Business

Technical appendix to the Managed Applications Service Description – Google Native Managed Service

Table des matières

1 Overview of the Service

1.1 Overall description

The document is an appendix to the Managed Applications Service Description. It provides service description and further details for the:

- Managed Business Application provided on Google Cloud Platform
- Managed Cloud Native Services provided on Google Cloud Platform

The description adds to the managed services already described in the other document called Managed Applications Service Description:

- MANAGED OS FOR CLOUD SERVERS
- MANAGED DATABASE
- MANAGED MIDDLEWARE
- MANAGED KUBERNETES
- MANAGED CONTAINER
- MANAGED APPLICATION
- MANAGED SAP
- CAASCAD SERVICE
- LOG AS A SERVICE (LAAS)

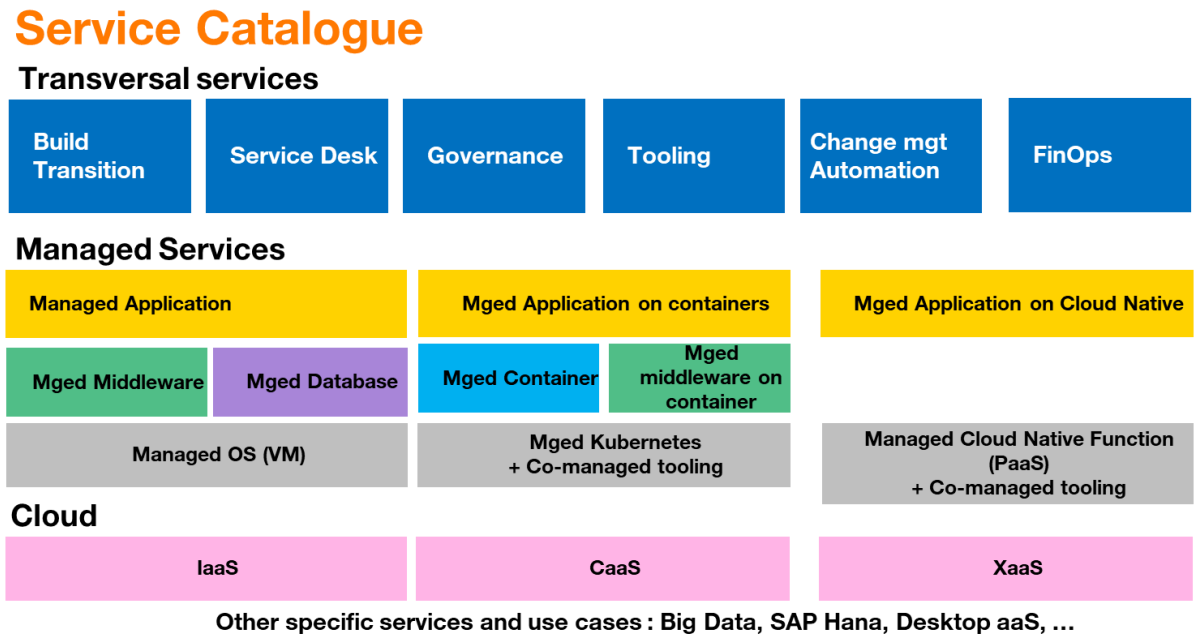

2 Managed Cloud Native Services on GCP

Customer's business application deployed on GCP are dependent on GCP Cloud Native Services (IaaS, PaaS). Orange Business Services provides the managed services necessary to ensure service assurance and change management for those dependences, as well as the configuration and deployment for building and recovering them.

2.1 Description

Orange Business Services provides technical operation and monitoring of the Customer's native GCP, as well as optimization/upgrade activities through the implementation of a network interconnection between the Orange Business Services "service area" and the third party cloud provider's IaaS platform.

The following table lists the summarized services provided as part of the "Cloud Native Services":

Table 1: Description "Cloud Native Services"

Phase	Activities
Cloud Native Services Implementation Phase	▪ Review and validation of the RACI of the GCP Customer application services by Orange Business Services ▪ Review and adjustment of the reflex sheets (MOP on incident) provided by the Customer's business to OBS (When applicable in transition where customer environment exists and in case of managed applications) ▪ Takeover and/or elaboration of the documentation for the use of Orange Business Services teams ▪ Co-definition and/or review of application alarms and thresholds ▪ Creation of accesses for Orange Business Services administrators ▪ Configuration of the VPN operation (when required) ▪ Configuration and testing of alarms in the centralized monitoring system OBS ▪ Customer training on the Cloud Store for change/incident request access
Cloud Native Services Operation Phase	▪ Supervision and Operation ○ Reading and analysis of alarms ○ Correction of faulty configurations ○ Joint review then update of security group and access control ○ Event management (changes & incidents) and interfacing with GCP support if needed and application operation ○ Supervision of the service 24/7

2.2 The cloud native services

One can typically distinguish 3 categories of services:

- The user plane services: if a business application depends on it, the business application is likely to be affected by a defect of it. The service does not have persistent data, therefore the recovery does not necessitate data restore.
- The data services: if a business application depends on a data service, the business application is likely to be affected by a defect of it. The service has persistent data, therefore a recovery may necessitate data restore. Data loss, data corruption may affect the business application as well.
- The other services: the business application does not depend on them. Most of those services are used for automation, observation, migration. The loss of the service is not likely to affect the business application. Some of the services are used for managing the user plane and data plane services of the business application, some others have specific usage for which a scope of work shall be established would the customer requires OBS to leverage them as part of the managed service provided.

User plane services	Data services	Other services
Compute □ Compute Engine □ App Engine **Networking** □ Cloud CDN □ Cloud Interconnect □ Cloud DNS □ Cloud Load Balancing □ Cloud VPN □ Cloud NAT □ Virtual Private Cloud □ Storage Transfer Service	**Storage** □ Persistent Disk □ Cloud Storage **Databases** □ Datastore □ Cloud Bigtable □ Cloud SQL **Data Analysis** □ BigQuery □ Pub/Sub □ Pub/Sub Lite □ Dataproc □ Dataflow □ Cloud Composer	**Containers** □ Google Kubernetes Engine **Operation** □ Cloud Audit Logs □ Cloud Logging □ Cloud Monitoring

GCP Cloud Native services by category

2.3 Tasks involved Cloud Native service management

The tasks involved for the management of a cloud native service depends on the service. They consist in:

- Configuring and deploying the service: Infrastructure as Code is leveraged in order to configure the service, the observability, the backup. Level 3 expertise on the service is leveraged for proper implementation thanks to the scope of work (refer to detailed description of build and SRE services)
- Applying the security group and access control policy defined by the customer.
- Service recovery thanks to Infrastructure as Code: in case of failure, most of the services requires to be recovered thanks to a redeployment. Re-configuring the service manually from scratch is not an efficient option: it takes time and is error prone. This is why recovery / redeployment from Infrastructure as Code is preferred.
- Supervision and remedial consists in watching for alarms raised on the service during the monitoring range (typically: 8x5 or 24x7). When an alarm occurs, an incident ticket is raised, a priority is assigned, the customer is notified. Then remedial action is taken thanks to the procedures made available to Level 2 / 1 by the Level 3. The remedial on a cloud native service may be necessary to restore the service of the business application. Would the procedure not remedy to the incident, then the incident is escaladed to the Level 3. Would the root cause be the CSP itself, then the incident is raised to the CSP by the Level 3.
- Backup and restore: depending on the service (if the service has persistence), it is necessary to backup the service data. The management service consists in configuring the backup solution and monitoring the proper run of it. Note: the backup solution has to be subscribed separately e.g. GCBDR. Restoring the service on incident may involve restoring the data from a backup.

- OS patching and anti-virus: keeping OS up to date and virus free is a managed service for Managed Virtual Machine / Managed OS. Please refer to the detailed description.
- Specifics: some cloud native services may have specific configuration or management tasks.
- Business application specifics: by default, standard alerts are watched. The configuration of alerts, logs on a cloud native service which are specific to a business application is subject to a specific scope of work.

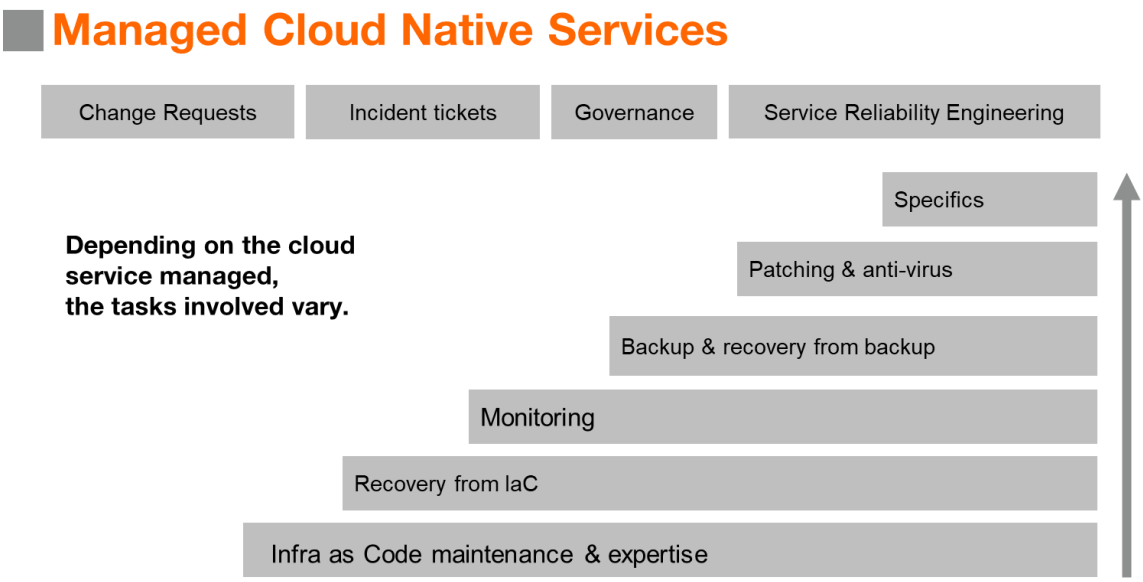

Tasks involved in managed services for cloud native service

Depending on the cloud native service managed, more or less management tasks are necessary and included in the managed service. This drives the complexity of the managed service.

The tasks involved typically depends on the category of the cloud native service, whether user plane, data plane on which the business application depends, or other services upon which the business application does not depend.

	Charging model	User plane services	Data plane services	Other services
Purpose		Used to support customer application	Used to support customer application	Used to operate user plane or data plane
Build	One-time charge based on SoW	IaC in Git, pushed via CI / CD	IaC in Git, pushed via CI / CD	IaC in Git, pushed via CI / CD
Maintaining IaC without changes	Monthly recurring charge	Yes	Yes	Yes
Monitoring & alerts	Monthly recurring charge	Yes	Yes	
Configuration restore on incident	Included in MRC	Yes, from IaC or export	Yes, from IaC or backup	Yes, from IaC when applicable
Data backup and restore on incident	Included in MRC		Yes	
Network and Security Management	Based on SoW	Optional: Based on SoW	Optional: Based on SoW	

Service Desk	Per incident ticket or percentage	Yes	Yes	Yes
Change Management	Per change, in Tokens vs complexity	Via IaC in Git, pushed via CI / CD.	Via IaC in Git, pushed via CI / CD.	Via IaC in Git, pushed via CI / CD
Disaster recovery	Specific design and quote	Optional: Based on SoW	Optional: Based on SoW	

2.4 Table of tasks involved in the management a Cloud Native service

GCP service	Type	Configuration	Monitoring and alerts configured in Google Cloud Monitoring	Backup configured in GCBDR	Recovery procedure	Patch management	Antivirus management	Specificities
Pre-requisite in case of		**Class 2, Class 4 when no GCBDR available for the service**	**Class 2**	**Class 2**	**Class 2 If different from a restore then Class 4, Class 5**	**Class 2**	**n/a**	
Cloud Load Balancing	Managed	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	From IaC	n/a	n/a	
Cloud DNS	Managed	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	From IaC	n/a	n/a	Private DNS should be included in the managed tenant
Content Delivery Network (CDN)	Managed	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring		From IaC	n/a	n/a	HA by design
Cloud NAT	Managed	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring		From IaC	n/a	n/a	Natively Redundant

Cloud Router	Managed	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	From IaC	n/a	n/a	Complex Recovery From Iac or by Operation Team actions
Cloud VPN	Managed	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	From IaC	n/a	n/a	HA in option by design
Cloud SQL	Managed	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring	GCBDR	From IaC	n/a	n/a	
Cloud Storage	Managed	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	From IaC	n/a	n/a	
Storage Transfer Service	Managed	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	From IaC	n/a	n/a	
Google Kubernetes Engine (Std)	Managed	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	From IaC	n/a	n/a	Patch management is included in the service
Google Kubernetes Engine (Autopilot)	Managed	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	From IaC	n/a	n/a	Patch management is included in the service
Compute Engine	Managed	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring	GCBDR	From Backup	Google VM Manager	OBS Sophos	Only supported OS versions
Virtual Private Cloud	Change mgt	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	From IaC	n/a	n/a	
Persistent Disk	Change mgt	Terraform or Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	From IaC	n/a	n/a	

Table of tasks involved in the management of cloud services (extract of services)

2.5 Tooling used for cloud native managed services

GCP tooling and OBS backend operations tooling are leveraged to deliver the managed services. Would the customer require the use of a different tooling, the feasibility shall be confirmed with OBS and the RACI and work-units may be revised.

Process	Tools used by OBS "Managed Applications" delivery teams
Configuration of the infrastructure	Terraform script or Google Cloud Deployment Manager40 GIT referential CI / CD
Supervision solution	Cloud Monitoring with connector to OBS supervision
Backup	GCBDR (incl snapshots)
OS patching solution	VM Manager for OS Patch management OBS MA patching tool (BRAC) OBS OS factory
Antivirus solution	OBS MA Sophos tool
Logging solution	Google Cloud Logging
Recovery	From backup when it exists From Terraform script in GIT if required Ideally from up-to-date Infra as code with CI/CD
Admin connectivity	VPN to OBS Bastion Zone – IaaS connection through CyberArk ou ORUN / Other : Internet through API
Portal for access to MA contract, incident & change ITSM	Portail Cloudstore (ECCS)

2.6 General pre-requisites to the run of managed services

The following pre-requisites are necessary to all managed services:

- The Customer shall have defined a valid architecture. (OBS can optionally provide Professional Services for architecture definition).
- The Customer shall have **a valid subscription to GCP including subscription to GCP Support plan and procure the GCP resources and GCP support plan. OBS can optionally supply this subscription inclusive of GCP support (ref to Multi-Cloud Ready offer for GCP), however, the subscription, the IaaS resources, the GCP support are not part of the Managed Services.** The Managed Services will leverage this support contract to escalate incident to GCP CSP.
- GCP platform for the Customer shall be urbanized alongside best practices of GCP's landing zone or shall offer comparable services.
- OBS proposes a default RACI depending on the class of transition and the resource managed. As a pre-requisite to the project, OBS and the Customer shall agree on the RACI.
- Agreement on the tooling used for GIT, CI / CD chain, Monitoring, Logging and Alerting solution.
- Additional pre-requisites are required when transition is not the entire responsibility of OBS (not Full Build, ref to Build chapter of the document)

In the case of Fully Managed service, OBS is using its own Git, CI / CD chain, Monitoring, Logging and Alerting solution.

In the case of a Co-managed service, OBS and the Customer agree on the Git, CI / CD chain, Monitoring, Logging and Alerting solution to be used. By default, the tooling is

- Either based on GCP tools i.e. Google Cloud Deployment Manager, Google Cloud Monitoring
- Or based on generic multi-cloud tooling proposed by OBS e.g. CaasCad (Prometheus, Grafana, …)

This tooling not included in the Managed Applications work units and can be purchased separately as part of GCP Subscription or as a multi-cloud tooling proposal made by OBS.

2.7 Criteria for the run of a managed cloud native service component

Criteria shall be met with an approval by Level 2 before turning a cloud native component to an active managed service (i.e. Run) by the Level 2 / Level 1 operations. The owner of the Build and of the Level 3 support owns the responsibility of making sure that the criteria are met:

- The architecture and deployment of the service shall be defined.
- The service shall be deployed thanks to Infrastructure-as-Code and tested prior to transitioning to the run team. Typically, successful testing in pre-production, with a pre-production environment iso-production. Note: IaC is necessary to recover the services in case of major failure.
- The use of the service shall be explained to the operation team
- The security policies and access control shall have been configured.
- The access shall have been configured allowing OBS Level 2 teams access.
- The service shall export the necessary metrics towards Google Cloud Monitoring.
- The data backup shall be configured in GCBDR when backup is applicable.
- The disaster recovery shall be configured when applicable.
- The troubleshooting and service restoration procedures shall be provided to Level 2.
- Whereas a procedure requires logs or dashboard those shall have been developed and deployed prior to transferring to run phase.
- A remedial procedure on incident shall not last more than 15 minutes. Beyond, that time amount, the effort would be charged on time base.

3 The build of services & managed services on GCP

3.1 Criteria for qualifying as “backend build” model SoW for a resource:

The “backend build” scope of work model for a resource is used for:

- a resource/service in scope for managed service for which the infrastructure is already built and deployed by the customer leveraging Infrastructure-as-Code.
- And, for which GCP tooling is fully configured and operational prior to transition under customer’s responsibility. The tooling used shall be:
  - Cloud Monitoring & Cloud logging for supervision with proper alerts defined
  - GCBDR properly configured and functional
  - VM Manager & Patch Management configured for VM patching
  - Remedial and troubleshooting procedures on known incident are defined and provided
  - Recovery procedures to be used are defined and provided by the customer
- And, customer provides documentation i.e. schema, HLD and DAT/LLD, architecture explaining how availability & HA, monitoring, security policies and access control, backup, disaster recovery, baseline security, SLA are achieved.

The build effort provided by OBS in the “backend build” includes integrating the alarms from GCP Monitoring to the OBS backend systems, capturing the procedural guides provided by the customer into the OBS knowledge repository of operations, and operations readiness. It includes

as well getting the administrative backend, the OBS ITSM, the portal and billing readiness for operations.

3.2 Criteria for qualifying as “operations build” model SoW for a resource:

The “operations build” scope of work model for a resource is used for:

- a resource/service in scope for managed service for which the infrastructure is already built and deployed by the customer leveraging Infrastructure-as-Code.
- And, customer provides documentation i.e. schema, HLD and DAT/LLD, architecture explaining how availability & HA, monitoring, backup, disaster recovery, baseline security, SLA are achieved.
- And, agreement reached between the customer and OBS to use the GCP and OBS backend tooling.

The build effort provided by OBS in the “operations build” includes that of the “backend build” plus the configuration and deployment of GCP tooling thanks to Infrastructure as Code and of OBS backend i.e.:

- Cloud Monitoring & Cloud logging for supervision with alerts
- GCBDR configuration and deployment
- Update Manager configuration for VM patching
- Anti-virus configuration for VM
- Use of standard remedial and troubleshooting procedures on known incident for the cloud native service.
- Use of standard recovery procedures for the cloud native service.

For further details on the operations per service, please refer to **Chapter 9: detailed description per cloud service.**

3.3 Criteria for qualifying as “full build” model SoW for a resource:

The “full build” scope of work model for a resource is used for:

- a resource/service in scope for managed service not yet built and deployed.
- And, customer provides documentation i.e. schema, HLD and DAT/LLD, architecture explaining how availability & HA, monitoring, backup, disaster recovery, baseline security, SLA are achieved.
- And, agreement reached between the customer and OBS to use the GCP and OBS backend tooling.

The build effort provided by OBS in the “full build” includes that of the “backend build” plus that of the “operational build” plus

- The configuration of the Landing Zone and the infrastructure of the resource leveraging Infrastructure as Code.

For further details on the operations per service, please refer to **Chapter 9: detailed description per cloud service.**

For further details of Infrastructure as Code for full build model, **please refer to chapter Infrastructure as code methodology.**

3.4 Mitigation in case of pre-requisites or criteria not met:

The assessment may reveal that criteria are not met for qualifying to a given build model. Then 3 options are possible:

- the scope of work shall be revisited with a more appropriate build model. This may affect the duration of the project and efforts.
- the customer may remedy to the missing criteria. This may affect the duration of the project and project management and coordination efforts.
- the customer and OBS may agree to live with some limitations in the management capabilities and responsibilities due to the missing criteria.

Would the project be delayed and would resources effort be overspent by OBS as result of pre-requisites and criteria under customer’s responsibility not being met, then OBS would be entitled to charge the overspent effort based on time and material.

3.5 Charging model for build

Service	Work Unit
Project management	Time and material
Service Implementation Coordination	Time and material
Service Reliability Engineer	Time and material
Technical Architect	Time and material (when necessary for documentation)
Full build model - 1st Resource Unit*	One Time Charge per resource
Full build model - subsequent Resource Unit of same type*	OTC per resource
Operations build model - 1st Resource Unit*	OTC per resource
Operations build model - subsequent Resource Unit same type*	OTC per resource
Backend build model - 1st Resource Unit *	OTC per resource
Backend build model - subsequent Resource Unit same type*	OTC per resource

Resource unit*: please refer to Chapter 9: detailed description per cloud service for the definition of the Resource Unit per cloud native service.

4 Security Guideline

Within the framework of a managed offer by Orange Business Services we consider two cases:

1. The deployment of workloads in the Customer's GCP Organization
2. The deployment of workloads in OBS's GCP organization

We have then 3 points to take into account when deploying GCP Security Native Services:

- How do we manage Identification & Authentication for OBS and for the Customer?
  a. User identity,
  b. Service account identity.
- How do we give access rights on these identities (group identities, user, and service account identities)?
- How do we make sure that we have a periodic audit on the 2 previous points? Performing an access review on a periodical basis, and giving an account to the customer, adds value to the offer..

As part of the NGOT (Next Generation Operation Tools) project, a perennial and automatic identity provider infrastructure is currently developed, in order to ensure the assignment of rights to Managed Services' users.

This IAM infrastructure leverages on Orange Business Services' HR processes to manage OBS users and their rights. Customer identities and rights can be managed as well using OBS' IAM infrastructure. Otherwise, co-managed identities and rights can be federated from OBS' IAM infrastructure, on the one hand, and from the customer's identity provider, on the other hand.

4.1 Resources hierarchy

The clients landing zone includes resource hierarchy in the client's organization, and in the gcp.orange-business.com organization to be able to manage the client separately by OCB and by the client in case of need. The resources in gcp.orange-business.com are also used to automate the client's infrastructure deployment via IaC.

Resource hierarchy will be kept simple and synthetic, giving us a good scoping of controls per environment types, then at last level, the projects themselves. **The 4 Levels in bold** hereunder are the standard design, with respect Orange Groups security standard. In case of more levels, the rules apply to all levels. Entities can have more than 4 levels.

1. Organization
2. **Customer Folder (in case of a full managed) (1st level Folder)**
3. **Landing Zone management folder (2nd level Folder)**
4. **Environment and common folders (3rd level Folder)**
5. **Entity folder (4th level Folder)**
6. Projects (5th level Folder)

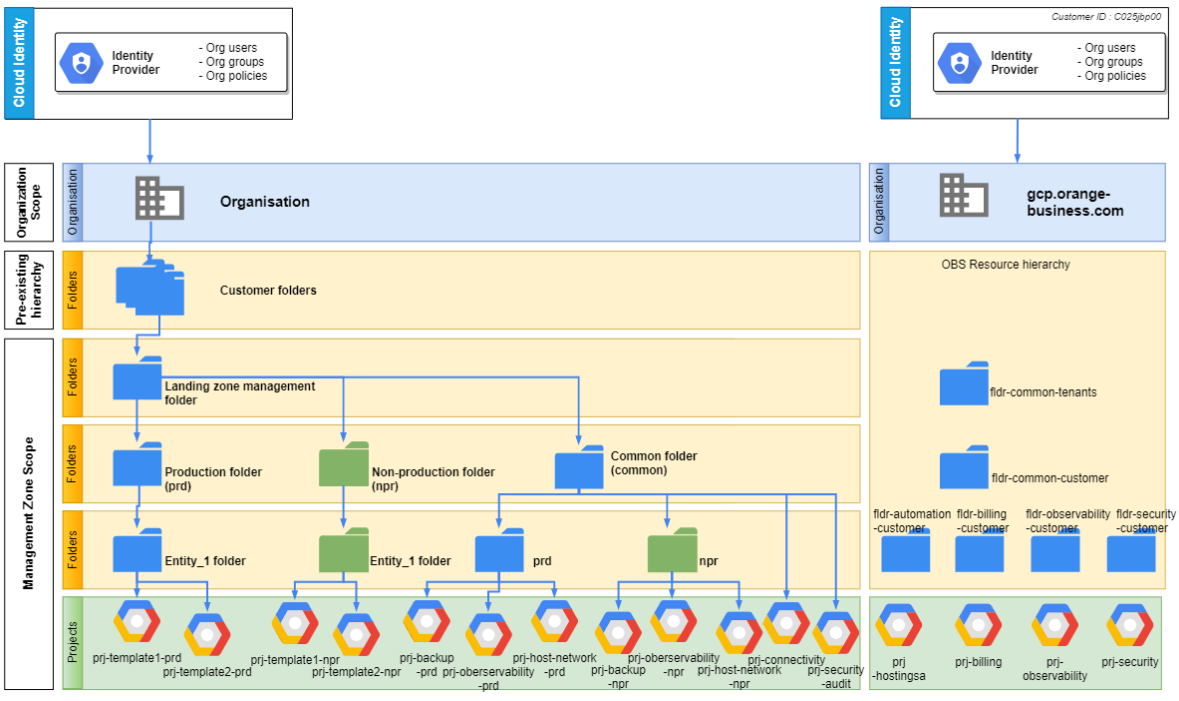

4.2 IAM policy

Identity and access management (IAM) covers **products, processes, and policies used to manage user identities and regulate user access within an organization**.

The implementation of the IAM policy follows the three principles hereunder:

1. Least privilege: ensure that a user has the necessary and sufficient permissions according to their roles
2. Isolation: make sure there is no access or even visibility possible from one identity provider to another
3. Autonomy: limit the loss of time related to the acquisition of permissions

4.2.1 Groups (organization level)

- ✓ Each IAM policy rule is supported by groups defined at the level to which policy is relevant. The group manager of each group is responsible to determine and fill out members.
- ✓ The group name is prefixed depending on the level to which he IAM policy rule belongs.

The following kinds of groups are currently defined:

4 privileged groups

- admins (admins resources): who can create projects and folders
- security-admins: who can give access by giving IAM roles
- billing-account-users: who can assign a billing account to a project
- network administrators:

least privileged groups

- viewers: those who can see resources

- dev-ops: which can deploy as code. Mainly who can masquerade as dedicated service accounts
- dev-apps: users with limited rights
- security-reviewers: who can review cloud security, mainly see IAM role settings

IAM groups created by OBS IaC to efficiently manage privileges

Privileged groups	Customer groups		OBS groups
yes	grp-gcp-iam-ocb-<tenant-name>-**tenant-resources-admins**@<customer-domain>		grp-gcp-iam-ocb-<tenant-name>-**tenant-resources-admins**@gcp.orange-business.com
yes	grp-gcp-iam-ocb-<tenant-name>-**tenant-security-admins**@<customer-domain>		grp-gcp-iam-ocb-<tenant-name>-**tenant-security-admins**@gcp.orange-business.com
yes	grp-gcp-iam-ocb-<tenant-name>-**tenant-billing-account-users**@<customer-domain>		grp-gcp-iam-ocb-<tenant-name>-**tenant-billing-account-users**@gcp.orange-business.com
yes	grp-gcp-iam-ocb-<tenant-name>-**tenant-dev-ops**@<customer-domain>		grp-gcp-iam-ocb-<tenant-name>-**tenant-dev-ops**@gcp.orange-business.com
no	grp-gcp-iam-ocb-<tenant-name>-**tenant-viewers**@<customer-domain>		grp-gcp-iam-ocb-<tenant-name>-**tenant-viewers**@gcp.orange-business.com
no	grp-gcp-iam-ocb-<tenant-name>-**tenant-dev-apps**@<customer-domain>		grp-gcp-iam-ocb-<tenant-name>-**tenant-dev-apps**@gcp.orange-business.com
no	grp-gcp-iam-ocb-<tenant-name>-**tenant-security-reviewers**@<customer-domain>		grp-gcp-iam-ocb-<tenant-name>-**tenant-security-reviewers**@gcp.orange-business.com

4.2.2 Permissions of groups

Privileged IAM groups and service accounts permissions

resource bound <customer-domain> (customer organisation)	member <ocb> is here the entity reseller	role definition	roles	implementation state
organization	grp-gcp-iam-ocb-<tenant-name>-**tenant-admins**@<customer-domain>	Members are responsible for organizing the structure of the resources used by the organization.	roles/cloudsupport.admin roles/resourcemanager.folderCreator roles/resourcemanager.projectCreator roles/securitycenter.admin	up to date
organization	sa-node-creator-<CUSTOMER>-prd	This service account can create nodes	roles/resourcemanager.folderCreator	up to date

		(folders or projects) in the <CUSTOMER> organisation.	roles/resourcemanager.projectCreator	
organization	grp-gcp-iam-ocb-<tenant-name>-**tenant-security-admins**@<customer-domain>	Members are responsible for establishing and managing security policies for the entire organization, including access management.	roles/compute.viewer roles/container.viewer roles/resourcemanager.organizationAdmin roles/iam.organizationRoleViewer roles/iam.securityReviewer roles/logging.configWriter roles/logging.privateLogViewer roles/orgpolicy.policyAdmin roles/orgpolicy.policyViewer roles/resourcemanager.folderIamAdmin roles/resourcemanager.tagAdmin roles/resourcemanager.tagUser roles/securitycenter.admin	up to date
organization	grp-gcp-iam-ocb-<tenant-name>-**tenant-billing-account-users**@<customer-domain>	Members can link the Customer sub billing account to a project in the Customer organisation	roles/billing.projectManager	up to date
<CUSTOMER> sub billing account	grp-gcp-iam-ocb-<tenant-name>-**tenant-billing-account-users**@<customer-domain>	Members can use the Customer sub billing account	roles/billing.user	up to date
<CUSTOMER> sub billing account	sa-node-creator-<tenant-name>-prd	This service account can use the <CUSTOMER> sub billing account	roles/billing.user	up to date
organization	grp-gcp-iam-ocb-<tenant-name>-tenant-viewers@<customer-domain>	Can view the organisation	roles/resourcemanager.organizationViewer	up to date
organization		Can view <TENANT> folder and sub folders	roles/resourcemanager.folderViewer	up to date
<CUSTOMER> sub billing account	grp-gcp-iam-ocb-<tenant-name>-tenant-viewers@<customer-domain>	Members can see all <CUSTOMER> sub billing	roles/billing.viewer	to add

organization	grp-gcp-iam-ocb-<tenant-name>-tenant-dev-apps@<customer-domain>	to be define	to be define	not implemented
organization	grp-gcp-iam-ocb-<tenant-name>-tenant-security-reviewers@<customer-domain>	Members are part of the security team responsible for reviewing cloud security in OCB folder(s)	roles/iam.securityReviewer roles/cloudasset.viewer	up to date

4.3 Management of secrets for Deployment

For reasons of reliability, performance and capitalization, the **deployment is automated** and implemented via Terraform and Ansible.

Automated deployment requires the use of dedicated service accounts, with special privileges. Access to these accounts is protected by specific secrets, the use of which must be restricted to the agents making up the automated deployment tools.

4.3.1 Standard OBS deployment tools

The following tools are used in support to Terraform and Ansible for automation deployment:

Tools	Description	Rules	External reference
GitLab	Code management Execution of CI/CD process (GitLab runners)	Best practices, Gitflow and naming convention	https://sourcehub.orange-business.com/
Hashicorp Vault	Secrets management	Best practices documentation	https://www.hashicorp.com/products/vault
Artifactory	Artifacts storage	Best practices documentation	https://multirepo.orangeapplicationsforbusiness.com/ui/

Required privileged service accounts

resource bound <customer-domain> (customer organisation)	member <ocb> is here the entity reseller	role definition	roles
Organization	sa-tenant-deployer-groups-mgmt@deployment-core-3c64.iam.gserviceaccount.com		Organization Viewer
customer cloud identity	sa-tenant-deployer-groups-mgmt@deployment-core-3c64.iam.gserviceaccount.com	This service account can manage groups in <CUSTOMER> organisation	Admin groups

resource bound <customer-domain> (customer organisation)	member <ocb> is here the entity reseller	role definition	roles
Organization	sa-tenant-deployer-ocb-dev@deployment-ocb-dev-3ea1.iam.gserviceaccount.com	This service account can • modify IAM on the entier organisation • create and delete projects	Organization Administrator Project Creator Project Deleter

4.3.2 Customer-defined deployment tools

According to the service class bought by the customer, the automated deployment scripts can be tailored to customer requirements. These solutions are designed and implemented through dedicated workshops with Orange Business Services.
Design of automated deployment processes and scripts encompasses the topic of the management of the secrets protecting dedicated service accounts.

4.4 Accounts and access review

Accounts and access rights review is an essential component of our IAM policy, closely linked to identity lifecycle management and account and rights provisioning.
It involves ensuring that the access rights of users of the information system are in conformity with what they should be, and certifying them, or - if necessary - carrying out remediation operations in the event of non-compliance with the client's authorization policy.
This activity is therefore part of a governance and control logic for authorizations, in order to provide the expected guarantees of compliance.
The accounts and access review is performed every six months, under the responsibility of the Business Security Officer. If the sold service does not include a BSO service, then the Service Delivery Manager is responsible of the accounts and access review.

4.5 Patch Management

Standard payloads running on traditional VMs are based on hardened OS images. Centralized patch management is performed upon them. When the payload is containerized, OBS can perform the equivalent to patch management, provided that it has responsibility and control on the building of the baseline image.

5 Detailed responsibilities and accountabilities

The following tables describe the standard default responsibilities between OBS and the customer depending **on the build model**.

The following tables describe the standard default responsibilities between OBS and the customer depending on classes of service. Those may be amended with mutuel consent depending on project.

- R stands for responsible
- A stands for Accountable
- C stands for Contributor
- I stands for Informed

5.1 RACI for Managed OS

Service Implementation	OBS	Customer	OBS	Customer	OBS	Customer
OS Server infrastructure implementation	Full build		Operations build		Backend build	
Deployment of the infrastructure	R, A	I	I	R, A	I	R, A
Deployment of LAN components	R, A	I	I	R, A	I	R, A
Deployment of DNS and NTP services	R, A	I	R, A	I	I	R, A
Backup tools for operations (GCBDR & GCP Snapshots)	R, A	I	R, A	I	I	R, A
Deployment of the OS patching solution (Google VM Manager) or OBS Patch management	R, A	I	R, A	I	I	R, A
Deployment of the Antivirus solution	R, A	I	R, A	I	SoW	SoW
Deployment of the supervision solution (Google Cloud Monitoring)	R, A	I	R, A	I	I	R, A
Deployment of the supervision solution (Google Cloud Monitoring)	R, A	I	R, A	I	I	R, A
Deployment of security groups and firewall rules	R, A	I	SoW	SoW	I	R, A
Recovery procedure (Infra as Code, restore, other…)	R, A	I	I	R, A	I	R, A
Testing and validation of infrastructure implementation	R	A	I	R, A	I	R, A
Testing and validation of GCP tooling implementation and lifecycle management	R	A	R	A	I	R, A
OS Server Implementation						
Evaluation or deployment of the operating system	R, A	I	R, A	I	I	R, A
Deployment of new packages	R, A	I	R, A	I	R, A	I
Test and validation of operating system implementation for new packages	R, A	I	R, A	I	R, A	I
Service implementation documentation						
Conception, architecture and low-level design for infrastructure	I	R, A	I	R, A	I	R, A
Implementation and operation documentation for infrastructure	R, A	I	I	R, A	I	R, A
Conception and low-level design for tooling (GCP)	R, A	I	R, A	I	I	R, A
Implementation & operation documentation for tooling (GCP)	R, A	I	R, A	I	I	R, A

5.2 RACI for Database as a Service

Service Implementation	OBS	Customer	OBS	Customer	OBS	Customer
Database aaS services conception and implementation	Full Build		Operations Build		Backend Build	
Maintenance of Infrastructure architecture referential	R, A	I	I	R, A	I	R, A
Maintenance of tooling configuration referential	R, A	I	R, A	I	I	R, A
Deployment of the infrastructure	R, A	I	I	R, A	I	R, A
Deployment of the supervision solution (Google Cloud Monitoring)	R, A	I	R, A	I	I	R, A
Deployment of the logging solution (Google Cloud Logging) (optional)	R, A	I	R, A	I	I	R, A
Deployment of the backup solution (GCBDR, Snapshot)	R, A	C, I	R, A	C, I	I	R, A
Recovery procedure for infrastructure from referential (Infra as code, restore from backup, other…)	R, A	C, I	I	R, A	I	R, A
Recovery procedure for tooling from referential (Infra as code, restore, other…)	R, A	C, I	R, A	C, I	I	R, A
Testing and validation of infrastructure implementation	R, A	I	I	R, A	I	R, A
Testing and validation of tooling implementation and lifecycle management	R, A	I	R, A	C, I	I	R, A
Customer provided script execution on DB instance	R	A, I	R	A, I	R	A, I
OBS script execution on DB instance	R, A	C, I	R, A	C, I	R, A	C, I
Service implementation documentation						
Conception, architecture, and low-level design for infrastructure	C, I	R, A	I	R, A	I	R, A
Implementation and operation documentation for infra	R, A	C, I	I	R, A	I	R, A
Conception and low-level design for tooling (GCP)	R, A	C, I	R, A	C, I	I	R, A
Implementation & operation documentation for tooling (GCP)	R, A	C, I	R, A	C, I	I	R, A

Service Operation	OBS	Customer	OBS	Customer	OBS	Customer
Database aaS services operations	Full Build		Operations Build		Backend Build	
Monitoring through Google Cloud Monitoring	R	I	R	I	R*	I
Investigation through Google Cloud Monitoring & Google Cloud Logging	R, A	C,I	R, A	C,I	R*	A
Restore from Infra as Code and backup	R, A	C,I	R, A	C,I	R*	A
Changing capacity of database instance	R, A	C,I	C, I	R, A	C, I	R, A
ITSM operations						
Change Management	R	A	R	A	R	A
Incident Management	R, A	R**,I	R, A	R**,I	R, A	R**,I
Event management	R, A	I	R, A	I	R, A	I
Baseline security management	R	A	SoW	SoW	SoW	SoW
Configuration management	R, A	C, I	R	A	R	A
Report management via SDM service	R, A	C, I	R, A	C, I	R, A	C, I
Invoicing management	R, A	I	R, A	I	R, A	I

R*: within the limitations of tooling provided by the Customer
R**: in co-management model, customer may have joint responsibilities related to the activity & incident

5.3 RACI for other Native Services managed

Service Implementation	OBS	Customer	OBS	Customer	OBS	Customer
Native service infrastructure implementation	Full Build		Operations Build		Backend Build	
Deployment of the infrastructure	R, A	I	I	R, A	I	R, A
Backup tools for operations (GCBDR) (1)	R, A	I	R, A	I	I	R, A
Deployment of the supervision solution (Google Cloud Monitoring) (1)	R, A	I	R, A	I	I	R, A
Deployment of the logging solution (Google Cloud Logging) optional (1)	R, A	I	R, A	I	I	R, A
Deployment of security groups and firewall rules	R, A	I	SoW	SoW	I	R, A
Recovery procedure (Infra as Code, restore, other…)	R, A	I	I	R, A	I	R, A
Testing and validation of infrastructure implementation	R	A	I	R, A	I	R, A
Testing and validation of GCP tooling implementation	R	A	R	A	I	R, A
Packages						
Deployment of new packages (1)	I	R, A	I	R, A	I	R, A
Service implementation documentation						
Conception, architecture, and low-level design for infrastructure	C, I	R, A	I	R, A	I	R, A
Implementation and operation documentation for infrastructure	R, A	I	I	R, A	I	R, A
Conception and low-level design for tooling (GCP)	R, A	I	R, A	I	I	R, A
Implementation & operation documentation for tooling (GCP)	R, A	I	R, A	I	I	R, A

Service Operation	OBS	Customer	OBS	Customer	OBS	Customer
Native service operations	Full Build		Operations build		Backend build	
Monitoring (1)	R, A	I	R, A	I	R*	A
Backup (1)	R	A	R	A	R*	A
Restore from Infra as Code and backup (1)	R, A	C, I	R, A	C, I	R*	A
Security groups, Firewall rules setting	R	A	SoW	SoW	I	R, A
ITSM operations						
Change Management	R	A	R	A	R*	A
Incident Management	R, A	R**, I	R, A	R**, I	R*, A	R**, I
Event management	R, A	I	R, A	I	R*	A
Baseline security management	R	A	SoW	SoW	SoW	SoW
Report management via SDM service	R, A	I	R, A	I	R, A	I
Invoicing management	R, A	I	R, A	I	R, A	I

R*: within the limitations of tooling provided by the Customer
R**: in co-management model, customer may have joint responsibilities related to the activity & incident
(1) When applicable as per detailed description per service

6 Detailed description per service (Extract)

6.1 Cloud Load Balancing

6.1.1 Description

Google Cloud Load Balancing operates at layer 4 or 7 of the Open Systems Interconnection (OSI) model. Google Cloud Load Balancing is a software-based managed service for distributing traffic in a single or multiple region across multiple instances of applications It's the single point of contact for clients. Load balancer distributes inbound flows that arrive at the load balancer's front end to backend pool instances. These flows are according to configured load-balancing rules and health probes. The backend pool instances can be GCP Virtual Machines or instances in a virtual machine scale set.

A public load balancer can provide outbound connections for virtual machines (VMs) inside your virtual network. These connections are accomplished by translating their private IP addresses to public IP addresses. Public Load Balancers are used to load balance internet traffic to your VMs.

An internal (or private) load balancer is used where private IPs are needed at the frontend only. Internal load balancers are used to load balance traffic inside a virtual network. A load balancer frontend can be accessed from an on-premises network in a hybrid scenario.

6.1.2 Build to run service included in the OTC

6.1.2.1 Build service pre-requisite

- Refer to generic description.

6.1.2.2 Build to run service

- Refer to generic description.

6.1.3 RUN services included in the MRC

6.1.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the load balancer.
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.1.3.2 Co-manage option

Yes, if CI/CD shared with the customer

6.1.3.3 KPI & alerts

Monitoring

Yes, Insights, Metrics, New metric possible with logs, Health probes

KPI monitored

L4/TCP

- l3/external/rtt_latencies >= xms
- l3/internal/rtt_latencies >= xms

L7 / HTTP(s)

- https/backend_latencies >= xms
- https/internal/backend_latencies >= xms

HTTP codes ratio (on demand)

- https/backend_request_count response_code_class = 500 / https/backend_request_count response_code_class = 200 >= x%

- https/backend_request_count response_code_class = 400 / https/backend_request_count response_code_class = 200 >= x%"

Alerts observed

L3/TCP
- l3/external/rtt_latencies >= xms
- l3/internal/rtt_latencies >= xms

L4 / HTTP(s)
- https/backend_latencies >= xms
- https/internal/backend_latencies >= xms

HTTP codes ratio (on demand)
- https/backend_request_count response_code_class = 500 / https/backend_request_count response_code_class = 200 >= x%
- https/backend_request_count response_code_class = 400 / https/backend_request_count response_code_class = 200 >= x%"

6.1.3.4 Backup and restore

Data backup and restore

Not applicable. Load balancer does not store data persistently.

Service restore

The Continuous Deployment chain is used to redeploy the Load Balancer from the configuration file of reference for production environment committed in the Git.

6.1.3.5 GCP SLA High Availability and Disaster Recovery inter-region

GCP ensures High Availability of the Load Balancer with standard SKU.

Maintaining a cross region Disaster Recovery requires specific design and subject to a specific additional charging.

6.1.4 Charging model

Work Unit
Per Load Balancer instance

6.1.5 Changes catalogue – in Tokens, per act

Changes examples	Effort	Impact on MRC
Setup / modify / delete URI	1 token	
Change health probes / Add new backend	2 tokens	
Other changes	Estimation in tokens based on time spent	

6.2 Cloud DNS

6.2.1 Description

Cloud DNS host your Domain Name System (DNS) domains in GCP. Cloud DNS offers both public zones and private managed DNS zones. A public zone is visible to the public internet, while a private zone is visible only from one or more Virtual Private Cloud (VPC) networks that you specify.

6.2.2 Build to run service included in the OTC

6.2.2.1 Build service pre-requisite

- Refer to generic description.

6.2.2.2 Build to run service

- Refer to generic description.

6.2.3 RUN services included in the MRC

Run a managed Cloud DNS service is optional. Depending on Customer's interest, the Customer may request the service. By default, there is no recurring task proposed on Cloud DNS service, but on demand changes and on demand investigations.

6.2.3.1 Run service pre-requisite

- A referential file exists in the Git used by OBS which includes the reference configuration of the DNS.
- This file can be executed with a CI/CD used by OBS and the execution has been tested successfully.

6.2.3.2 Co-manage option

For the Public part, OBS work with the customer for the publics domain naming context.
For the private Part, a RACI must be done.

6.2.3.3 KPI & alerts

Monitoring

Yes, Metrics,

KPI monitored

Number of changes in the DNS database.

Alerts observed

Number of changes in the DNS rules

6.2.3.4 Backup and restore

Data backup and restore

Yes. Backup is proposed based on regular export.

Service restore

The CI/CD chain is used to redeploy the records from a backup zone into the native DNS service or from an export

6.2.3.5 GCP SLA High Availability and Disaster Recovery inter-region

Cloud DNS is a high-performance, resilient, global Domain Name System (DNS) service that publishes your domain names to the global DNS.
In case of public DNS the customer should be responsible for the host mastering (registration)

6.2.4 Charging model

Work Unit
Per resource group

6.2.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create / update/ delete zone (one zone including reverse)	1 token
Create / update/ delete record (up to 10 records)	1 token
Zone delegation*	1 token
Configure Firewall DNS	2 tokens
Other changes	Estimation in tokens based on time spent

Zone Delegation*: Specification should be received as a prerequisite.

6.3 Content Delivery Network (CDN)

6.3.1 Description

Google Cloud CDN is a fast, reliable, and secure content delivery network that ensures the delivery of data without any latency. Google Cloud CDN delivers content peers to peers securely over the cloud. Google Cloud CDN optimizes your static content on its fast and reliable servers for delivering your static assets quickly and efficiently and gives you the option to keep our data public or private. Through Google cloud CDN it allows to load very easily, faster and securely, the website of our organizations in a simple and secure way for customer as well as for us.

6.3.2 Build to run service included in the OTC

6.3.2.1 Build service pre-requisite

- Refer to generic description.

6.3.2.2 Build to run service

- Refer to generic description.

6.3.3 RUN services included in the MRC

Run a managed Cloud DNS service is optional. Mandatory if offer is Managed applications, optional if offer is managed infrastructure.

6.3.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the CDN.
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.3.3.2 Co-manage option

Yes based on RACI determined during pre-sales or build.

6.3.3.3 KPI & alerts

Monitoring

Yes: Metrics and diagnostic logs

KPI monitored

- Byte Hit Ratio
- Request Count
- Response Size
- Total Latency
- Customized ping page per zone

Alerts observed

- Customized ping page per zone
- Latency per zone,
- log analysis on métrics

6.3.3.4 Backup and restore

Data backup and restore

Can be exported from CI/CD Pipeline.

Service restore

The Continuous Deployment chain is used to redeploy the CDN from the configuration file of reference for production environment committed in the Git.

6.3.3.5 GCP SLA High Availability and Disaster Recovery inter-region

Based on design SOW, the service can be built in multiple regions.

6.3.4 Charging model

Work Unit
Per Endpoint

6.3.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Purge CDN	1 Token
Add URL	1 Token
Other changes	Estimation in tokens based on time spent

6.4 Cloud NAT

6.4.1 Description

Cloud NAT is a distributed, software-defined managed service. Cloud NAT configures the Andromeda software that powers your Virtual Private Cloud (VPC) network so that it provides source network address translation (source NAT or SNAT) for VMs without external IP addresses. Cloud NAT also provides destination network address translation (destination NAT or DNAT) for established inbound response packets.

6.4.2 Build to run service included in the OTC

6.4.2.1 Build service pre-requisite

- Refer to generic description.

6.4.2.2 Build to run service

- Refer to generic description.

6.4.3 RUN services included in the MRC

6.4.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the Cloud NAT.
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.4.3.2 Co-manage option

No, OBS manages the Cloud NAT

6.4.3.3 KPI & alerts

Monitoring

Yes, Metrics,

KPI monitored

- nat_allocation_failed = 1
- dropped_sent_packets_count >= x%
- dropped_received_packets_count >= x%

Alerts observed

- nat_allocation_failed = 1
- dropped_sent_packets_count >= x%
- dropped_received_packets_count >= x%

6.4.3.4 Backup and restore

Data backup and restore

Can be exported from CI/CD Pipeline.

Service restore

The Continuous Deployment chain is used to redeploy the CDN from the configuration file of reference for production environment committed in the Git.

6.4.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA by design. Based on design SOW, the service can be built in multiple regions

6.4.4 Charging model

Work Unit
Per Endpoint

6.4.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create / update/ delete (including reverse)	1 token
Configure Firewall NAT	2 tokens
Other changes	Estimation in tokens based on time spent

6.5 Cloud Router

6.5.1 Description

Cloud Router is a fully distributed and managed Google Cloud service that uses the Border Gateway Protocol (BGP) to advertise IP address ranges. It programs custom dynamic routes based on the BGP advertisements that it receives from a peer. Instead of a physical device or appliance, each Cloud Router is implemented by software tasks that act as BGP speakers and responders. A Cloud Router also serves

as the control plane for Cloud NAT. Cloud Router provides BGP services for the following Google Cloud products:

- Dedicated Interconnect
- Partner Interconnect
- HA VPN
- Router appliance

6.5.2 Build to run service included in the OTC

6.5.2.1 Build service pre-requisite

- Refer to generic description.

6.5.2.2 Build to run service

- Refer to generic description.

6.5.3 RUN services included in the MRC

6.5.3.1 Run service pre-requisite

- A referential file exists in the Git used by OBS which includes the reference configuration of the Cloud Router.
- This file can be executed with a CI/CD used by OBS and the execution has been tested successfully.

6.5.3.2 Co-manage option

No, Orange Business Services manages the Cloud Router service.

6.5.3.3 KPI & alerts

Monitoring

Yes, Metrics, Logs, Probes

Cloud Router can be monitorized by using Cloud Monitoring using Alerts and Metrics. Realtime Native reporting from GCP (Cloud Monitoring, Cloud Logging) can be used by OBS and Specific reporting on quote.

KPI monitored

gcp.router.best_received_routes_count	Current number of best routes received by router.
gcp.router.bgp.received_routes_count	Current number of routes received on a bgp session.
gcp.router.bgp.sent_routes_count	Current number of routes sent on a bgp session.
gcp.router.bgp.session_up	Indicator for successful bgp session establishment.
gcp.router.bgp_sessions_down_count	Number of BGP sessions on the router that are down.
gcp.router.bgp_sessions_up_count	Number of BGP sessions on the router that are up.
gcp.router.nat.allocated_ports	The number of ports allocated to all VMs by the NAT gateway
gcp.router.nat.closed_connections_count	The number of connections to the NAT gateway that are closed
gcp.router.nat.dropped_received_packets_count	The number of received packets dropped by the NAT gateway
gcp.router.nat.new_connections_count	The number of new connections to the NAT gateway
gcp.router.nat.open_connections	The number of connections open to the NAT gateway
gcp.router.nat.port_usage (gauge)	The highest port usage among all VMs connected to the NAT gateway
gcp.router.nat.received_bytes_count	The number of bytes received by the NAT gateway

gcp.router.nat.received_packets_count	The number of packets received by the NAT gateway
gcp.router.nat.sent_bytes_count	The number of bytes sent by the NAT gateway
gcp.router.nat.sent_packets_count	The number of packets sent by the NAT gateway
gcp.router.router_up	Router status, up or down
gcp.router.sent_routes_count	Current number of routes sent by router.

Alerts observed

Orange Business Services will set alert depending on the SOW of the Customer.

6.5.3.4 Backup and restore

Data backup and restore

The backup is based on demand Export Template IaC

Service restore

Recovery will be from Infra as Code or by Orange Business Services Operation Team actions.
The Continuous Deployment chain is used to redeploy the Cloud Router service from the configuration file of reference for production environment committed in the Git.

6.5.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA and non HA are provided by Google Cloud Platform depending on the design and service parameter configuration
Recovery after regions loss is Based on design SOW, the service can be built in multiple regions.

6.5.4 Charging model

Work Unit
Per router

6.5.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Modify/delete router Simple modification router	1 token
Create router Complex modification router	2 tokens
Other changes	Estimation in tokens based on time spent

6.6 Cloud VPN

6.6.1 Description

Cloud VPN securely extends your peer network to Google's network through an IPsec VPN tunnel. Traffic is encrypted and travels between the two networks over the public internet.

6.6.2 Build to run service included in the OTC

6.6.2.1 Build service pre-requisite

- Refer to generic description.

6.6.2.2 Build to run service

- Refer to generic description.

6.6.3 RUN services included in the MRC

6.6.3.1 Run service pre-requisite

- This file can be executed with a CI/CD and the execution has been tested successfully.

6.6.3.2 Co-manage option

No, Orange Business Services manages the Cloud VPN service.

6.6.3.3 KPI & alerts

Monitoring

Yes, Metrics, Logs, Probes

Cloud VPN can be monitorized by using Cloud Monitoring using Alerts and Metrics. Realtime Native reporting from GCP (Cloud Monitoring, Cloud Logging) can be used by OBS and Specific reporting on quote.

KPI monitored

gcp.vpn.network.dropped_received_packets_count	Ingress packets dropped for tunnel.
gcp.vpn.network.dropped_sent_packets_count	Egress packets dropped for tunnel.
gcp.vpn.network.received_bytes_count	Ingress bytes for tunnel.
gcp.vpn.network.sent_bytes_count	Egress bytes for tunnel.
gcp.vpn.tunnel_established	Indicates successful tunnel establishment if greater than 0.
gcp.router.best_received_routes_count	Number of best routes received by router.
gcp.router.bgp.received_routes_count	Number of routes received on a bgp session.
gcp.router.bgp.sent_routes_count	Number of routes sent on a bgp session.
gcp.router.bgp.session_up	Indicator for successful bgp session establishment.
gcp.router.bgp_sessions_down_count	Number of BGP sessions on the router that are down.
gcp.router.bgp_sessions_up_count	Number of BGP sessions on the router that are up.
gcp.router.router_up	Router status up or down
gcp.router.sent_routes_count	Number of routes sent by router.

Alerts observed

Orange Business Services will set alert depending on the SOW of the Customer.

6.6.3.4 Backup and restore

Data backup and restore

The backup is based on demand Export Template IaC

Service restore

Recovery will be from Infra as Code or by Orange Business Services Operation Team actions.
The Continuous Deployment chain is used to redeploy the Cloud VPN service from the configuration file of reference for production environment committed in the Git.

6.6.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA are provided by Google Cloud Platform by default.

HA VPN is a high-availability (HA) Cloud VPN solution that lets you securely connect your on-premises network to your VPC network through an IPsec VPN connection in a single region. HA VPN provides an SLA of 99.99% service availability.
Recovery after regions loss is Based on design SOW, the service can be built in multiple regions.

6.6.4 Charging model

Work Unit
Per Tunnel VPN

6.6.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Modify/delete tunnel	1 token
Create tunnel	2 tokens
Other changes	Estimation in tokens based on time spent

6.7 Cloud SQL

6.7.1 Description

Cloud SQL is a fully-managed database service that helps you set up, maintain, manage, and administer your relational databases on Google Cloud Platform. You can use Cloud SQL with MySQL, PostgreSQL, or SQL Server. Cloud SQL provides a cloud-based alternative to local MySQL, PostgreSQL, and SQL Server databases. Many applications running on Compute Engine, App Engine and other services in Google Cloud use Cloud SQL for database storage.

Each Cloud SQL instance is powered by a virtual machine (VM) running on a host Google Cloud server. Each VM operates the database program, such as MySQL Server, PostgreSQL, or SQL Server, and service agents that provide supporting services, such as logging and monitoring.

6.7.2 Build to run service included in the OTC

6.7.2.1 Build service pre-requisite

- Refer to generic description.

6.7.2.2 Build to run service

- Refer to generic description.

6.7.3 RUN services included in the MRC

6.7.3.1 Run service pre-requisite

- This file can be executed with a CI/CD and the execution has been tested successfully.

6.7.3.2 Co-manage option

Yes if CI/CD shared with the customer (IaC Part)

6.7.3.3 KPI & alerts

Monitoring

Yes, Metrics, SlowQuery Log (MySQL)

KPI monitored

- CPU utilization
- Storage usage

- Memory usage
- Read/write operations
- Ingress/Egress bytes
- MySQL queries
- MySQL questions
- Read/write InnoDB pages
- InnoDB data fsyncs
- InnoDB log fsyncs
- Active connections

Alerts observed

- CPU and memory utilization
- Disk utilization
- MySQL connections
- Auto-failover requests and replication lag

6.7.3.4 Backup and restore

Data backup and restore

The backup is based on regular export.

Service restore

Recovery will be from Infra as Code + Backup of the data.

6.7.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA and non HA are provided by Google Cloud Platform depending on the design and service parameter configuration
Recovery after regions loss is Based on design SOW, the service can be built in multiple regions.

6.7.4 Charging model

Work Unit
Per Instance

6.7.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create / update/ delete instance Create/update/delete Database (MySQL, MySQL, PostgreSQL, or SQL Server) Run script SQL	1 token
Clonage Database	2 tokens
Other changes	Estimation in tokens based on time spent

6.8 Cloud Storage

6.8.1 Description

Google Cloud Storage is a RESTful online file storage web service for storing and accessing data on Google Cloud Platform infrastructure. The service combines the performance and scalability of Google's cloud with advanced security and sharing capabilities.

6.8.2 Build to run service included in the OTC

6.8.2.1 Build service pre-requisite

- Refer to generic description.

6.8.2.2 Build to run service

- Refer to generic description.
- In addition, build to run service for Cloud Storage service will include lifecycle rules, IAM policies.

6.8.3 RUN services included in the MRC

Run a managed Cloud Storage service is optional. Depending on Customer's interest in monitoring the storage KPIs, in alerting based on KPIs, the Customer may request the service. By default, there is no recurring task proposed on Cloud Storage service, but on demand changes and on demand investigations.

6.8.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the Cloud Storage service.
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.8.3.2 Co-manage option

Yes

6.8.3.3 KPI & alerts

Monitoring

Yes, Metrics

Cloud Storage service is monitored through Cloud Monitoring. Orange Business Services will examines Cloud Storage usage (e.g., how many bytes are stored, how many download requests are coming from your applications) and will set alerts according to your SOW.

Orange Business Service will collect metrics from Google Storage to:

- Visualize the performance of your Storage services
- Correlate the performance of your Storage services with your applications

Métriques

gcp.storage.api.request_count	The number of API calls
gcp.storage.authn.authentication_count	The number of HMAC/RSA signed requests
gcp.storage.authz.acl_based_object_access_count	The number of requests that result in an object being granted access solely due to object ACLs.
gcp.storage.authz.acl_operations_count	The usage of ACL operations
gcp.storage.authz.object_specific_acl_mutation_count	The number of changes made to object specific ACLs
gcp.storage.network.received_bytes_count	The number of bytes received over the network
gcp.storage.network.sent_bytes_count	The number of bytes sent over the network
gcp.storage.storage.object_count	The total number of objects per bucket
gcp.storage.storage.total_byte_seconds	The total daily storage in byte seconds used
gcp.storage.storage.total_bytes	The total size of all objects in the bucket

6.8.3.4 Backup and restore

Data backup and restore

No backup.

Service restore

Recovery will be from Infra as Code + Backup of the data.

6.8.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA and non HA are provided by Google Cloud Platform by default for Cloud Storage service.

6.8.4 Charging model

Work Unit
Per Bucket

6.8.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Modify life cycle rules/ Chargement de données	1 token
Bucket synchronization	2 tokens
Other changes	Estimation in tokens based on time spent

6.9 Storage Transfer Service

6.9.1 Description

Storage Transfer Service is a Google Cloud product that enables you to:

- Move or backup data to a Cloud Storage bucket either from other cloud storage providers or from your on-premises storage.
- Move data from one Cloud Storage bucket to another, so that it is available to different groups of users or applications.
- Periodically move data as part of a data processing pipeline or analytical workflow.

With Storage Transfer Service, you can transfer data from other clouds, HTTP(S) and filesystems in private data centers, as well as transfer data between Google Cloud Storage buckets.

6.9.2 Build to run service included in the OTC

6.9.2.1 Build service pre-requisite

- Refer to generic description.

6.9.2.2 Build to run service

- Refer to generic description.

6.9.3 RUN services included in the MRC

Run a managed Storage Transfer Service is optional. Depending on Customer's interest in monitoring the storage KPIs, in alerting based on KPIs, the Customer may request the service. By default, there is no recurring task proposed on Storage Transfer Service, but on demand changes and on demand investigations.

6.9.3.1 Run service pre-requisite

- This file can be executed with a CI/CD and the execution has been tested successfully.

6.9.3.2 Co-manage option

No, Orange Business Services fully managed OBS manages the Storage Transfer Service.

6.9.3.3 KPI & alerts

Monitoring

Yes, Metrics, Logs, Probes

KPI monitored

- CPU
- Disk
- HTTP request and response status
- Memory
- Network
- Number of active instances

Alerts observed

- CPU
- Disk
- HTTP request and response status
- Memory
- Network
- Number of active instances

6.9.3.4 Backup and restore

Data backup and restore

The backup is based on demand Export Template IaC.
Using Google data transfer services you can easily backup data from another cloud storage provider to Google Storage Transfer Service.

Service restore

Recovery will be from Infra as Code + Backup of the data.

6.9.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA and non HA are provided by Google Cloud Platform by default.

6.9.4 Charging model

Work Unit
Per Job

6.9.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Modify/delete Job	1 token
Create Job	2 tokens
Other changes	Estimation in tokens based on time spent

6.10 Google Kubernetes Engine (Std)

6.10.1 Description

Google Kubernetes Engine (GKE) is a Google Cloud Platform (GCP) service. It is a hosted platform that allows you to run and orchestrate containerized applications. GKE manages Docker containers deployed on a cluster of machines.

GKE offers two modes of operation:

- **Standard:** You manage the underlying infrastructure of the cluster, which provides greater flexibility in configuring nodes.
- **Autopilot:** Google provisions and manages all of the underlying cluster infrastructure, including nodes and node pools. This gives you a cluster that is optimized for autonomous operation.

6.10.2 Build to run service included in the OTC

6.10.2.1 Build service pre-requisite

- Refer to generic description.

6.10.2.2 Build to run service

- Refer to generic description.

6.10.3 RUN services included in the MRC

6.10.3.1 Run service pre-requisite

- This file can be executed with a CI/CD and the execution has been tested successfully.

6.10.3.2 Co-manage option

Yes if CI/CD shared with the customer KPI & alerts

Monitoring

Yes, Insights, Metrics, logs, Health probes.

Orange Business Services will collect metrics from Docker, Kubernetes, and your containerized applications

KPI monitored

- Disk I/O
- CPU and memory usage
- Container and pod events
- Network throughput
- Individual request traces

Alerts observed

- Disk I/O
- CPU and memory usage
- Container and pod events
- Network throughput

6.10.3.3 Backup and restore

Data backup and restore

The backup is based on backup of IaC + resources k8s + data

Service restore

Recovery will be from Infra as Code + Backup of the data.

6.10.3.4 GCP SLA High Availability and Disaster Recovery inter-region

HA and non HA are provided by Google Cloud Platform depending on the design and service parameter configuration.
Recovery is based on design SOW, need actions from Operation teams of Orange Business Services.

6.10.4 Charging model

Work Unit
Per Cluster

6.10.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Add/delete node	1 token
Update Cluster	2 tokens
Modify network ranges Modify autoscaling parameters	4 tokens
Other changes	Estimation in tokens based on time spent

6.11 Google Kubernetes Engine (Autopilot)

6.11.1 Description

Google Kubernetes Engine (GKE) is a Google Cloud Platform (GCP) service. It is a hosted platform that allows you to run and orchestrate containerized applications. GKE manages Docker containers deployed on a cluster of machines.

GKE offers two modes of operation:

- **Standard:** You manage the underlying infrastructure of the cluster, which provides greater flexibility in configuring nodes.
- **Autopilot:** Google provisions and manages all of the underlying cluster infrastructure, including nodes and node pools. This gives you a cluster that is optimized for autonomous operation.

6.11.2 Build to run service included in the OTC

6.11.2.1 Build service pre-requisite

- Refer to generic description.

6.11.2.2 Build to run service

- Refer to generic description.

6.11.3 RUN services included in the MRC

6.11.3.1 Run service pre-requisite

- This file can be executed with a CI/CD and the execution has been tested successfully.

6.11.3.2 Co-manage option

Yes if CI/CD shared with the customer KPI & alerts

Monitoring

Yes, Insights, Metrics, logs, Health probes.

Orange Business Services will collect metrics from Docker, Kubernetes, and your containerized applications

KPI monitored

- Disk I/O
- CPU and memory usage
- Container and pod events
- Network throughput
- Individual request traces

Alerts observed

- Disk I/O
- CPU and memory usage
- Container and pod events
- Network throughput

6.11.3.3 Backup and restore

Data backup and restore

The backup is based on backup of IaC + resources k8s + data

Service restore

Recovery will be from Infra as Code + Backup of the data.

6.11.3.4 GCP SLA High Availability and Disaster Recovery inter-region

HA and non HA are provided by Google Cloud Platform depending on the design and service parameter configuration.
Recovery is based on design SOW, need actions from Operation teams of Orange Business Services.

6.11.4 Charging model

Work Unit
Per Cluster

6.11.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Force update cluster	1 token
Other changes	Estimation in tokens based on time spent

6.12 Compute Engine

6.12.1 Description

The Managed Service for Compute Engine is called Managed OS. OBS manages both the OS and the Compute Engine.

Orange Business Services can managed service units like OS, Middleware, Database in the Managed Compute Engine.

4 possible Managed services:

- Managed OS only
- Managed OS + Managed MW
- Managed OS + Managed DB
- Managed OS + Managed MW + Managed DB

Compute Engine is a computing and hosting service that lets you create and run virtual machines on Google infrastructure. Compute Engine offers scale, performance, and value that lets you easily launch large compute clusters on Google's infrastructure.

6.12.2 Build to run service included in the OTC

6.12.2.1 Build service pre-requisite

- Refer to generic description.

6.12.2.2 Build to run service

- Refer to generic description.

6.12.3 RUN services included in the MRC

6.12.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the Compute Engine.
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.12.3.2 Co-manage option

Yes but need to be careful to the RACI between OBS & Customer

6.12.3.3 KPI & alerts

Monitoring is performed through configuration and activation of Cloud Monitoring.
OBS backend supervision system is collecting alerts from Cloud Monitoring & Cloud Logging.

Monitoring

Yes, Insights, Metrics, logs, Health probes.

Metrics do not require installation of the Monitoring or Logging agent, but you must enable the Container-Optimized OS Health Monitoring feature.

KPI monitored for Instances:

- CPU Utilization
- Count of disk read/write bytes
- Count of disk read/write operations
- Count of throttled read/write operations
- Count of sent bytes/received bytes
- Count of incoming bytes dropped due to firewall policy
- Count of incoming packets dropped due to firewall policy

Alerts observed:

Alert on CPU, Memory Usage and Disk Usage.

Project metrics:

Like most cloud service providers, Google Compute Engine has limits on the number of resources a project may consume. If the customer are approaching (or have reached) his quota for a specific resource, OBS will tune the quota metrics for the customer if needed.

Activating Detailed Monitoring will be charged by GCP.

6.12.3.4 OS patching

GCP VM Manager

For managed OS, OBS leverages GCP VM Manager for the patching of the Operating System (OS).

Behavior: With GCP VM Manager, patches are decided by Google and all patches are to be applied if mandatory for the Compute Engine for Windows and Linux.

Additional reporting could be asked by the Customer and extra fees will be charged.

6.12.3.5 Antivirus

For managed OS, OBS leverages its central anti-virus system based on Sophos. This requires the installation of the anti-virus agent on the OS for each Compute Engine as well as the VPN connectivity to OBS Centralized Administration Zone. OBS systems allows for central reporting on Malware from its backend console system.

Would the Customer desire to keep its own Antivirus system, then OBS shall not be taken responsible for protection against viruses.

6.12.3.6 Backup and restore

Data backup and restore

By default, OBS leverages GCBDR on the Compute Engine for Managed OS. The configuration of GCBDR pattern as well as retention period shall be agreed with the Customer prior to the RUN. The first backup is full. The following backups are incremental. You can the frequency of the backup. As example: 1 x backup per week, 1x incremental backup per day per Compute Engine. The retention period depends on customer request. GCP charges will be calculated based on change rate.

Restore of Compute Engine are performed from the backup.

- In case of incident, latest version of backup can be restored
- Upon change request, a previous version of backup can be restored.

6.12.3.7 GCP SLA High Availability and Disaster Recovery inter-region

Service is Highly Available within a single Availability Zone. HA can be configured using instance group.

Multi-Availability Zones design requires specific design and subject to a specific additional charging.

This service is covered by GCBDR which enables the creation of backup copies across GCP Regions.

If this option is activated, traffic between regions and storage will be charged by GCP.

6.12.3.8 Administration tasks tracing

Actions performed by OBS managed teams on the managed OS are done from OBS Administration Zone through an access controlled by a CyberArk bastion. OBS CyberArk bastion protects the access and keep trace of the actions performed by the maintenance team allowing for audit.

The VPN connectivity to the OBS Administration Zone necessary for the management.

6.12.3.9 Login on to the Virtual Machine

For Windows OS based Compute Engine, access shall be granted by the Customer to OBS managed application operations staff through a domain account configured with proper privilege groups.

For Linux OS based Compute Engine, an encrypted key is created and provided to OBS managed application operations staff to log onto the VM.

For Applications, in case of managed application: a secret stored in a safe.

6.12.3.10 Logs

Log management is not included in the managed OS / managed Compute Engine service.
Optionally it can be activated through GCP Cloud Logging through Change Request process.

6.12.3.11 Security

By default, the MRC includes the use of security policies and groups as per customer's configuration request.
The MRC does not cover security recommendations. Security recommendations can be part of an optional security scope of work based on customer request.

6.12.3.12 Limitations

Managed Applications services is provided only for OS versions supported by the CSP vendor.

6.12.4 Charging model

Work Unit
Per Virtual Machine instance

6.12.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create a Virtual Machine	2 Tokens
Attach a Disk to a Virtual Machine	2 Tokens
Restore a Virtual Machine from a snapshot	1 Token
Backup a Virtual Machine	1 Token
Create and Deploy VMs in a Instance Group	2 Tokens
Start/Stop/Restart Virtual Machine	2 Tokens
Create/modify/delete Storage Accounts	2 Tokens
Other changes	Estimation in tokens based on time spent

6.13 Virtual Private Cloud

6.13.1 Description

Virtual Private Cloud (VPC) provides networking functionality to Compute Engine virtual machine (VM) instances, Google Kubernetes Engine (GKE) clusters, and the App Engine flexible environment. VPC provides networking for your cloud-based resources and services that is global, scalable, and flexible.

A VPC network is a global resource which consists of a list of regional virtual subnetworks (subnets) in data centers, all connected by a global wide area network. VPC networks are logically isolated from each other in the Google Cloud Platform.

At the basic level, managing Virtual Private Cloud consists in building, deploying, and maintaining the Infra as Code for it and managing the changes.

OBS has 2 prices for Managed Virtual Private Cloud depending on the number of subnets of the customer projects:

- VPC with 1 to 2 subnets
- VPC with more then 3 subnets

The management of Virtual Private Cloud is included as part of a larger bundle of Network and Security Managed services which provides network and security design, maintenance, network watching, intrusion detection, troubleshooting depending on an agreed Scope of Work.

6.13.2 Build to run service included in the OTC

6.13.2.1 Build service pre-requisite

- Refer to generic description.

6.13.2.2 Build to run service

- Refer to generic description.

6.13.3 RUN services included in the MRC

6.13.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the Virtual Private Cloud
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.13.3.2 Co-manage option

No, Orange Business Services manages the Virtual Private Cloud service.

6.13.3.3 KPI & alerts

Monitoring

Yes, Metrics, Logs (option)

Alerts observed:

Packet loss, up/down network

6.13.3.4 Backup and restore

Data backup and restore

Can be exported from Infra as Code.

Service restore

Recovery will be from Infra as Code + Backup

6.13.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA by design.
No Recovery after region loss, need to run the IaC on another region only for subnet

6.13.4 Charging model

Work Unit
Per Virtual Private Cloud instance

6.13.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Add subnet/add range IP on subnet/reservation of static address	1 token
Creation network peering	2 tokens
Other changes	Estimation in tokens based on time spent

6.14 Persistent Disk

6.14.1 Description

Persistent disks are durable network storage devices that your instances can access like physical disks in a desktop or a server. The data on each persistent disk is distributed across several physical disks.

Store data from VM instances running in Compute Engine or GKE, Persistent Disk is an Google's Cloud block storage offering.

OBS proposed 4 types of Persistent Disk:

1. Managed Standard Persistent Disk
2. Managed Balanced Persistent Disk
3. Managed SSD Persistent Disk
4. Managed Extreme Persistent Disk

6.14.2 Build to run service included in the OTC

6.14.2.1 Build service pre-requisite

- Refer to generic description.

6.14.2.2 Build to run service

- Refer to generic description.

6.14.3 RUN services included in the MRC

Run a managed Persistent Disk service is optional. Depending on Customer's mandatory if Persistent Disk is attached to managed services, the Customer may request the service.

6.14.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the Persistent Disk.
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.14.3.2 Co-manage option

No, Orange Business Services manages the Persistent Disk service.

6.14.3.3 KPI & alerts

Monitoring

Yes, Metrics

Persistent Disk service is monitored through Cloud Monitoring. Orange Business Services will examines Persistent Disk usage (e.g., how many bytes are stored, how many download requests are coming from your applications) and will set alerts according to your SOW.

Orange Business Service will collect metrics from Cloud Monitoring to:

- Graph multiple persistent disk performance metrics with **Metrics Explorer** page
- Graph average IOPS by using the **Disk read operations** metric
- Graph average throughput rates by using the **Disk read bytes** metric
- Graph maximum per second read operations by using the **Peak disk read operations** metric
- Graph average throttled operations rates by using the **Throttled read operations** metric
- Graph average throttled bytes rates by using the **Throttled read bytes** metric

6.14.3.4 Backup and restore

Data backup and restore

Backup of Iac + Disk + Data

Service restore

Recovery will be from Infra as Code + Backup of the data.

6.14.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA by design but not DR by design.
Regional Persistent Disk depending on application need, need to run the IaC on another region and restore (option)

6.14.4 Charging model

Work Unit
Per Disk

6.14.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create Disk Attach Disk to a VM	1 token
Extend Disk Mount/Format Disk	2 tokens
Enable Encryption	4 tokens
Other changes	Estimation in tokens based on time spent

6.15 Cloud Interconnect

6.15.1 Description

Cloud Interconnect provides low latency, high availability connections that enable you to reliably transfer data between your on-premises and Google Cloud Virtual Private Cloud (VPC) networks. Also, Interconnect connections provide internal IP address communication, which means internal IP addresses are directly accessible from both networks.

Cloud Interconnect offers two options for extending your on-premises network:

- Dedicated Interconnect provides a direct physical connection between your on-premises network and Google's network.
- Partner Interconnect provides connectivity between your on-premises and VPC networks through a supported service provider.

6.15.2 Build to run service included in the OTC

6.15.2.1 Build service pre-requisite

- Refer to generic description.

6.15.2.2 Build to run service

- Refer to generic description.

6.15.3 RUN services included in the MRC

6.15.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the Cloud Interconnect service.
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.15.3.2 Co-manage option

No, Orange Business Services manages the Cloud Interconnect service.

6.15.3.3 KPI & alerts

Monitoring

Yes, Insights, Metrics, Health probes

Metric

gcp.interconnect.network.attachment.capacity	Network capacity of the attachment
gcp.interconnect.network.attachment.received_bytes_count	Number of inbound bytes received.
gcp.interconnect.network.attachment.received_packets_count	Number of inbound packets received.
gcp.interconnect.network.attachment.sent_bytes_count	Number of outbound bytes sent.
gcp.interconnect.network.attachment.sent_packets_count	Number of outbound packets sent.
gcp.interconnect.network.interconnect.capacity	Active capacity of the interconnect.
gcp.interconnect.network.interconnect.dropped_packets_count	Number of outbound packets dropped due to link congestion.
gcp.interconnect.network.interconnect.link.operational	Whether the operational status of the circuit is up.
gcp.interconnect.network.interconnect.link.rx_power	Light level received over physical circuit.
gcp.interconnect.network.interconnect.link.tx_power	Light level transmitted over physical circuit.
gcp.interconnect.network.interconnect.operational	Whether the operational status of the interconnect is up.
gcp.interconnect.network.interconnect.receive_errors_count	Number of errors encountered while receiving packets.
gcp.interconnect.network.interconnect.received_bytes_count	Number of inbound bytes received.
gcp.interconnect.network.interconnect.received_unicast_packets_count	Number of inbound unicast packets received.
gcp.interconnect.network.interconnect.send_errors_count	Number of errors encountered while sending packets. Shown as error
gcp.interconnect.network.interconnect.sent_bytes_count	Number of outbound bytes sent.

gcp.interconnect.network.interconnect.sent_unicast_packets_count	Number of outbound unicast packets sent.

6.15.3.4 Backup and restore

Data backup and restore

Backup of Iac

Service restore

Recovery will be from Infra as Code and actions from Operation Team.

6.15.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA (SLA 99,9% or 99,99%) by design depending of the chosen options.
Recovery after region loss based on WAN Architecture requirement from the customer.

6.15.4 Charging model

Work Unit
Per Cloud Interconnect

6.15.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Disable my interconnect connection	1 token
Restrict interconnect usage	2 tokens
Create interco with customer configuration	> 9 tokens
Other changes	Estimation in tokens based on time spent

6.16 Big Query

6.16.1 Description

Google BigQuery is a Data Warehouse designed to allow companies to perform SQL queries very quickly thanks to the processing power of the Google Cloud infrastructure. Thus, it is part of the Infrastructure as a Cloud Service (IaaS) family. Designed for Big Data, this platform can analyze billions of rows of data.

Google BigQuery is the Big Data analysis platform offered by Google via the Cloud.

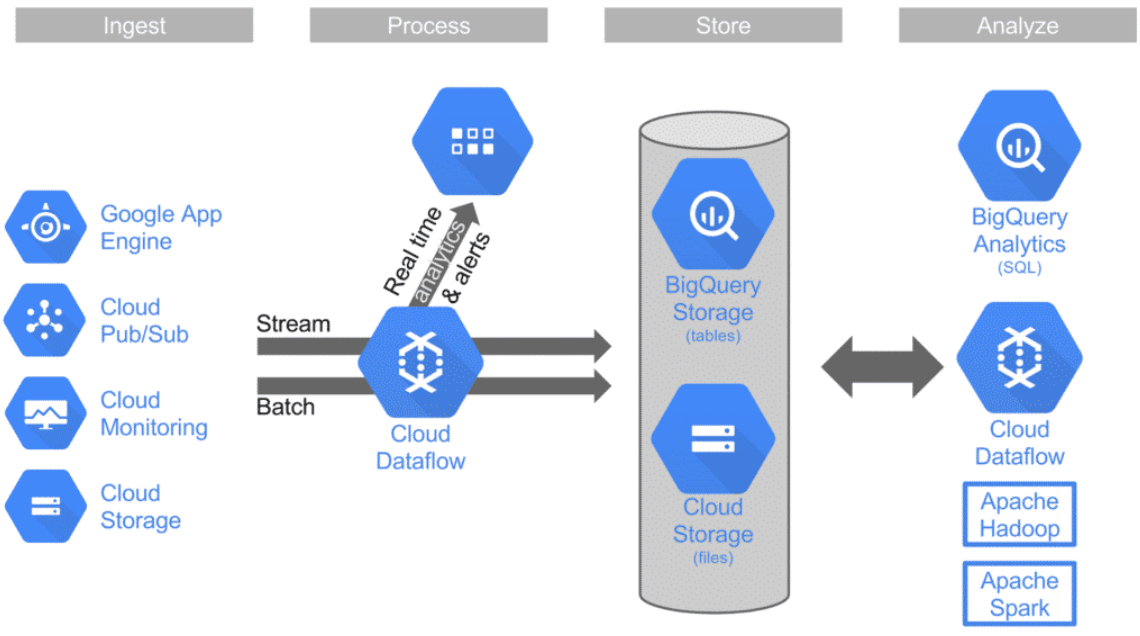

6.16.2 Build to run service included in the OTC

6.16.2.1 Build service pre-requisite

- Refer to generic description.
- Interaction loop necessary with the customer at each Build

6.16.2.2 Build to run service

- Refer to generic description.

6.16.3 RUN services included in the MRC

6.16.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the BigQuery service.
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.16.3.2 Co-manage option

No by default, IAC is fully managed by OBS, we are responsible of the CI/CD up to the dataset (the customer can have access to the tables modifications case by case. The requests for table changes are through tokens.

6.16.3.3 KPI & alerts

Monitoring

Yes, Metrics, Logs

Alerts observed:

Alerts on KPI customer per customer :

- Slot usage
- Job Concurrency
- Job performance
- Failed jobs
- Bytes processed by default in BigQuery

6.16.3.4 Backup and restore

Data backup and restore

Yes, Template IaC, Backup Regional Tables

Service restore

Recovery from Snapshot - Log - Ingestion Code -

6.16.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA and non HA are provided by Google Cloud Platform by default for BigQuery service.
BigQuery does not automatically provide a backup or replica of your data in another geographic region. You can create cross-region dataset copies to enhance your disaster recovery strategy."

6.16.4 Charging model

Work Unit
Per Table

6.16.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create table/modify table/delete table Add/modify/update/delete user with policies Copy table	1 token
Charge data from a bucket	2 tokens
Other changes	Estimation in tokens based on time spent

6.17 Pub/Sub

6.17.1 Description

Create scalable messaging and ingestion for event-driven systems and streaming analytics. Ingest events for streaming into BigQuery, data lakes or operational databases.

Pub/Sub offers a broader range of features, per-message parallelism, global routing, and automatically scaling resource capacity.

Pub/Sub allows services to communicate asynchronously, with latencies on the order of 100 milliseconds. Pub/Sub is used for streaming analytics and data integration pipelines to ingest and distribute data. It is equally effective as a messaging- oriented middleware for service integration or as a queue to parallelize tasks.

6.17.2 Build to run service included in the OTC

6.17.2.1 Build service pre-requisite

- Refer to generic description.

6.17.2.2 Build to run service

- Refer to generic description.

6.17.3 RUN services included in the MRC

6.17.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the Pub/Sub service.
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.17.3.2 Co-manage option

No by default, Iac is fully managed by Orange Business Services.

6.17.3.3 KPI & alerts

Monitoring

Yes, Metrics

Monitoring/alarm on :

- Publisher Status,
- Troughput,
- Publish Requests size,
- Topic,
- Access right

Alerts observed:

Alerts on KPI customer per customer :

- pubsub_snapshot
- pubsub_subscription
- pubsub_topic

6.17.3.4 Backup and restore

Data backup and restore

Yes, from IaC and snapshot.

Service restore

Recovery will be from snapshot.

6.17.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA are provided by Google Cloud Platform by default for Pub/Sub service. Pub/Sub is global/multi-regional with SLAs guaranteed by Google. For the highest degree of redundancy OBS can create Pub/Sub publisher clients in different GCP regions. Pub/Sub keeps any given message in a single region, although, replicated across zones

6.17.4 Charging model

Work Unit
Per instance

6.17.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create/modify/delete instance	1 token
Create snapshot msg	2 tokens

Other changes	Estimation in tokens based on time spent

6.18 Pub/Sub Lite

6.18.1 Description

Pub/Sub and Pub/Sub Lite are both horizontally scalable and managed messaging services. Pub/Sub is usually the default solution for most application integration and analytics use cases. Pub/Sub Lite is only recommended for applications where achieving extremely low cost justifies some additional operational work.

Pub/Sub Lite is a cost-effective solution that trades off operational workload, availability, and features for cost efficiency. Pub/Sub Lite requires you to manually reserve and manage resource capacity. Within Pub/Sub Lite, you can choose either zonal or regional Lite topics. Regional Lite topics offer the same availability SLA as Pub/Sub topics. However, there are reliability differences between the two services in terms of message replication.

6.18.2 Build to run service included in the OTC

6.18.2.1 Build service pre-requisite

- Refer to generic description.

6.18.2.2 Build to run service

- Refer to generic description.

6.18.3 RUN services included in the MRC

6.18.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the Pub/Sub Lite service.
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.18.3.2 Co-manage option

No by default, Iac is fully managed by Orange Business Services.

6.18.3.3 KPI & alerts

Monitoring

Yes, Metrics

Monitoring/alarm on :

- Publisher Status,
- Troughput,
- Publish Requests size,
- Reservation

Alerts observed:

Alerts on KPI customer per customer :

- pubsublite_reservation
- pubsublite_subscription_partition
- pubsublite_topic_partition

6.18.3.4 Backup and restore

Data backup and restore

Yes, from IaC and snapshot.

Service restore

Recovery will be from snapshot.

6.18.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA are provided by Google Cloud Platform by default for Pub/Sub Lite service with less resiliency & low reliability then Pub/Sub Lite service.

6.18.4 Charging model

Work Unit
Per instance

6.18.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create/modify/delete instance Reservation gestion Throughput capacity	1 token
Create snapshot msg	2 tokens
Other changes	Estimation in tokens based on time spent

6.19 Dataproc

6.19.1 Description

Dataproc is a managed Spark and Hadoop service that lets you take advantage of open source data tools for batch processing, querying, streaming, and machine learning. Dataproc automation helps you create clusters quickly, manage them easily, and save money by turning clusters off when you don't need them. With less time and money spent on administration, you can focus on your jobs and your data.

6.19.2 Build to run service included in the OTC

6.19.2.1 Build service pre-requisite

- Refer to generic description.

6.19.2.2 Build to run service

- Refer to generic description.

6.19.3 RUN services included in the MRC

6.19.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the Dataproc service.
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.19.3.2 Co-manage option

No by default, Iac is fully managed by Orange Business Services.

6.19.3.3 KPI & alerts

Monitoring

Yes, Metrics

Metric

gcp.dataproc.cluster.hdfs.datanodes	Indicates the number of HDFS DataNodes that are running inside a cluster.
gcp.dataproc.cluster.hdfs.storage_capacity	Indicates capacity of HDFS system running on cluster in GB.
gcp.dataproc.cluster.hdfs.storage_utilization	The percentage of HDFS storage currently used.
gcp.dataproc.cluster.hdfs.unhealthy_blocks	Indicates the number of unhealthy blocks inside the cluster.
gcp.dataproc.cluster.job.completion_time.avg	The time jobs took to complete from the time the user submits a job to the time Dataproc reports it is completed.
gcp.dataproc.cluster.job.completion_time.samplecount	Sample count for cluster job completion time
gcp.dataproc.cluster.job.completion_time.sumsqdev	Sum of squared deviation for cluster job completion time
gcp.dataproc.cluster.job.duration.avg	The time jobs have spent in a given state.
gcp.dataproc.cluster.job.duration.samplecount	Sample count for cluster job duration
gcp.dataproc.cluster.job.duration.sumsqdev	Sum of squared deviation for cluster job duration
gcp.dataproc.cluster.job.failed_count	Indicates the number of jobs that have failed on a cluster.
gcp.dataproc.cluster.job.running_count	Indicates the number of jobs that are running on a cluster.
gcp.dataproc.cluster.job.submitted_count	Indicates the number of jobs that have been submitted to a cluster.
gcp.dataproc.cluster.operation.completion_time.avg	The time operations took to complete from the time the user submits a operation to the time Dataproc reports it is completed.
gcp.dataproc.cluster.operation.completion_time.samplecount	Sample count for cluster operation completion time
gcp.dataproc.cluster.operation.completion_time.sumsqdev	Sum of squared deviation for cluster operation completion time
gcp.dataproc.cluster.operation.duration.avg	The time operations have spent in a given state.
gcp.dataproc.cluster.operation.duration.samplecount	Sample count for cluster operation duration
gcp.dataproc.cluster.operation.duration.sumsqdev	Sum of squared deviation for cluster operation duration
gcp.dataproc.cluster.operation.failed_count	Indicates the number of operations that have failed on a cluster.
gcp.dataproc.cluster.operation.running_count	Indicates the number of operations that are running on a cluster.
gcp.dataproc.cluster.operation.submitted_count	Indicates the number of operations that have been submitted to a cluster.
gcp.dataproc.cluster.yarn.allocated_memory_percentage	The percentage of YARN memory is allocated.

gcp.dataproc.cluster.yarn.apps	Indicates the number of active YARN applications.
gcp.dataproc.cluster.yarn.containers	Indicates the number of YARN containers.
gcp.dataproc.cluster.yarn.memory_size	Indicates the YARN memory size in GB.
gcp.dataproc.cluster.yarn.nodemanagers	Indicates the number of YARN NodeManagers running inside cluster.
gcp.dataproc.cluster.yarn.pending_memory_size	The current memory request, in GB, that is pending to be fulfilled by the scheduler.
gcp.dataproc.cluster.yarn.virtual_cores	Indicates the number of virtual cores in YARN.

6.19.3.4 Backup and restore

Data backup and restore

Yes, from IaC.

Service restore

Recovery will be from Infra as Code

6.19.3.5 GCP SLA High Availability and Disaster Recovery inter-region

Standard, Single node and HA are provided by Google Cloud Platform for Dataproc service.

6.19.4 Charging model

Work Unit
Per Cluster

6.19.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create/delete cluster	1 token
Bench/config cluster	4 tokens
Other changes	Estimation in tokens based on time spent

6.20 Dataflow

6.20.1 Description

Google Cloud Dataflow is a fully managed service for executing Apache Beam pipelines within the Google Cloud Platform ecosystem.

6.20.2 Build to run service included in the OTC

6.20.2.1 Build service pre-requisite

- Refer to generic description.

6.20.2.2 Build to run service

- Refer to generic description.

6.20.3 RUN services included in the MRC

6.20.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the Dataflow service.
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.20.3.2 Co-manage option

No by default, Iac is fully managed by Orange Business Services.

6.20.3.3 KPI & alerts

Monitoring

Yes, Metrics, Lgs

Overview metrics:

- Autoscaling
- Throughput
- CPU utilization
- Worker error log count

Streaming metrics (streaming pipelines only):

- Data freshness (with and without Streaming Engine)
- System latency (with and without Streaming Engine)
- Backlog bytes (with and without Streaming Engine)
- Parallelism (Streaming Engine only)
- Duplicates (Streaming Engine only)

Input metrics:

- Pub/Sub read, BigQuery read, etc.

Output metrics:

- Pub/Sub write, BigQuery write, etc.

6.20.3.4 Backup and restore

Data backup and restore

Yes, From Iac + Backup Pipeline by Customer

Service restore

Recovery From Iac or by Operation Team actions (Restoration). ingestion by the Customer or by OBS with procedure

6.20.3.5 GCP SLA High Availability and Disaster Recovery inter-region

Not HA by design for Dataflow service.
Dataflow does not automatically provide a backup or replica of your data in another geographic region ==> need actions from Operation teams.

If there are no grouping/time-windowing operations, a failover to another Dataflow job in another zone or region by reusing the subscription leads to no data loss in pipeline output data.

1. Job fails if region fails over : deploy 2 or more dataflow for streaming purposes
2. Streaming from PubSub (no grouping / time-windowing) : messages are acked only when persisted in destination

3. Streaming from PubSub (windowing + not rely on data before the outage) : PubSub Seek functionnality
4. Streaming from PubSub (grouping + rely on data after the outage) : Dataflow Snapshot functionnality (in preview)

6.20.4 Charging model

Work Unit
Per Job

6.20.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Delete Job	1 token
Deploy/Create Job	1 Business Hour day
Other changes	Estimation in tokens based on time spent

6.21 Cloud Composer

6.21.1 Description

Cloud Composer is a managed Apache Airflow service that helps you create, schedule, monitor and manage workflows. Cloud Composer automation helps you create Airflow environments quickly and use Airflow-native tools, such as the powerful Airflow web interface and command line tools, so you can focus on your workflows and not your infrastructure.

6.21.2 Build to run service included in the OTC

6.21.2.1 Build service pre-requisite

- Refer to generic description.

6.21.2.2 Build to run service

- Refer to generic description.

6.21.3 RUN services included in the MRC

6.21.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the Cloud Composer service.
- This file can be executed with a CI/CD and the execution has been tested successfully.

6.21.3.2 Co-manage option

Yes only if CI/CD shared with the customer

6.21.3.3 KPI & alerts

Monitoring

Yes, Metrics, logs, Health probes

Metric

gcp.composer.environment.api.request_count	Number of Composer API requests seen so far.

gcp.composer.environment.api.request_latencies.avg	Distribution of Composer API call latencies.
gcp.composer.environment.api.request_latencies.samplecount	Sample count for API request latencies
gcp.composer.environment.api.request_latencies.sumsqdev	Sum of squared deviation for API request latencies
gcp.composer.environment.dagbag_size	The current DAG bag size
gcp.composer.environment.dag_processing.parse_error_count	Number of errors raised during parsing DAG files
gcp.composer.environment.dag_processing.processes	Number of currently running DAG parsing processes
gcp.composer.environment.dag_processing.total_parse_time	Number of seconds taken to scan and import all DAG files once
gcp.composer.environment.database_health	Healthiness of Composer Airflow database
gcp.composer.environment.database.cpu.reserved_cores	Number of cores reserved for the database instance
gcp.composer.environment.database.cpu.usage_time	CPU usage time of the database instance, in seconds
gcp.composer.environment.database.cpu.utilization	CPU utilization ratio (from 0.0 to 1.0) of the database instance
gcp.composer.environment.database.disk.bytes_used	Used disk space on the database instance, in bytes
gcp.composer.environment.database.disk.quota	Maximum data disk size of the database instance, in bytes
gcp.composer.environment.database.disk.utilization	Disk quota usage ratio (from 0.0 to 1.0) of the database instance
gcp.composer.environment.database.memory.bytes_used	Memory usage of the database instance in bytes
gcp.composer.environment.database.memory.quota	Maximum RAM size of the database instance, in bytes
gcp.composer.environment.database.memory.utilization	Memory utilization ratio (from 0.0 to 1.0) of the database instance
gcp.composer.environment.executor.open_slots	Number of open slots on executor
gcp.composer.environment.executor.running_tasks	Number of running tasks on executor
gcp.composer.environment.finished_task_instance_count	Overall number of finished task instances
gcp.composer.environment.healthy	Healthiness of Composer environment.
gcp.composer.environment.num_celery_workers	Number of Celery workers.
gcp.composer.environment.num_workflows	Number of workflows.
gcp.composer.environment.scheduler_heartbeat_count	Scheduler heartbeats
gcp.composer.environment.task_queue_length	Number of tasks in queue.
gcp.composer.environment.web_server.cpu.reserved_cores	Number of cores reserved for the web server instance
gcp.composer.environment.web_server.cpu.usage_time	CPU usage time of the web server instance, in seconds
gcp.composer.environment.web_server.memory.bytes_used	Memory usage of the web server instance in bytes
gcp.composer.environment.web_server.memory.quota	Maximum RAM size of the web server instance, in bytes
gcp.composer.environment.worker.pod_eviction_count	Number of Airflow worker pods evictions

gcp.composer.workflow.run_count	Number of workflow runs completed so far.
gcp.composer.workflow.run_duration	Duration of workflow run completion.
gcp.composer.workflow.task.run_count	Number of workflow tasks completed so far.
gcp.composer.workflow.task.run_duration	Duration of task completion.

6.21.3.4 Backup and restore

Data backup and restore

From Iac + GitLab for Application Part

Service restore

Recovery from Logs and actions from Operation Team.

6.21.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA and non HA are provided by Google Cloud Platform depending on the design and service parameter configuration.
Recovery after region loss are based on design SOW, need actions from Operation teams.

6.21.4 Charging model

Work Unit
Per instance GKE

6.21.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create/modify/delete instance GKE	1 token
Add node	2 tokens
Other changes	Estimation in tokens based on time spent

6.22 Cloud Big Table

6.22.1 Description

Bigtable is a NoSQL database service, a concept that, by moving away from traditional relational databases, allows it to adapt to the needs of the modern web. These databases are indeed able to run on several different machines simultaneously, which allows to scale up and manage huge volumes of data. It is a system with horizontal scalability.

Bigtable is exposed to applications through multiple client libraries, including a supported extension to the Apache HBase library for Java. Then Bigtable integrates with the existing Apache ecosystem of open source big data software.

6.22.2 Build to run service included in the OTC

6.22.2.1 Build service pre-requisite

- Refer to generic description.

6.22.2.2 Build to run service

- Refer to generic description.

6.22.3 RUN services included in the MRC

6.22.3.1 Run service pre-requisite

- This file can be executed with a CI/CD and the execution has been tested successfully.

6.22.3.2 Co-manage option

No by default, IAC is fully managed by OBS, we are master of the CI/CD up to the table ((Customer can have access to the modifications of the column families on a case by case basis/request for change via tickets)

6.22.3.3 KPI & alerts

Monitoring

Yes, Insights, Metrics, logs, Key Visualizer

Orange Business Services monitors Cloud Bigtable using graphs available in Google Cloud Console or automatically by programming using the Cloud Monitoring API

Orange Business Services uses native tools for logs. Google Bigtable logs are collected with Google Cloud Logging and sent to a Cloud Pub/Sub via a Push HTTP forwarder.

KPI monitored

- Average CPU usage
- Storage usage
- Memory usage
- Read/write operations
- Reading latency

Alerts observed

- CPU and memory utilization
- Disk utilization

Metric

gcp.bigtable.backup.bytes_used	Backup storage used.
gcp.bigtable.autoscaling.max_node_count	Maximum number of nodes in an autoscaled cluster.
gcp.bigtable.autoscaling.min_node_count	Minimum number of nodes in an autoscaled cluster.
gcp.bigtable.autoscaling.recommended_node_count_for_cpu	Recommended number of nodes in an autoscaled cluster based on CPU usage.
gcp.bigtable.autoscaling.recommended_node_count_for_storage	Recommended number of nodes in an autoscaled cluster based on storage usage.
gcp.bigtable.cluster.cpu_load	CPU load of a cluster.
gcp.bigtable.cluster.cpu_load_by_app_profile_by_method_by_table	CPU load of a cluster split by app profile, method, and table.

gcp.bigtable.cluster.cpu_load_hottest_node	CPU load of the busiest node in a cluster.
gcp.bigtable.cluster.disk_load	Utilization of HDD disks in a cluster.
gcp.bigtable.cluster.node_count	Number of nodes in a cluster.
gcp.bigtable.cluster.storage_utilization	Storage used as a fraction of total storage capacity.
gcp.bigtable.disk.bytes_used	Amount of compressed data for tables stored in a cluster.
gcp.bigtable.disk.storage_capacity	Capacity of compressed data for tables that can be stored in a cluster.
gcp.bigtable.replication.latencies.avg	Distribution of replication request latencies for a table.
gcp.bigtable.replication.latencies.samplecount	Sample count for replication request latencies.
gcp.bigtable.replication.latencies.sumsqdev	Sum of squared deviation for replication request latencies.
gcp.bigtable.replication.max_delay	Upper bound for replication delay between clusters of a table.
gcp.bigtable.server.error_count	Number of server requests for a table that failed with an error.
gcp.bigtable.server.latencies.avg	Distribution of replication request latencies for a table.
gcp.bigtable.server.latencies.samplecount	Sample count for replication request latencies.
gcp.bigtable.server.latencies.sumsqdev	Sum of squared deviation for replication request latencies.
gcp.bigtable.server.modified_rows_count	Number of rows modified by server requests for a table.
gcp.bigtable.server.multi_cluster_failovers_count	Number of failovers during multi-cluster requests.
gcp.bigtable.server.received_bytes_count	Number of uncompressed bytes of request data received by servers for a table.
gcp.bigtable.server.request_count	Number of server requests for a table.
gcp.bigtable.server.returned_rows_count	Number of rows returned by server requests for a table.
gcp.bigtable.server.sent_bytes_count	Number of uncompressed bytes of response data sent by servers for a table.
gcp.bigtable.table.bytes_used	Amount of compressed data stored in a table.

6.22.3.4 Backup and restore

Data backup and restore

The backup is based From IaC + Snapshot from table in same zone in same cluster

Service restore

Recovery will be from other table.

6.22.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA by design.
Replication of tables on other regions necessary for recovery after region loss.

6.22.4 Charging model

Work Unit
Per Instance

6.22.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create/modify/delete table Add/mofdify/update/delete user with policies Copy table	1 token
Strategy for making optimal insertion keys	3 tokens
Reclustering table	More than 1 day
Other changes	Estimation in tokens based on time spent

6.23 Cloud Datastore

6.23.1 Description

Datastore is a NoSQL database that offers great scalability for your applications. This database automatically manages data segmentation and replication so that you have a sustainable, high-availability database that can dynamically scale to handle the load of your applications. Datastore offers a multitude of features such as ACID transactions, SQL queries, indexes and more.

- ✓ Applications can use Datastore to execute SQL-like queries that support filtering and sorting.
- ✓ Datastore replicates data across multiple data centers, providing a high level of read/write availability.
- ✓ Datastore also provides automatic scalability, high consistency for read and ancestor queries, eventual consistency for all other queries, and atomic transactions. The service has no scheduled downtime.

6.23.2 Build to run service included in the OTC

6.23.2.1 Build service pre-requisite

- Refer to generic description.

6.23.2.2 Build to run service

- Refer to generic description.

6.23.3 RUN services included in the MRC

6.23.3.1 Run service pre-requisite

- This file can be executed with a CI/CD and the execution has been tested successfully.

6.23.3.2 Co-manage option

No by default, IAC is fully managed by OBS, we are master of the CI/CD up to the table ((Customer can have access to the modifications of the column families on a case by case basis/request for change via tickets)

6.23.3.3 KPI & alerts

Monitoring

Yes, Insights, Metrics, logs

Orange Business Service collects metrics from Google Datastore to :

- Visualize the performance of your datastores
- Correlate the performance of your datastores with your applications

Orange Business Services uses native tools for logs. Cloud Datastore logs are collected with Google Cloud Logging and sent to a Cloud Pub/Sub via a Push HTTP forwarder.

Metric

gcp.datastore.api.request_count	Datastore API calls.
gcp.datastore.index.write_count	Datastore index writes.
gcp.datastore.entity.read_sizes.avg	Average of sizes of read entities.
gcp.datastore.entity.read_sizes.samplecount	Sample Count for sizes of read entities.
gcp.datastore.entity.read_sizes.sumsqdev	Sum of Squared Deviation for sizes of read entities.
gcp.datastore.entity.write_sizes.avg	Average of sizes of written entities.
gcp.datastore.entity.write_sizes.samplecount	Sample Count for sizes of written entities.
gcp.datastore.entity.write_sizes.sumsqdev	Sum of Squared Deviation for sizes of written entities.

6.23.3.4 Backup and restore

Data backup and restore

The backup is based From IaC + Snapshot from table in same zone in same cluster

Service restore

Recovery will be from other table.

6.23.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA by design.
Replication of tables on other regions necessary for recovery after region loss.

6.23.4 Charging model

Work Unit
Per Instance

6.23.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create/modify/delete table Add/mofdify/update/delete user with policies Copy table	1 token
Strategy for making optimal insertion keys	3 tokens
Reclustering table	More than 1 day
Other changes	Estimation in tokens based on time spent

6.24 Memorystore

6.24.1 Description

The Cloud Memorystore service is a data storage service in RAM entirely managed by Google and compatible with Redis. Redis is a cache management system compatible with the main CMS such as WordPress, Drupal, Magento or Prestashop. Enabling a Redis service for these applications will dramatically speed up your users' browsing experience. With the Cloud Memorystore service you can easily achieve sub-millisecond latencies and the service is calibrated to support loads consistent with the largest cache requirements.

The Cloud Memorystore service is completely isolated inside your VPC network. And only your virtual server instances have access to it. By using Cloud Memorystore you relieve your virtual server instances of redundant and unnecessary computations.

6.24.2 Build to run service included in the OTC

6.24.2.1 Build service pre-requisite

- Refer to generic description.

6.24.2.2 Build to run service

- Refer to generic description.

6.24.3 RUN services included in the MRC

6.24.3.1 Run service pre-requisite

- This file can be executed with a CI/CD and the execution has been tested successfully.

6.24.3.2 Co-manage option

No by default, IAC is fully managed by OBS, we are master of the CI/CD up to the table ((Customer can have access to the modifications of the column families on a case by case basis/request for change via tickets)

6.24.3.3 KPI & alerts

Monitoring

Yes, Metrics, Hit Logs,

Orange Business Service collects metrics from Cloud Memorystore to :

- Visualize the performance of your datastores
- Correlate the performance of your datastores with your applications

Orange Business Services uses native tools for logs. Cloud Memorystore logs are collected with Google Cloud Logging and sent to a Cloud Pub/Sub via a Push HTTP forwarder.

Metric

gcp.redis.clients.blocked	Number of blocked clients
gcp.redis.clients.connected	Number of client connections
gcp.redis.commands.calls	Total number of calls for this command
gcp.redis.commands.total_time	The amount of time in microseconds that this command took in the last second
gcp.redis.commands.usec_per_call	Average time per call over 1 minute by command
gcp.redis.keyspace.avg_ttl	Average TTL for keys in this database
gcp.redis.keyspace.keys_with_expiration	Number of keys with an expiration in this database
gcp.redis.keyspace.keys	Number of keys stored in this database
gcp.redis.persistence.rdb.bgsave_in_progress	Flag indicating a RDB save is on-going
gcp.redis.replication.master.slaves.lag	The number of bytes that replica is behind.
gcp.redis.replication.master.slaves.offset	The number of bytes that have been acknowledged by replicas.
gcp.redis.replication.master_repl_offset	The number of bytes that master has produced and sent to replicas. To be compared with replication byte offset of replica.
gcp.redis.replication.offset_diff	The number of bytes that have not been replicated to the replica. This is the difference between replication byte offset (master) and replication byte offset (replica).
gcp.redis.replication.role	Returns a value indicating the node role. 1 indicates master and 0 indicates replica.
gcp.redis.server.uptime	Uptime in seconds
gcp.redis.stats.cache_hit_ratio	Cache Hit ratio as a fraction

gcp.redis.stats.connections.total	Total number of connections accepted by the server
gcp.redis.stats.cpu_utilization	CPU, in seconds of utilization, consumed by the Redis server broken down by System/User and Parent/Child relationship
gcp.redis.stats.evicted_keys	Number of evicted keys due to max memory limit
gcp.redis.stats.expired_keys	Total number of key expiration events
gcp.redis.stats.keyspace_hits	Number of successful lookup of keys in the main dictionary
gcp.redis.stats.keyspace_misses	Number of failed lookup of keys in the main dictionary
gcp.redis.stats.memory.maxmemory	Maximum amount of memory Redis can consume
gcp.redis.stats.memory.system_memory_usage_ratio	Memory usage as a ratio of maximum system memory
gcp.redis.stats.memory.usage_ratio	Memory usage as a ratio of maximum memory
gcp.redis.stats.memory.usage	Total number of bytes allocated by Redis
gcp.redis.stats.network_traffic	Total number of bytes sent to/from redis (includes bytes from commands themselves, payload data, and delimiters)
gcp.redis.stats.pubsub.channels	Global number of pub/sub channels with client subscriptions
gcp.redis.stats.pubsub.patterns	Global number of pub/sub pattern with client subscriptions
gcp.redis.stats.reject_connections_count	Number of connections rejected because of max clients limit

6.24.3.4 Backup and restore

Data backup and restore

The backup is based From IaC + Snapshot from table in same zone in same cluster

Service restore

Recovery will be from Dump of the DB

6.24.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA by design.
Replication of tables on other regions necessary for recovery after region loss.

6.24.4 Charging model

Work Unit

Per Instance

6.24.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create/modify/delete table Add/mofdify/update/delete user with policies Copy table	1 token
Optimisation of index	4 tokens
Other changes	Estimation in tokens based on time spent

6.25 Cloud Firestore

6.25.1 Description

Cloud Firestore is a document-oriented NoSQL database that automatically manages data partitioning and replication to ensure reliability, while being able to scale up according to application needs. And it does so automatically.

Google Cloud Firestore is also a flexible and scalable database for mobile, web and server development from Firebase and Google Cloud Platform.

6.25.2 Build to run service included in the OTC

6.25.2.1 Build service pre-requisite

- Refer to generic description.

6.25.2.2 Build to run service

- Refer to generic description.

6.25.3 RUN services included in the MRC

6.25.3.1 Run service pre-requisite

- This file can be executed with a CI/CD and the execution has been tested successfully.

6.25.3.2 Co-manage option

Yes if CI/CD shared with the customer (IaC Part)

6.25.3.3 KPI & alerts

Monitoring

Yes, Metrics, SlowQuery Log (FireStore)

Orange Business Services uses native tools for logs. Cloud Firestore logs are collected with Google Cloud Logging and sent to a Cloud Pub/Sub via a Push HTTP forwarder.

Metric

gcp.firestore.document.delete_count	The number of successful document deletes.
gcp.firestore.document.read_count	The number of successful document reads from queries or lookups.
gcp.firestore.document.write_count	The number of successful document

6.25.3.4 Backup and restore

Data backup and restore

The backup based on regular export.

Service restore

Recovery will be from Infra as Code + Backup of the data.

6.25.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA and non HA are provided by Google Cloud Platform depending on the design and service parameter configuration
Recovery after regions loss is Based on design SOW, the service can be built in multiple regions.

6.25.4 Charging model

Work Unit
Per Instance

6.25.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create/update/delete instance Create/update/delete BD Run script FireStore	1 token
Index refactoring	4 tokens
Other changes	Estimation in tokens based on time spent

6.26 Cloud Spanner

6.26.1 Description

Cloud Spanner is a sensitive relational database service, which is fully managed and designed to offer transactional consistency on a global scale. It provides schemas, SQL (ANSI 2011 with extensions) and automatic synchronous replication to guarantee high availability.

Cloud Spanner benefits:

- High consistency, including highly consistent secondary indexes,
- SQL compatibility with ALTER statements for schema modifications,
- Managed instances guarantee high availability through integrated, transparent and synchronous data replication.

Cloud Spanner offers regional and multi-regional instance configurations.

6.26.2 Build to run service included in the OTC

6.26.2.1 Build service pre-requisite

- Refer to generic description.

6.26.2.2 Build to run service

- Refer to generic description.

6.26.3 RUN services included in the MRC

6.26.3.1 Run service pre-requisite

- This file can be executed with a CI/CD and the execution has been tested successfully.

6.26.3.2 Co-manage option

Yes if CI/CD shared with the customer (IaC Part)

6.26.3.3 KPI & alerts

Monitoring

Yes, Metrics

Orange Business Services uses native tools for logs. Cloud Spanner logs are collected with Google Cloud Logging and sent to a Cloud Pub/Sub via a Push HTTP forwarder.

Metrics

gcp.spanner.api.received_bytes_count	Uncompressed request bytes received by Cloud Spanner.
gcp.spanner.api.sent_bytes_count	Uncompressed response bytes sent by Cloud Spanner.
gcp.spanner.api.api_request_count	Cloud Spanner API requests.
gcp.spanner.api.request_count	Rate of Cloud Spanner API requests.
gcp.spanner.api.request_latencies.avg	Average server request latencies for a database.
gcp.spanner.api.request_latencies.samplecount	Sample count of server request latencies for a database.
gcp.spanner.api.request_latencies.sumsqdev	Sum of Squared Deviation of server request latencies for a database.
gcp.spanner.api.request_latencies_by_transaction_type	Distribution of server request latencies by transaction types.
gcp.spanner.instance.cpu.utilization	Utilization of provisioned CPU, between 0 and 1.
gcp.spanner.instance.cpu.smoothed_utilization	24-hour smoothed utilization of provisioned CPU between 0.0 and 1.0.
gcp.spanner.instance.cpu.utilization_by_operation_type	Percent utilization of provisioned CPU, by operation type between 0.0 and 1.0.
gcp.spanner.instance.cpu.utilization_by_priority	Percent utilization of provisioned CPU, by priority between 0.0 and 1.0.
gcp.spanner.instance.node_count	Total number of nodes.
gcp.spanner.instance.session_count	Number of sessions in use.
gcp.spanner.instance.storage.used_bytes	Storage used in bytes.

gcp.spanner.instance.storage.limit_bytes	Storage limit for instance in bytes
gcp.spanner.instance.storage.limit_bytes_per_processing_unit	Storage limit per processing unit in bytes.
gcp.spanner.instance.storage.utilization	Storage used as a fraction of storage limit.
gcp.spanner.instance.backup.used_bytes	Backup storage used in bytes.
gcp.spanner.instance.leader_percentage_by_region	Percentage of leaders by cloud region between 0.0 and 1.0.
gcp.spanner.instance.processing_units	Total number of processing units.
gcp.spanner.lock_stat.total.lock_wait_time	Total lock wait time for lock conflicts recorded for the entire database.
gcp.spanner.query_count	Count of queries by database name, status, query type, and used optimizer version.
gcp.spanner.query_stat.total.bytes_returned_count	Number of data bytes that the queries returned
gcp.spanner.query_stat.total.cpu_time	Number of seconds of CPU time Cloud Spanner spent on operations to execute the queries.
gcp.spanner.query_stat.total.execution_count	Number of times Cloud Spanner saw queries during the interval.
gcp.spanner.query_stat.total.failed_execution_count	Number of times queries failed during the interval.
gcp.spanner.query_stat.total.query_latencies	Distribution of total length of time, in seconds, for query executions within the database.
gcp.spanner.query_stat.total.returned_rows_count	Number of rows that the queries returned.
gcp.spanner.query_stat.total.scanned_rows_count	Number of rows that the queries scanned excluding deleted values.
gcp.spanner.read_stat.total.bytes_returned_count	Total number of data bytes that the reads returned excluding transmission encoding overhead.
gcp.spanner.read_stat.total.client_wait_time	Number of seconds spent waiting due to throttling.

gcp.spanner.read_stat.total.cpu_time	Number of seconds of CPU time Cloud Spanner spent execute the reads excluding prefetch CPU and other overhead. *Shown as second*
gcp.spanner.read_stat.total.execution_count	Number of times Cloud Spanner executed the read shapes during the interval.
gcp.spanner.read_stat.total.leader_refresh_delay	Number of seconds spent coordinating reads across instances in multi-region configurations.
gcp.spanner.read_stat.total.locking_delays	Distribution of total time in seconds spent waiting due to locking.
gcp.spanner.read_stat.total.returned_rows_count	Number of rows that the read returned.
gcp.spanner.row_deletion_policy.deleted_rows_count (count)	Count of rows deleted by the policy since the last sample.
gcp.spanner.row_deletion_policy.processed_watermark_age	Time between now and the read timestamp of the last successful execution.
gcp.spanner.row_deletion_policy.undeletable_rows	Number of rows in all tables in the database that can't be deleted.
gcp.spanner.transaction_stat.total.bytes_written_count	Number of bytes written by transactions.
gcp.spanner.transaction_stat.total.commit_attempt_count	Number of commit attempts for transactions.
gcp.spanner.transaction_stat.total.commit_retry_count	Number of commit attempts that are retries from previously aborted transaction attempts.
gcp.spanner.transaction_stat.total.participants	Distribution of total number of participants in each commit attempt.
gcp.spanner.transaction_stat.total.transaction_latencies	Distribution of total seconds taken from the first operation of the transaction to commit or abort.

6.26.3.4 Backup and restore

Data backup and restore

Yes, backup automatic include

Service restore

Recovery will be from Infra as Code + Backup of the data.

6.26.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA and Multiregional
Recovery after regions loss : Managed Service, Serverless, everything is managed by Google

6.26.4 Charging model

Work Unit
Per Instance

6.26.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create/update/delete BD	1 token
Modification of the DB schema	4 tokens
Other changes	Estimation in tokens based on time spent

6.27 Cloud Run

6.27.1 Description

Cloud Run is a serverless solution for hosting containers.

Cloud Run enables services to be hosted in containers using a serverless approach. This service can be categorized as CaaS (Container as a Service). To exploit the code and develop an application, DevOps can use GCP Cloud Run.

Cloud Run is a platform that lets you run your code directly on Google's infrastructure. Through GCP Cloud Run, your various services can be hosted in containers, all in a serverless approach. It's a container-based tool that guarantees separation between your tasks and the platform.

6.27.2 Build to run service included in the OTC

6.27.2.1 Build service pre-requisite

- Refer to generic description.

6.27.2.2 Build to run service

- Refer to generic description.

6.27.3 RUN services included in the MRC

6.27.3.1 Run service pre-requisite

- This file can be executed with a CI/CD and the execution has been tested successfully.

6.27.3.2 Co-manage option

Yes if CI/CD shared with the customer (IaC Part)

6.27.3.3 KPI & alerts

Monitoring

Yes, Logs, Metrics

Orange Business uses native tools for logs. Cloud Run logs are collected with Google Cloud Logging. Audit logs are also available for Cloud Run with Cloud Audit Logs.

Metrics

gcp.run.container.billable_instance_time	Billable time aggregated from all container instances of the revision (ms/s).
gcp.run.container.cpu.allocation_time	Container CPU allocation of the revision in seconds.
gcp.run.container.cpu.utilizations.avg	The average distribution of container CPU utilization distribution across all container instances of the revision.
gcp.run.container.cpu.utilizations.p95	The 95th percentile distribution of container CPU utilization distribution across all container instances of the revision.
gcp.run.container.cpu.utilizations.p99	The 99th percentile distribution of container CPU utilization distribution across all container instances of the revision.
gcp.run.container.cpu.utilizations.samplecount	Sample count of the distribution of service request times in milliseconds.
gcp.run.container.instance_count	The number of container instances that exist, broken down by state.
gcp.run.container.max_request_concurrencies.avg	Average of the maximum number of concurrent requests being served by each container instance over a minute.
gcp.run.container.max_request_concurrencies.p95	95th percentile distribution of the maximum number of concurrent

	requests being served by each container instance over a minute.
gcp.run.container.max_request_concurrencies.p99	99th percentile distribution of the maximum number of concurrent requests being served by each container instance over a minute.
gcp.run.container.max_request_concurrencies.samplecount	Sample count of the distribution of the maximum number of concurrent requests being served by each container instance over a minute.
gcp.run.container.memory.allocation_time	Container memory allocation of the revision in Gigabytes-seconds.
gcp.run.container.memory.utilizations.avg	Average of the container memory utilization distribution across all container instances of the revision.
gcp.run.container.memory.utilizations.p95	95th percentile distribution of the container memory utilization distribution across all container instances of the revision.
gcp.run.container.memory.utilizations.p99	99th percentile distribution of the container memory utilization distribution across all container instances of the revision.
gcp.run.container.memory.utilizations.samplecount	Sample count of the container memory utilization distribution across all container instances of the revision.
gcp.run.container.network.received_bytes_count	The incoming socket and HTTP response traffic of revision, in bytes.
gcp.run.container.network.sent_bytes_count	The outgoing socket and HTTP response traffic of revision, in bytes.

gcp.run.request_count	The number of service requests.
gcp.run.request_latencies.avg (gauge)	Average distribution of service request times in milliseconds. Shown as millisecond
gcp.run.request_latencies.p95	The 95th percentile distribution of service request times in milliseconds.
gcp.run.request_latencies.p99	The 99th percentile distribution of service request times in milliseconds.
gcp.run.request_latencies.samplecount	Sample count of the distribution of service request times in milliseconds.
gcp.run.request_latencies.sumsqdev	Sum of squared deviation of the distribution of service request times in milliseconds.

6.27.3.4 Backup and restore

Data backup and restore

The backup is based on backup of IaC + data

Service restore

Recovery after region loss will be based on design SOW, need actions from Operation teams.

6.27.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA and non HA are provided by GCP depending on the design and service parameter configuration

6.27.4 Charging model

Work Unit
Per lines of code

6.27.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create/Modify/Remove Cloud Native Users	1 token
Create/Modify/Remove a Container Image	3 tokens
Other changes	Estimation in tokens based on time spent

6.28 Cloud Functions

6.28.1 Description

Google Cloud Functions is a serverless compute platform that allows you to run code in response to events without having to provision or manage servers. Because Cloud Function is a fully managed service, it is a great way to efficiently automate tasks, build microservices, and connect your applications to other cloud products and services, both on and off Google Cloud.

Cloud Functions falls into the Functions as a Service (FaaS) category of computing. FaaS is all about the code–and just the code. With Google Cloud Functions, you have your choice of working with a range of runtimes: Go, Java, .NET Core, Node.js, PHP, Python or Ruby.

6.28.2 Build to run service included in the OTC

6.28.2.1 Build service pre-requisite

- Refer to generic description.

6.28.2.2 Build to run service

- Refer to generic description.

6.28.3 RUN services included in the MRC

6.28.3.1 Run service pre-requisite

- This file can be executed with a CI/CD and the execution has been tested successfully.

6.28.3.2 Co-manage option

Yes if CI/CD shared with the customer (IaC Part)

6.28.3.3 KPI & alerts

Monitoring

Yes, Logs, Metrics

Orange Business uses native tools for logs. Cloud Functions logs are collected with Google Cloud Logging.

Metrics

gcp.cloudfunctions.function.active_instances	The number of active function instances
gcp.cloudfunctions.function.execution_count (count)	The number of function executions.
gcp.cloudfunctions.function.execution_times.avg	Average of functions execution times.
gcp.cloudfunctions.function.execution_times.p95	95th percentile of functions execution times.
gcp.cloudfunctions.function.execution_times.p99	99th percentile of functions execution times.
gcp.cloudfunctions.function.execution_times.samplecount	Sample count for functions execution times.
gcp.cloudfunctions.function.execution_times.sumsqdev	Sum of squared deviation for functions execution times.
gcp.cloudfunctions.function.instance_count	The number of function instances broken down by state

gcp.cloudfunctions.function.network_egress	The outgoing network traffic of a function
gcp.cloudfunctions.function.user_memory_bytes.avg	The average function memory usage during execution
gcp.cloudfunctions.function.user_memory_bytes.p95	The 95th percentile of function memory usage during execution
gcp.cloudfunctions.function.user_memory_bytes.p99	The 99th percentile of function memory usage during execution
gcp.cloudfunctions.function.user_memory_bytes.samplecount	The sample count for a function's memory usage.
gcp.cloudfunctions.function.user_memory_bytes.sumsqdev	The sum of squared deviation for function's memory usage.
gcp.cloudfunctions.function.active_instances	The number of active function instances
gcp.cloudfunctions.function.execution_count	The number of function executions.
gcp.cloudfunctions.function.execution_times.avg	Average of functions execution times.
gcp.cloudfunctions.function.execution_times.p95	95th percentile of functions execution times.
gcp.cloudfunctions.function.execution_times.p99	99th percentile of functions execution times.
gcp.cloudfunctions.function.execution_times.samplecount	Sample count for functions execution times.
gcp.cloudfunctions.function.execution_times.sumsqdev	Sum of squared deviation for functions execution times.
gcp.cloudfunctions.function.instance_count	The number of function instances broken down by state
gcp.cloudfunctions.function.network_egress	The outgoing network traffic of a function
gcp.cloudfunctions.function.user_memory_bytes.avg	The average function memory usage during execution

6.28.3.4 Backup and restore

Data backup and restore

The backup is based on backup of IaC + data

Service restore

Recovery after region loss will be based on design SOW, need actions from Operation teams.

6.28.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA and non HA are provided by GCP depending on the design and service parameter configuration

6.28.4 Charging model

Work Unit
Per lines of code

6.28.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create/Modify/Remove Cloud Native Users Request log files	1 token

Create/Modify/Remove Cloud Functions	2 tokens
Other changes	Estimation in tokens based on time spent

6.29 Cloud Scheduler

6.29.1 Description

Google Cloud Scheduler is a fully managed serverless task scheduling service that enables users to schedule, automate and manage tasks on various Google Cloud services. Google Cloud Scheduler enables developers to define time-based event triggers, which can execute tasks, workflows or even easily call external services.

Google Cloud Scheduler offers robust integration with other Google Cloud services such as Cloud Functions, Cloud Run and AppEngine, enabling developers to create complex data-driven applications without the need for extensive infrastructure management or maintenance. This seamless integration allows developers to concentrate on the design and implementation of application logic, while Google Cloud Scheduler takes care of the execution of scheduled tasks and the management of underlying infrastructure resources. What's more, Cloud Scheduler guarantees secure task execution by offering authentication support for tasks, enabling developers to control access to underlying services and thus maintaining overall application security.

Google Cloud Scheduler represents a powerful, scalable and reliable solution for automating and managing time-based tasks in serverless computing environments. Google Cloud Scheduler enables developers to concentrate on application logic and functionality, while infrastructure planning and management aspects are handled seamlessly behind the scenes, guaranteeing a consistent and enjoyable development experience.

6.29.2 Build to run service included in the OTC

6.29.2.1 Build service pre-requisite

- Refer to generic description.

6.29.2.2 Build to run service

- Refer to generic description.

6.29.3 RUN services included in the MRC

6.29.3.1 Run service pre-requisite

- This file can be executed with a CI/CD and the execution has been tested successfully.

6.29.3.2 Co-manage option

No by default, IAC is fully managed by OB, we are master of the CI/CD.

6.29.3.3 KPI & alerts

Monitoring

Yes, Logs

Orange Business uses native tools for logs. Cloud Scheduler logs are collected with Google Cloud Logging.

6.29.3.4 Backup and restore

Data backup and restore

The backup is based on backup of IaC

Service restore

Recovery after region loss will be based on design SOW, need actions from Operation teams.

6.29.3.5 GCP SLA High Availability and Disaster Recovery inter-region

HA and non HA are provided by GCP depending on the design and service parameter configuration

6.29.4 Charging model

Work Unit
Per jobs

6.29.5 Changes catalogue – in Tokens, per act

Changes examples	Effort
Create/Modify/Remove Cloud Native Users	1 token
Create/Modify/Remove a scheduling rules	2 tokens
Other changes	Estimation in tokens based on time spent

7 End of the document