ALCATEL-LUCENT RAINBOW™

Rainbow Edge Offer

TECHNICAL OVERVIEW Ed 05

JUNE 2020

Author: Operations - Cloud Services

Alcatel·Lucent
Enterprise

Contents

Glossary

ALE: Alcatel-Lucent Enterprise

CPE: Customer Premises Equipment

DC: Data Center

DMZ: Demilitarized Zone

FQDN: Fully Qualified Domain Name

GK: GateKeeper

HA: High-Availability

MPLS: Multiprotocol Label Switching

NAT: Network Address Translation

PBX: Private Branch Exchange

PoP: Point of Presence

QoS: Quality of Service

RTT: Round Time Trip

(S)DRS: (Storage) Distributed Resource Scheduler

SLA: Service Level Agreement

SSL: Secure Sockets Layer

TURN: Traversal Using Relays around NAT – RFC 5766

VPN: Virtual Private Network

1   Introduction

This guide provides some technical overview of Rainbow Edge offer. It provides customer various inputs on what is the best Edge offer to be selected and how summarize the global architecture, infrastructure and underlying network environments to be put in place.

Before moving forward with Rainbow edge, the End Customer should consider two options of the public multi-tenant cloud service. The shared technology offers a secure solution while benefiting from a very competitive total cost of ownership.

a.   Rainbow manages several visibility profiles. See the complete description here https://support.openrainbow.com/hc/en-us/articles/115001397924-How-to-Change-the-Visibility-of-My-Company- . With an "isolated" company the End Customer can get a fully protected environment. And he could decide if certain selected employees could be connected to the open world.

b.   On top of the previous setting, the company could also connect a private file storage system to Rainbow to ensure the shared files are stored in a data center controlled by the end customer. More information here https://support.openrainbow.com/hc/en-us/articles/115001019370-Customer-Premises-Storage-Solution-Brief

2   Overview

Alcatel-Lucent Enterprise (ALE) is introducing Alcatel-Lucent Rainbow, an overlay cloud service operated by ALE. Rainbow offers contact management, presence, persistent messaging, audio/video, screen and file sharing, with PSTN termination and API openness to integrate with existing customer PBXs, machines and apps.

Rainbow is a global Cloud services available in all countries with hardened methods implemented to ensure data privacy and geographical data isolation and segregation. For specific use cases, Rainbow features and extended variant called "Edge", allowing for Rainbow services and data to be hosted in a customer-based or partner-based datacenter. Edge is not a commercial acronym or denomination of a new Rainbow service but a local instance of Rainbow services, part of a broader mesh.

3 History

Modifications	Date	Edition
World map update (new ANZ region) and Swift requirement for standalone low-scale deployments.	30/06/2020	Ed 05
Renamed "Rainbow Spot" to "Rainbow Edge". Added extra section related to extra "Rainbow Voice" features.	12/03/2020	Ed 04
Cosmetic changes, note related to Rainbow Voices features not being part of the Rainbow Edge scope.	13/02/2020	Ed 03
Reviewed, Extended Introduction, Smart Hands, Mail Server, Debian OS	10/02/2020	Ed 02
Initial Revision	04/02/2020	Ed 01

4 Related Documents

- Rainbow – Solution Brief – Infrastructure Abstract
- Rainbow – Solution Brief – Security Abstract
- Rainbow - Network Requirements

5 Rainbow Edge Variants

Rainbow is a global UCaaS and CPaaS service provider, sustained by datacenters in multiple regions around the globe, as described in the diagram below.
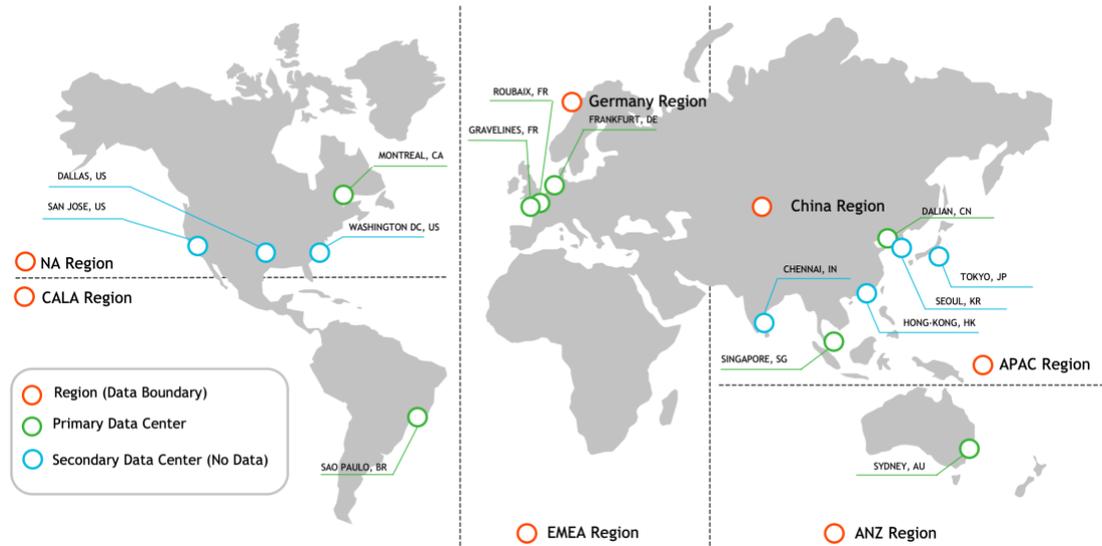
Doing so, Rainbow guarantees:

- Performances to end users, ensuring to have primary or secondary datacenters as close as possible from its users, as to minimize global latency.

- Data privacy and data geographical isolation, ensuring that user data are stored in the region of users' belonging and not replicated globally.

For specific use cases, Rainbow Edge Extensions can be offered to provide more data isolation. Amongst the various possibilities are:

- **Region/Country Local DC**, a purely ALE hosted and operated datacenter. ALE is regularly adding up new datacenters in various locations, based in business perspectives.

- **Connected Edge**, a partner or customer hosted datacenter, operated by ALE, offering private on-premises data storage while using global Rainbow applicative infrastructure.

- **Standalone Edge**, a partner or customer hosted datacenter, operated by ALE, offering private on-premises data storage and standalone applicative infrastructure.

Each approach features different pros and cons and needs to be carefully chosen by end customer.

	Region/Country Local DC	Connected Edge	Standalone Edge
Rainbow Relationship Engine & Services	Global	Global	Restricted to partner's scope
Coverage (Perfs, Latency, Mobility)	Global Worldwide	Global Worldwide	Edge-Local Only Access
Runtime Operations	ALE 24x7x365	ALE 24x7x365	ALE 24x7x365
Availability	HA, Guaranteed by ALE	Limited, Shared with Partner	Limited, Shared with Partner
Infrastructure Resources	Mutualized	Dedicated	Dedicated
Access, Security, Logs	Mutualized	Dedicated	Dedicated
Data Location	Region-Bound	Edge-Bound	Edge-Bound
Data Storage	Mutualized / Segregated	Dedicated / Private	Dedicated / Private
Infrastructure & Network Costs	ALE Mutualized	Partner-Supported	Partner-Supported (>TCO)

Except from first option which is basically a 100% ALE operated datacenter, the first question to be addressed is should the customer Edge be connected or remain standalone.

Connected Edge will ensure users' connectivity (through federation) to Rainbow's global service, allowing interaction with global companies and will inherently benefit from our global coverage when it comes to performances.

Standalone Edge will provide customers with a private Cloud Rainbow instances (with or without public Internet connectivity) with no possible interaction with the global Rainbow network.

In both cases, the underlying infrastructure will be selected and provided by the customer, who will be in charge of its well behavior. Both the network and application setup, upgrades rollout and day to day runtime operations will remain held by ALE Cloud Operations staff.

Please bear in mind that at least partial Internet connectivity is required by ALE Operations team to handle 24x7x365 support.

6 Rainbow Edge Global Architecture

6.1 Highly Available Infrastructure

Rainbow internal datacenter's architecture comes with a generic "Frontend / Backend / Data" architecture, with infinite horizontal and vertical scalability and high availability, as described in the diagram below:
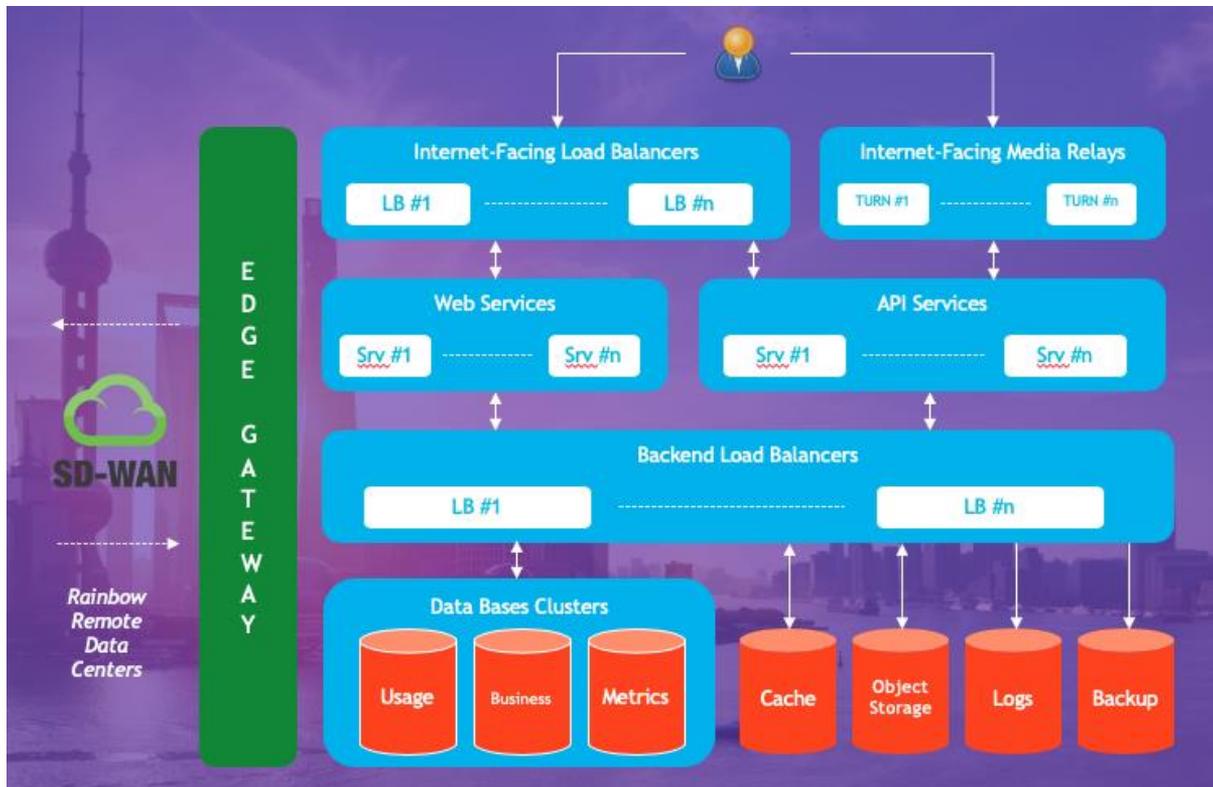
All datacenter services are connected to global Rainbow infrastructure through GateKeeper Edge Gateway (see next section for additional details). The link allows data and service federation with global Rainbow in case of Connected Edge offer and is restricted to ALE operational and administration maintenance in case of Standalone Edge.

Based on customer's service expectation, various kind of databases can co-exist as to handle infrastructure metrics (used for monitoring and serviceability), object storage (for file transfer features), logs (for application debugging and/or legal compliance) and backup (for resilience).

The matrix below depicts a mapping between data type and parts of the Open Source storage technology underneath. When it comes to Connected Edges, data are split between Global and Edge shards. When it comes to Standalone Edges, all data are stored in the same local databases.

Family	Type	Examples	Storage
User	Business	*Profile, Credentials, Settings, Companies, ACLs, Service Consumption …*	**Global MongoDB Shard for Anonymous Data**
			Edge-Local MongoDB Shard for Private Data
		Avatar/Picture	**Edge-Local SWIFT**
	Usage	*Messages, Contacts, Relationship, Call Logs, Activity, …*	**Edge-Local MariaDB**
		Files	**Edge-Local SWIFT**
	Activity Logs	*Connection timestamp, origin, destination, action …*	**Edge-Local ELK**
Service	Application Logs	*Server-side actions and behavior*	**Edge-Local ELK**
	Analytics	*Anonymous end-user usages correlation*	**Global MongoDB**
			Google Analytics
Infrastructure	Metrics	*System and network resources usage and consumption*	**Edge-Local + Globally Federated Prometheus**

Note that regardless of the data type and underlying storage technology, all data are encrypted at rest using AES256 mechanisms.

6.2 GateKeeper Edge Gateway

The GateKeeper (GK) component acts as Rainbow datacenters' edge gateway. It a full software-based network equipment that provides the following features:

- Local region (Edge) VLANs manager
- Local region (Edge) router
- Local region (Edge) firewall
- Local region (Edge) DHCP / DNS server
- Local region (Edge) software packages repository / mirror.
- Global SD-WAN like encrypted VPN tunnel for remote GK mesh and connectivity.
- Fully manageable through REST APIs

As stated, the GateKeeper components provides the basic networking services to all software components within the Edge.

It also provides network connectivity to global Rainbow platform through a WireGuard-based secure tunnel between other GateKeeper instances.

In the event of a Connected Edge, the tunnel will be used to exchange and federate service usage data between the Edge and the various datacenters from global Rainbow.

In the event of Standalone Edge, the sole usage of the tunnel will be to provide connectivity to remote administration tools and will require a limited bandwidth.

Network requirements and recommendation are to:

- Allow unrestricted outbound traffic to the Internet (to allow NTP / DHCP Relay / Packages Mirroring … to work)
- Block everything but WireGuard for inbound traffic from the Internet (UDP port to be chosen).

It is strongly advised to have GateKeeper redounded (as any other parts of the application are) as to ensure connectivity in case of failure.

For performances, QoS, quality and security reasons, GateKeeper inter-connected can be achieved through public Internet routing (default) or private segments (dedicated MPLS, mutualized or dedicated cross-connections through various PoPs like Equinix / Megaport …).

6.3 Software Licenses

Rainbow service is made out of multiple Open Source software and ALE proprietary applications. It does not involve any third-party software requiring any proprietary license to be purchased.

All servers make use of Debian GNU/Linux stable distribution as their native operating system.

The customer use of proprietary software solution for infrastructure hosting (e.g. VMware …) is however subject to license and is the sole responsibility of the customer.

7   Rainbow Edge Requirements

7.1 Generic Hardware Requirements

Rainbow services are hardware-agnostic and can be deployed on multiple IaaS providers:

- Bare Metal servers
- Virtual Machines on VMware vSphere
- Virtual Machines on Linux KVM
- Raw VM services on AWS (e.g. EC2), Azure, GCP …

As a rule of thumb, it is recommended for high-availability considerations to use (when technically and financially possible) two to three geographically distinct datacenters, yet inter-connected, and deploy the various servers' instances on each datacenter. When considering such an option, services are deployed in an active-active state, meaning that both datacenters will be actively used (by opposition to active-passive scenarios) and resources usage should not exceed half of the processing capabilities as to cope with failover mechanisms and high load.

The IaaS selection is bound to the customer and can even be a mix of both technologies providing that network requirements are met. The customer is responsible from the IaaS infrastructure, both from setup and runtime phases and must deal with any issues related to hardware or network layers.

Should the customer opt for virtualization technology, ALE recommends the following to be put in place. This relates to VMware's terminology but equivalent exists for other providers:

- Virtual Machines should run on modern hypervisors (minimum recommended is vSphere 6.5)
- VM and Storage High-Availability should be provided by vSphere DRS / SDRS technology.
- ALE recommends running at least 2 ESX hypervisors for High-Availability DRS support
- ALE recommends using distinct datastores: "standard-perf" and "high-perf" ones, with High-Availability SDRS support.
- Low-Level (Physical resources, i.e. hardware, disks, network connectivity, Internet bandwidth, power …) fully managed by Customer.

7.2 Generic Network Requirements

As for hardware requirements, Rainbow Edges come with some network specific requirements (or at least recommendations), namely:

- Should the customer host the service on multiple distinct datacenters, it is mandatory that the DCs share the same L2 segment for the different services to be part of the same private network and VLAN.
- Should the customer host the service on multiple distinct datacenters, it is mandatory that the round time trip (RTT) between the different DCs never goes above 100ms.
- ALE requires private -LAN addressing on 10.20/16 subnet.
- Customer has to provide public IPv4 on WAN (Internet facing) and WAN connectivity to the various servers. It is highly recommended that WAN IPv4 addresses are directly associated to the various servers and not mapped through NAT.

- When Internet-facing services are expected, it is recommended to have a /28 subnet (16 IPs).

- All servers are expected to provide at least 1 Gbps NICs. For servers requiring both LAN and WAN network access, it is highly recommended that 2 distinct NICs are being used.

- For large scale deployments, when a complete service redundancy and clustering support is expected, it is recommended that data hosting services (databases, object storage, logs and backups) feature a 10 Gbps NIC on LAN. While not being mandatory, it is highly recommended in case of hardware failure, to allow service restoration as fast as possible.

7.3 Operational Requirements

Due to the nature of Edges, service operations will feature a split responsibility:

- Customer will remain in charge of the functional state of the underlying infrastructure (hardware or virtual environments well-being) and associated network (both LAN/WAN) connectivity. ALE's contractual SLA won't apply if the root cause is verified to be low-level infrastructure related. See section 7.5 ("Smart Hands") for additional details.

- ALE will remain in charge of high-level services (virtual machines, OS, network configuration, applications, monitoring …). This will be fully managed by ALE. As to manage and reach our commitments, no logical access to the servers will be authorized to customer's or partner's IT staff, as to prevent any possible failure cause due to external interaction with Rainbow core services.

- Note that due to the complex nature of the Rainbow Edge, ALE provides a 24x7x365 Operational support and coverage, as an extra professional service.

7.4 Generic Hardware Footprint

Rainbow Edges hardware requirements are highly dependent on:

- The "Connected" or "Standalone" variants
- The expected volume of users (total and concurrent)
- The expected scalability
- The expected "extra" services (logs, metrics, backup, file sharing …)

That being said, below is the map of invariant network and remote-administration services that are mandatory for all Rainbow Edges instances.

Name	Count #VMs	CPU vCPUs	Memory Size (GB)	Disk (per server) Size (GB)	Disk (per server) IOPS	Network (per server) NIC Type	Network (per server) IPv4
Edge Gateway	2	2	2	64	Standard	WAN	1
				256	Standard	LAN	1 + VIP
Logs	1	4	32	1024	Optimized	LAN	1
Backup	1	1	2	64	Standard	LAN	1
				2048	Standard		
				2048	Standard		
Metrics+Monitoring	1	2	2	64	Standard	LAN	1
				512	Standard		

Please refer to further sections for extra hardware footprint required for "Connected" and "Standalone" Edge variants.

7.5 Local "Smart Hands"

Due to the shared service responsibility (partner providing infrastructure vs. ALE providing over the top service), ALE Operations team require local "Smart Hands" from partner. "Smart Hands" are an around-the-clock, on-site, operational support service for remote management, installation and troubleshooting of the data center's equipment. This is absolutely key for ALE Ops team to be able to maintain the global expected SLA.

8 Connected Edge

8.1 Network Topology

A Rainbow "Connected Edge" is meant to ensure customer's data storage within customers premises. Most of Rainbow application-level services will remain mutualized and hosted on Rainbow global core datacenters and only a limited subset of data-hosting servers will remain installed on customer's premises.
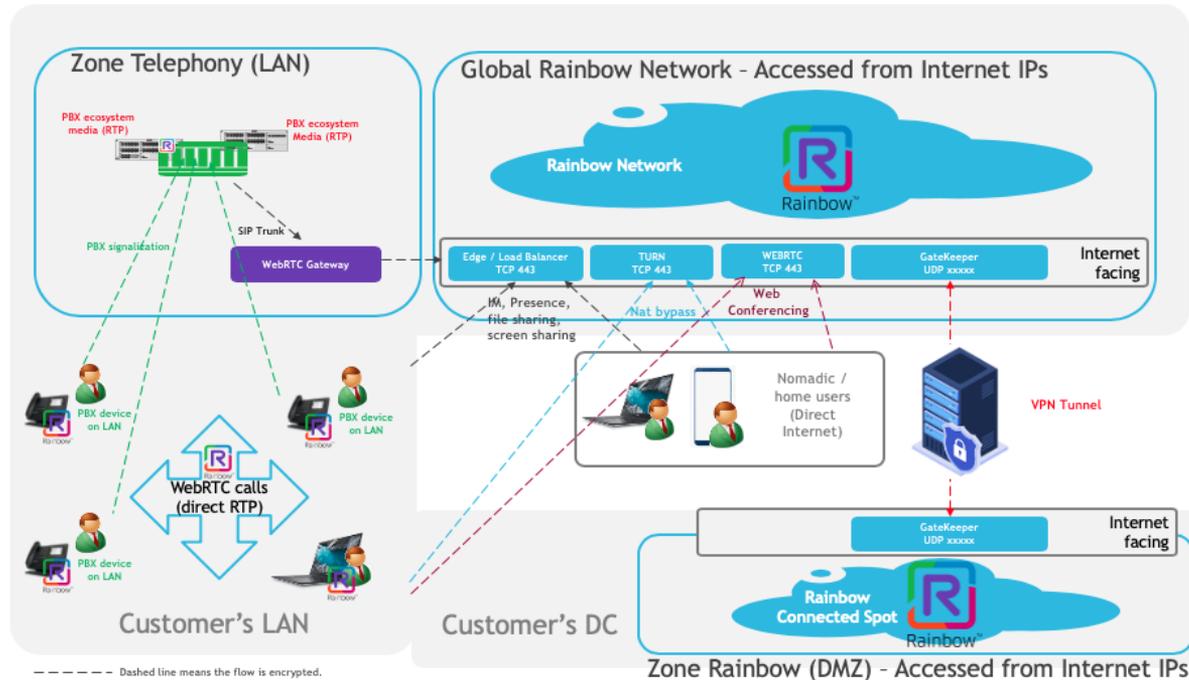
8.2 Network Requirements

Following the aforementioned architecture, users' connections will remain on Rainbow Internet-facing public servers (public load balancers, TURN media relays and WebRTC conferencing servers). Users connections will not be made against the Edge servers, hosting customers data.

As a consequence, the Edge will only feature a server-to-server (or Cloud-to-Cloud) connection between the GateKeepers Edge Gateways installed in customer's premises and Rainbow Cloud GateKeepers. As mentioned in chapter 6, this is done through a secure VPN tunnel using WireGuard technology.

The only requirement is then to have a public IPv4 address on each GateKeeper instance, with a UDP port (to be chosen) open for remote inter-GK tunneling. Data hosted with customer's Edge will be handled by global Rainbow Cloud application servers. As a consequence, all data generated or used by customer's users will transit (in a fully encrypted way) within the GK tunnel. Bandwidth may come as an issue. It is recommended to have between 100 Mbps and 1 Gbps of Internet bandwidth, depending on the expected volume of users to be addressed. A minimum of 50 Mbps is required by ALE to ensure a proper behavior of the service.

8.3 Hardware Footprint

All servers stated in section 7.4 (Generic Hardware Footprint) are required.

On top of that list are the following requirements:

Name	Count #VMs	HA Mode	CPU vCPUs	Memory Size (GB)	Disk (per server) Size (GB)	Disk (per server) IOPS	Network (per server) NIC Type	Network (per server) IPv4
XMPP Server	2	M / S	4	16	64	Standard	LAN	1
MySQL DB	2	M / S	4	64	64	Standard	LAN	1
					256	Optimized		
					256	Optimized		
Mongo DB	3	M / S / S	4	64	64	Standard	LAN	1
					256	Optimized		
					256	Optimized		
Private Load Balancer	2	M / M	4	4	64	Standard	LAN	1 + 2 VIP
SWIFT Object Storage	3	M / M / M	6	16	64	Standard	LAN	1
					1024	Standard		
					1024	Standard		
					1024	Standard		

N.B:

- M stands for Master (or Primary)
- S stands for Slave (or Secondary)

These are the minimum footprint requirements that can accommodate up to 20k concurrent users (100k registered users). It is important to note that these requirements are likely to scale up, should time go on (and so will the usage), or user base grow.

Going above this footprint is possible and recommended as to increase performances or "be ready for the future".

Should the usage scale up, so will the required footprint. Rainbow services have been designed to scale up both horizontally and vertically, making things easier for upgrades. As a rule of thumb, it is recommended to double up the footprint for every 20k concurrent (100-200k registers) users. The most adequate scale up strategy (horizontal / vertical) is to be discussed between customer and ALE Cloud Operations.

8.4 Hardware Footprint for Rainbow Voice

Adding "Rainbow Voice" feature is a possible scenario. While core "Rainbow Voice" services remain operated at ALE's main datacenters, Edge datacenters allows for local voice breakout, where both SIP Trunks and phones are directly connected to.

All servers stated in previous section (Connected Edge Hardware Footprint) are required.

On top of that list are the following requirements:

Name	Count	HA	CPU	Memory	Disk (per server)		Network (per server)	
	#VMs	Mode	vCPUs	Size (GB)	Size (GB)	IOPS	NIC Type	IPv4
In-Country Nodes (ICNs)	2	M / S	16	64	128	Standard	LAN	1 + VIP
							WAN	1 + VIP
Voice Cloud PBX Orchestrator	1	M	2	6	128	Standard	LAN	1
Voice Cloud PBX Gateway	5	M / M	2	6	128	Standard	LAN	1

N.B:

- M stands for Master (or Primary)
- S stands for Slave (or Secondary)
- VIP stands for "Virtual IP", a floating IP shared through VRPP.

These are the minimum footprint requirements that can accommodate up to 50k subscribed users and 2000 concurrent calls. It is important to note that these requirements are likely to scale up, should time go on (and so will the usage), or user base grow.

Going above this footprint is possible and recommended as to increase performances or "be ready for the future".

Should the usage scale up, so will the required footprint. As a rule of thumb, this topology can be horizontally scaled up by adding the same hardware specifications for each couple of 50000/2000 concurrent users/calls.

9 Standalone Edge

9.1 Network Topology

A Rainbow "Standalone Edge" is meant to provide customer with a completely isolated Rainbow service instance, with data storage within customers premises as well as application servers. It can be seen as private Cloud service, dedicated to customer with no possible interaction with the global Rainbow network.

The standalone Edge still behaves like a Cloud service, being accessible from public Internet, allowing user connection from both enterprise LAN and public Internet. Should the service be bound to enterprise's LAN only, it will require the various nomadic/home users (desktop or mobile) to always connect to enterprise's corporate network first, through a dedicated VPN. This is customer's responsibility and out of the scope of this document.
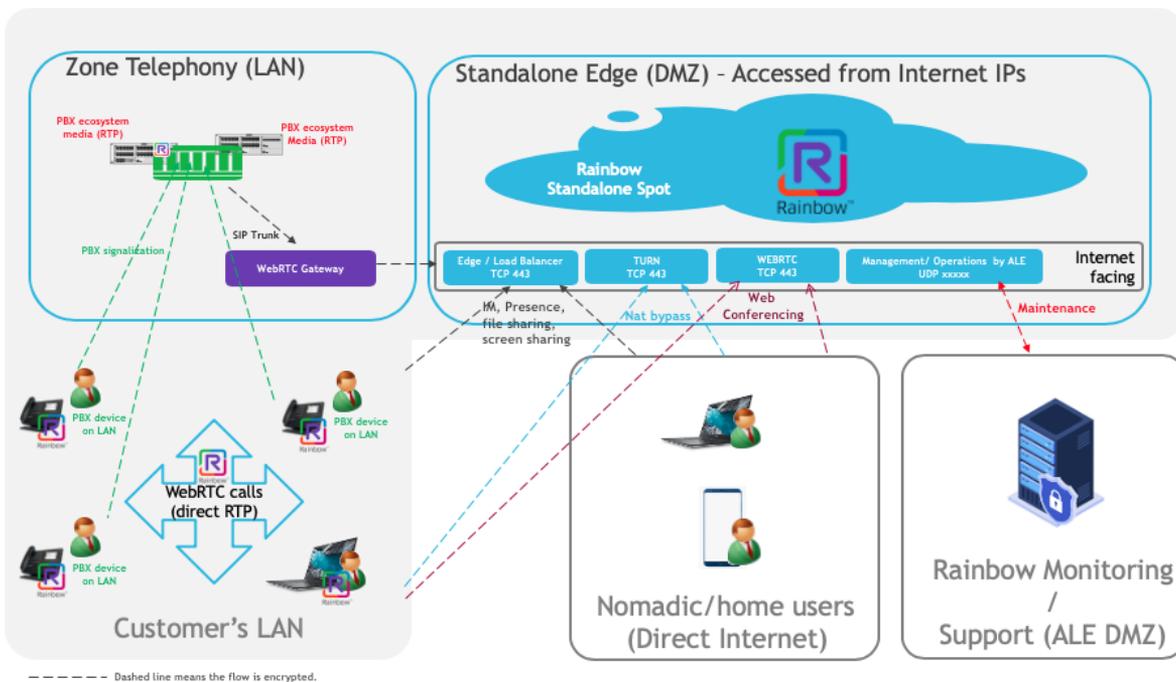
9.2 Network Requirements

Rainbow "Standalone Edge" is still considered as Cloud service, yet being private to customer. This allows nomadic/home users to reach their private Cloud service from virtually anywhere.

The direct consequences of this private Cloud are:

- The Edge will be Internet-facing. Servers parts of the customer's DMZ will then have services exposed to Internet and will potentially be subjects to network threats. ALE Cloud Operations team has developed all the necessary toolkits and will ensure the global security of the infrastructure. The network requirements exposed in the "Rainbow Network Requirements" document then need to be addressed, but from an inbound perspective (server-side instead of client-side).

- The Edge will be made available through its very own DNS and FQDN. Global Rainbow platform is accessible through openrainbow.com (and its associated subdomains). Customer will be requested by ALE to provide a domain for its service. It can be a standalone dedicated domain (e.g. customer-brand-new-service.com) or part of an existing corporate domain (e.g. service.customer-brand.com, where customer-brand.com is the customer's existing domain name). Be aware that ALE will not be held responsible for any legal copyright infringement associated with the selected domain and will be the sole responsibility of the customer. The requirement for ALE Cloud operations team is to be able to administrate the public DNS accordingly. ALE Ops team can either:

  o Provide the customer with the necessary technical information.

  o Buy and manage the domain for the customer.

  o Let the customer buy the domain and do the technical administration part.

- The Edge being publicly available over the Internet with its new unique domain name, it is up to the customer to provide associated SSL certificates to ensure encrypted traffic. Rainbow uses wildcard SSL certificates, signed by recognized Certificate of Authority using SHA256. ALE Ops can either:

  o Provide the customer with the necessary technical information and install the provided certificates back in return.

  o Buy and manage the certificates for the customer (transaction will still require domain owner's approval)

9.3 Hardware Footprint (Small Scale Deployments)

Small scale deployments (500-1000 concurrent users, 5000 registered users) may not require a full blown highly available and scalable infrastructure. Most of Rainbow services can then be mutualized within one single server instance codenamed AiO (All-In-One), where HA is covered through infrastructure's virtualization capabilities.

All servers stated in section 7.4 (Generic Hardware Footprint) are still required though, as to provide remote connectivity and operations of the service.

On top of that list are the following requirements:

Name	Count #VMs	CPU vCPUs	Memory Size (GB)	Disk (per server) Size (GB)	Disk (per server) IOPS	Network (per server) NIC Type	Network (per server) IPv4
All-In-One	1	8	32	OS	Standard	LAN	1
				512	Optimized	WAN	2
				512	Optimized		
WebRTC STUN/TURN Server	1	1	4	64	Standard	LAN	1
						WAN	1
WebRTC Conferencing Server	1	4	8	64	Standard	LAN	1
						WAN	1

				64	Standard		
SWIFT Object Storage	1	4	16	1024	Standard	LAN	1
				1024	Standard		
				1024	Standard		

9.4 Hardware Footprint (Large Scale Deployments)

For large scale deployments (more than 1000 concurrent users, 5000 registered), expected to scale up, a more exhaustive infrastructure is foreseen, as to allow further seamless scalability.

All servers stated in section 7.4 (Generic Hardware Footprint) and 8.3 (Connected Edge Hardware Footprint) are still required. A Standalone Edge footprint is basically a connected one (data only) with all the missing application parts to be fully standalone. On top of that list are the following requirements:

	Count	CPU	Memory	Disk (per server)		Network (per server)	
Name	#VMs	vCPUs	Size (GB)	Size (GB)	IOPS	NIC Type	IPv4
API Servers	2	4	16	64	Standard	LAN	1
WebRTC STUN/TURN Server	2	1	4	64	Standard	LAN	1
						WAN	1
WebRTC Conferencing Server	2	4	8	64	Standard	LAN	1
						WAN	1
Public Load Balancer	2	4	4	64	Standard	LAN	1
						WAN	2
Cloud PBX Gateway	2	4	4	128	Standard	LAN	1

These are the minimum footprint requirements that can accommodate up to 20k concurrent users (100k registered users). It is important to note that these requirements are likely to scale up, should time go on (and so will the usage), or user base grow.

The scalability detailed in section 8.3 is still valid (doubles up approximately every 20k concurrent users).

9.5 Hardware Footprint for Rainbow Voice (Large Scale Deployments)

Added "Rainbow Voice" feature is a possible scenario. Please note however, that, in the case of Standalone Edge, this solution is limited to large-scale deployments (section 9.4) and not small-scale deployments (section 9.3). As for other Rainbow services, Voice services are part of local Edge when it comes to Standalone scenario.

All servers stated in section 9.4 (Standalone Edge Hardware Footprint) 8.4 (Connected Edge Hardware Footprint for Rainbow Voice) are required.

On top of that list are the following requirements:

Name	Count	HA	CPU	Memory	Disk (per server)		Network (per server)	
	#VMs	Mode	vCPUs	Size (GB)	Size (GB)	IOPS	NIC Type	IPv4
Voice Web Server	2	M / S	16	64	128	Standard	LAN	1 + VIP
Voice Proxy	2	M / S	16	64	128	Standard	LAN	1 + VIP
Voice Database	2	M / S	16	64	2048	Optimized	LAN	1 + VIP

N.B:

- M stands for Master (or Primary)
- S stands for Slave (or Secondary)
- VIP stands for "Virtual IP", a floating IP shared through VRPP.

These are the minimum footprint requirements that can accommodate up to:

- 200k subscribed users for Voice Web and Database servers
- 50k subscribed users and 2000 concurrent calls for Voice Proxy servers.

It is important to note that these requirements are likely to scale up, should time go on (and so will the usage), or user base grow. Going above this footprint is possible and recommended as to increase performances or "be ready for the future".

Should the usage scale up, so will the required footprint. As a rule of thumb, this topology can be horizontally scaled up by adding the same hardware specifications for each couple of servers.

9.6 Mail Server

Please note that in case of standalone Edge, Rainbow will require a valid SMTP server to be used to send email notification, passwords recovery and so on. Different possibilities are offered:

- Using company's corporate mailing services if available
- Using external mail service provided by partner.
- Extending the Rainbow standalone infrastructure with a couple of instances, secured and dedicated to send emails using the appropriate Edge's domain (recommended scenario).

End of Document