ALCATEL-LUCENT RAINBOW™

RAINBOW SERVICE LEVEL INDICATORS AND OBJECTIVES

JULY 2021

Contents

Glossary

- **SLI** (Service Level Indicator) is a defined quantitative measure concerning a subset of the global service provided.
- **SLO** (Service Level Objective) is a target value for a service level that is measured by an SLI
- **SLA** (Service Level Agreement) is an explicit contract with end customers or partners on a global SLA which could be a combination of SLO. It may also include some financial penalties associated to actual SLI achievements.
- **Rainbow Region** maps to a geographical area and group of datacenters where the data of the customer's company is located (and usually accessed). Currently 6 regions are available:
  - North America (NA)
  - Caribbean and Latin America (CALA)
  - Europe and Middle East (EMEA – Default one)
  - Germany (DE)
  - Asia-Pacific (APAC)
  - Australia & New Zealand (ANZ)
- **Rainbow instance** defines a comprehensive Rainbow service operated at a given domain name. Some Rainbow instance examples are:
  - .com worldwide public solution
  - .health dedicated HDS instance in France

- **Scheduled Downtime:** Major Rainbow service upgrades are rolled-out on a bi-monthly basis in the most seamless way as possible. However, Rainbow Operations sometimes need to perform maintenance which have a business continuity impact to keep Rainbow working as smoothly as possible. In such case, "Scheduled Downtime" is necessary. The customer is given a 48 hours advance notice on Rainbow Help Center.

1 Scope

This document applies to:

☐ all ALE entities globally

☑ only the following ALE entities: Cloud Business Division

2 Introduction

The aim of this documentation is to present Rainbow Service Level Indicators (SLI) and Service Level Objectives (SLO) as defined by the Rainbow Operations team.

This document will not elaborate about Rainbow Service Level Agreement (SLA).

Unless specified, this document applies to each of the Rainbow instances.

3 Common definitions

- Data to build Indicators is based on server monitoring tools (API errors, sessions with unexpected terminations) and from black-box monitoring tools (healthcheck v3)
- Any core unavailability happening on EMEA region will be also counted as regions unavailability as it will affect hybrid PBX or free, guest and expatriate users.
- Scheduled maintenances down times will not be counted as unavailability.
- Unavailability lasting less than 5 minutes are not part of SLI computations.
- For regional SLIs, a local Rainbow user is fully located in a Rainbow region if all of its data: company, messages, contacts and files shared are on the same region, else the user is considered as a roaming/core user.
- Indicators based on API errors are computed by dividing the total of 5xx requests by total of 2xx, 3xx and 5xx requests. 4xx (client error) are excluded from computation.
- SLI are aggregated on a monthly basis. As our business evolves, so will our SLO. ALE reserves the right to change the terms of the SLO in accordance with the Rainbow Service Description at any time.

4 Service Level Indicators

4.1 [Local] Rainbow core

xx.xx% of local Rainbow users are able to connect to Rainbow, send messages to local contacts.

[Semi-auto] The percentage is the ratio of minutes not able to connect in a month:

To compute this metric, we work with healthcheck failed probes (>5 minutes), API error rate (>10% on a region on enduser, authentication, search and filesharing) and xmpp congestions (>5 minutes on mod_user, mod_sql).

If one of these conditions is met, we count the number of minutes the unavailability lasts and multiply it by the portion of potentially impacted users (eg: 25% if congestion on 1 server out of 4).

Thus, with current SLI computation, the unavailability cannot be targeted on a population like a company, a country or an application.

4.2 [Local] Probes core

xx.xx% of local Rainbow healthcheck users are able to connect to Rainbow and communicate.

[Auto] The ratio of successful scenarios per month:

Healthcheck is an automated agent, spread worldwide, that runs several end-user checks on rainbow, like login, send an IM or check public web site availability.

To compute this metric, we work with healthcheck probes:

- *Count only scenarios about: login, logout, roster, search user, chat, bubble join/leave/chat, file transfer*
- *Scenario execution time is not a cause for success/failure, except for timeouts.*

Divide the count of monthly failed scenarios on each local healthcheck instance by total of monthly scenarios

4.3 [Local] Conferencing

xx.xx% of local Rainbow users are able to interact in a Rainbow bubble, start or join a PSTN or webRTC conference.

[Semi-auto] The percentage is the ratio of successful monthly conferences usages:

To compute this metric, we work with healthcheck failed probes (>5 minutes), API error rate (>10% on conference portal local), Janus crashes, PGi unavailability and xmpp congestions (>5 minutes on mod_muc, mod_sql), as well as channels outages.

If one of these conditions is met, we count:

- *The number of failures to start and join a conference due to API errors (PSTN & WebRTC)*
- *The number of failures due to janus crash or restart. Count only audio sessions.*
- *In case of Rainbow core unavailability or xmpp mod_muc congestion the estimated number of missed webrtc conferences made during the same time during the same day on previous week.*

Finally divide the missed+failed conferences participations by the total (start+join+failed+missed).

4.4 [Global] Hybrid PBX Telephony

Hybrid PBX users are able to control their company phone xx.xx% of the time.

[Semi-auto] The percentage is the ratio of minutes when majority of users lost telephony services in the month:

To compute this metric, we work with PBX phones sessions on xmpp EMEA.

Following situations trigger a significant unavailability:

- *A network or xmpp failure occurs causing a significant drop in connected phone or pbx sessions in EMEA xmpp and possible congestions on xmpp or uncomplete connection on PCGs.*
- *A PCG host or public load balancer failure, disconnecting a large portion of PBXs.*

Each time an unavailability occurs, we count the number of minutes until 98% xmpp phone sessions are established on xmpp. The reference is the number of sessions just before the outage. This 2% margin is required to ignore PBX under maintenance and outliers.

Following situations do not participate to SLI measurement:

- *WebRTC gateway connections*
- *One PCG instance crash or restart as the PBX will automatically reconnect on another PCG and users recover their telephony services in less than 5 minutes. Also, no incidence on established calls.*

4.5 [Local] Media relay (turn)

xx.xx% of local Rainbow users are able to use a Rainbow media relay (turn) during Rainbow webRTC calls.

[Semi-auto] The percentage is the ratio of successful monthly turn usages:

To compute this metric, we work with turn crashes and unavailability

Following unavailability are counted:

- *The number of failures due to turn crash, restart or public network failure.*
- *In case of 2 turn servers failing in the same region count estimated number of missed webrtc calls made during the same time during the same day on previous week.*

Finally divide the missed+failed sessions by the total (start+join+failed+missed).

4.6 [Global] Rainbow Hub

phones SIP calls are possible %xx.xx of the time whatever the Hub service is used

[Semi-auto] The percentage is the ratio of SIP engine uptime monthly

This metric is the mean value of the following subdivided metrics:

SLI-V = (SLI-VT + SLI-VP + SLI-VD + SLI-VS) / 4

Service	SLI	Unit
Rainbow Hub SIP Trunks	SLI-VT	% of availability during a month
Rainbow Hub 3PCC	SLI-VP	% of availability during a month
Rainbow Hub SIP devices	SLI-VD	% of availability during a month
Rainbow Hub Softphony	SLI-VS	% of availability during a month

4.6.1 Rainbow Hub SIP Trunks

SIP Trunk accesses are available xx.xx % of the time per Rainbow Hub partner.

[Semi-auto] The percentage is the ratio of SIP trunk uptime monthly:

To compute this metric, we work with SIP trunk crashes and unavailability (metric ngcp_peer_host_status).

This metric is global and not per unique SIP trunk, and is a mean of all the trunks availability metrics

Following unavailability are counted:

- *The time of unavailability due to SIP trunk crash, restart or public network failure are discounted from operational time*

Finally divide the operational time from the monthly period.

Depending on the SIP trunk model, (Over The Top, Cloud Cross connection, or local edge partner) end to end SIP trunk availability may be affected by final SIP provider, or by intermediate networks that may impact the service. Work is going on to discriminate such condition and do not affect ALE SLO/SLI).

4.6.2 Rainbow Hub SIP devices

SIP phones can pass calls to SIP trunk xx.xx % of the time.

[Semi-auto] The percentage is the ratio of Kamailio-lb and kamailio-proxy availability monthly:

To compute this metric, we work with the availability of 2 centrals services (Kamailio-LoadBalancer and Kamailio-Proxy).

Following situations trigger a significant unavailability, we count the period:

For which any one of these services is down (service crashed or stopped) for more than 30s

When the API telephony rate is degraded (ratio of 5xx or 4xx relatively to 200 OK is higher than the acceptable value impacted clients)

When a significant drop of registered subscribers (physical devices only) is observed (2% loss). The reference is the number of registered subscribers in the previous minutes just before the outage. This 2% margin is required to ignore devices impacted by access network and local disturbances (local phone disconnect…)

Finally divide the availability time by the month duration

4.6.3 Rainbow Hub Softphony

Rainbow users can pass SIP trunk calls xx.xx % of the time.

[Semi-auto] The percentage is the ratio of monthly cloud-wrg uptime:

To compute this metric, we discount the time when the Cloud-wrg component is not registered on the kamailio load balancer.

Finally divide the availability time by the month duration

4.6.4 Rainbow Hub 3PCC

3rd party calls for end user SIP phones are possible xx.xx % of the time.

[Semi-auto] The percentage is the ratio of monthly rvcp-pcg uptime connection:

To compute this metric, we discount the time when:

A drop of NGCP subscribers successfully connected is observed, corresponding to the time to retrieve 90% of the number of subscribers connected before the drop

The ratio of rvcppcg xmpp connection failure is over 5%

The CSTA websocket connection is failed (1 ngcp CSTA cnx)

Finally divide the availability time by the month duration

4.7 [Global] Rainbow admin and subscriptions

xx.xx% of monthly successful Rainbow admin and subscription APIs calls.

[Auto] The percentage is the ratio of monthly successful API calls to admin and subscription portals:

To compute this metric, we work with following API metrics. Compute the error rate for following API routes:

- *API calls to admin, massprovisionning, subscription and invoicing portals*
- *All GET, POST, PUT and DELETE API calls*
- *All POST, PUT and DELETE on /api/rainbow/subscription/v1.0/xxx*

  This includes quotations and subscriptions update, but not the billing info read.

4.8 Rainbow compounded Regional SLI

xx.xx% of the time during a month
where Rainbow service are considered up and running.

The previous global and local SLI are compounded into Regional SLI according with:

SLI-Reg = (SLI-H + SLI-V + SLI-A + SLI-C + SLI-P + SLI-B + SLI-M) / 7

Service	SLI	Unit	Comments
Hybrid PBX telephony	SLI-H	% of availability during a month	Measured WW
HUB Services	SLI-V	% of availability during a month	
Rainbow admin and subscriptions	SLI-A	% of successful API calls	
Rainbow core	SLI-C	% of availability during a month	Measured for each Rainbow region
Rainbow probe	SLI-P	% of availability during a month	
Conferencing	SLI-B	% of successful conferences (SFU or PSTN)	
Media relay	SLI-M	% of successful usage of TURN	
Compounded Regional SLI	SLI-Reg	% of availability during a month	Measured for each Rainbow region

5 Service Level Objectives

Service	SLO	Unit	Comments
Hybrid PBX telephony	**99,75%**	% of availability during a month	Measured WW
Rainbow Hub SIP Trunks	**99,90%**	% of SIP trunks availability	
Rainbow Hub SIP Phones	**99,90%**	% of SIP phones availability	
Rainbow Hub softphone	**99,90%**	% of available min/month	
Rainbow Hub 3PCC	**99,90%**	% of available min/month	
Rainbow admin and subscriptions	**99,95%**	% of successful API calls	
Rainbow core	**99,95%**	% of availability during a month	Measured for each Rainbow region
Rainbow probes	**99,90%**	% of valid IM scenario /month	
Conferencing	**99,90%**	% of successful conferences (SFU or PSTN)	
Media relay	99.95%	% of successful usage of TURN	
Compounded Regional SLI	99.9%	% of availability during a month	Measured for each Rainbow region

- End of Document -