewalls to allow calls on the local gateway to be routed to other phones that will be registered with MiCloud Flex. This connection must support FQDNs.
- Embedded Voice Mail (EMEM): Please note that there are different capacities and functionality between what is supported by EMEM in MiCloud Flex and what is supported by EMEM in an onsite gateway. For further information on capacities, please refer to the MiVoice Business Engineering Guidelines.
    - Mitel 3300 ICP platforms support 30 or less Voice Mail ports, versus up to 120 in MiCloud Flex and other x86 platforms.
    - Mitel 3300 ICP platforms support 750 or less Voice Mail boxes (see MiVB Engineering Guidelines for individual platform limits), versus up to 5000 in MiCloud Flex and other x86 platforms.
    - The "Call Flows" feature is currently only available on MiCloud Flex on Google Cloud.
    - EMEM Auto Attendant and RAD programming on the gateway is a manual process.
        - The Hunt Group pilot number will most likely be different than the one initially setup via ICW on the resilient MiCloud Flex pair. Should a customer want the HG pilot number to be the same on ALL the systems in the cluster, they will need to mark the HG pilot number on ALL systems as Local Only.
- Local MBGs are not currently supported at the local gateway. All MBG functions are provided by MBG at hosted locations.
- Mitel Performance Analytics (MPA): The MPA functionality in MiCloud Flex can monitor the onsite gateway, subject to the maximum number of monitored systems by the MPA

probe that resides in the MiCloud Flex data center. Please refer to the MItel Performance Analytics Engineering Guidelines for the current limits. If the link between the gateway and MiCloud Flex fails, the MPA will report the loss of connection, even though the gateway may still be up and running.

- MiCollab Visual Voice Mail integration is not supported on the onsite gateway.
- It is recommended to update the version of the onsite MiVoice Business gateway to match the version of the MiVoice Business in MiCloud Flex on Google Cloud. At a minimum this must be running MiVB Release 9.1, and MiVB Release 9.1 SP1 when 5300 phones are used.