Business

Annexe technique au Descriptif de Service Managed Applications
Managed Applications pour GCP

Table des matières

1 Aperçu du service

1.1 Description Générale

Ce document est une annexe à la description du service Managed Applications. Il fournit une description du service et des détails supplémentaires pour :

- Application d'Entreprise gérée sur Google Cloud Platform
- Les Services Natifs dans le Cloud Managé sur Google Cloud Platform

La description s'ajoute aux services gérés déjà décrits dans l'autre document intitulé Managed Applications Service Description :

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

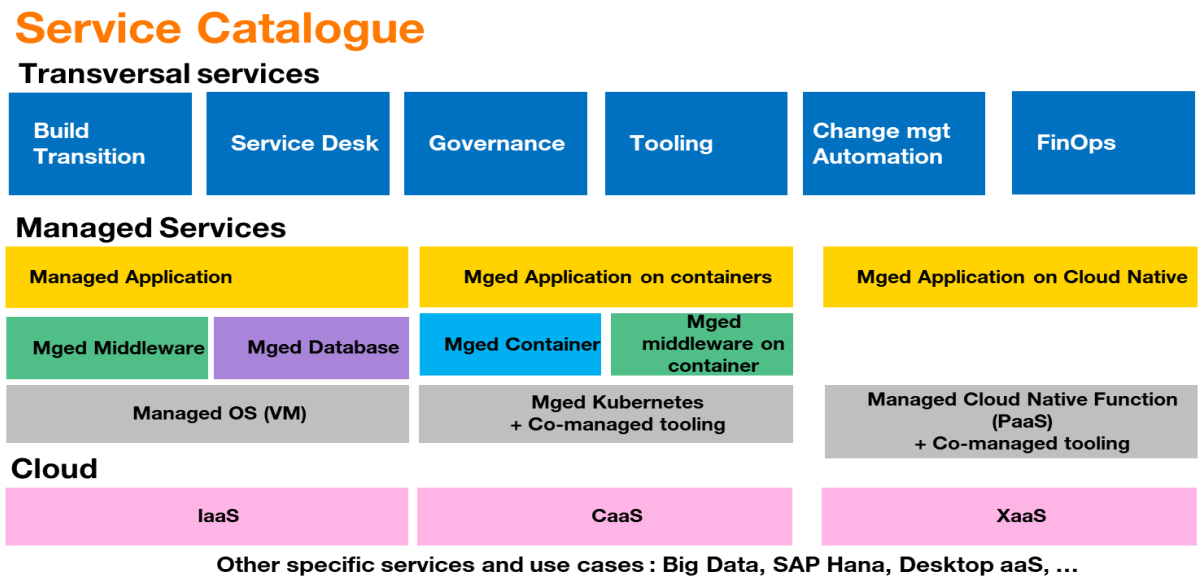

2 Managed Cloud Native Services sur GCP

Les applications métier du client déployées sur GCP sont dépendantes des Services Cloud Native GCP (IaaS, PaaS). Orange Business Services fournit les services managés nécessaires pour garantir l'assurance de service et la gestion des changements pour ces dépendances, ainsi que la configuration et le déploiement pour les construire et les restaurer.

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	Activités
Services Natifs Hyperscalers **Phase d'Implémentation**	▪ Revue et validation du RACI des services applicatifs du Client GCP par le Prestataire ▪ Création de l'infrastructure as code : selon la classe de transition ▪ Revue et ajustement des fiches réflexes (MOP sur incident) fournies par l'entreprise du Client au Prestataire (Lorsqu'applicable en transition où l'environnement du Client existe et en cas d'applications managées) ▪ Reprise et/ou élaboration de la documentation à l'usage des équipes du Prestataire ▪ Co-définition et/ou révision des alarmes et des seuils des applications ▪ Création d'accès pour les administrateurs du Prestataire ▪ Configuration du fonctionnement du VPN (si nécessaire) ▪ Configuration et test des alarmes dans le système de surveillance centralisé du Prestataire ▪ Formation des Clients sur le Cloud Store pour l'accès aux demandes de changement/incident.
Services Natifs Hyperscalers **Phase Opération**	▪ Supervision et exploitation ○ Lecture et analyse des alarmes ○ Maintien de l'IaC (hors changements) selon la classe de transition ○ Correction des configurations défectueuses ○ Revue conjointe puis mise à jour des groupes de sécurité et des contrôles d'accès ○ Gestion des événements (changements & incidents) et interfaçage avec le support GCP si nécessaire et l'exploitation des applications ○ Supervision du service 24/7

2.2 Les services cloud native

On peut typiquement distinguer 3 catégories de services :

1. Les services plan utilisateur : si une application métier en dépend, l'application métier est susceptible d'être affectée par un défaut de celui-ci. Le service n'a pas de données persistantes, donc la récupération ne nécessite pas de restauration de données.
2. Les services de données : si une application métier dépend d'un service de données, l'application métier est susceptible d'être affectée par un défaut de celui-ci. Le service a des données persistantes, donc une récupération peut nécessiter une restauration des données. La perte de données, la corruption de données peuvent également affecter l'application métier.
3. Les autres services : l'application métier ne dépend pas d'eux. La plupart de ces services sont utilisés pour les automatismes, l'observation, la migration. La perte du service n'est pas susceptible d'affecter l'application métier. Certains de ces services sont utilisés pour gérer les services du plan utilisateur et du plan de données de l'application métier, d'autres ont un usage spécifique pour lequel un cahier des charges sera établi si le client demande à OBS de les exploiter dans le cadre du service managé fourni.

Services utilisateur	Services de données	Autres services

Calcul □ Compute Engine □ App Engine **Réseau** □ Cloud CDN □ Cloud Interconnect □ Cloud DNS □ Cloud Load Balancing □ Cloud VPN □ Cloud NAT □ Virtual Private Cloud □ Storage Transfer Service	**Stockage** □ Persistent Disk □ Cloud Storage **Bases de données** □ Datastore □ Cloud Bigtable □ Cloud SQL **Analyse des données** □ BigQuery □ Pub/Sub □ Pub/Sub Lite □ Dataproc □ Dataflow □ Cloud Composer	**Containers** □ Google Kubernetes Engine **Opération** □ Cloud Audit Logs □ Cloud Logging □ Cloud Monitoring

Services Cloud natifs GCP par catégorie

2.3 Tâches liées à la gestion des services Cloud natifs

Les tâches nécessaires à la gestion d'un service cloud natif dépendent du service. Elles consistent à :

- Configurer et déployer le service : L'Infrastructure as Code est exploitée afin de configurer le service, l'observabilité, la sauvegarde. L'expertise de niveau 3 sur le service est mise à profit pour une mise en œuvre correcte grâce à l'étendue du travail (voir la description détaillée des services de Build et SRE).
- Application du groupe de sécurité et de la politique de contrôle d'accès définis par le client.
- Récupération du service grâce à l'Infrastructure as Code : en cas de panne, la plupart des services doivent être récupérés grâce à un redéploiement. Reconfigurer le service manuellement à partir de zéro n'est pas une option efficace : cela prend du temps et est source d'erreurs. C'est pourquoi la récupération / le redéploiement à partir d'Infrastructure as Code est préféré.
- La supervision et la remédiation consistent à surveiller les alarmes levées sur le service pendant la plage de surveillance (typiquement : 8x5 ou 24x7). Lorsqu'une alarme survient, un ticket d'incident est créé, une priorité est attribuée, le client est notifié. Ensuite, une action corrective est entreprise grâce aux procédures mises à disposition des niveaux 2/1 par le niveau 3. La remédiation sur un service cloud native peut être nécessaire pour rétablir le service de l'application métier. Si la procédure ne permet pas de remédier à l'incident, alors l'incident est remonté au niveau 3. Si la cause première est le CSP lui-même, l'incident est porté à l'attention du CSP par le niveau 3.
- Sauvegarde et restauration : selon le service (si le service a une persistance), il est nécessaire de sauvegarder les données du service. Le service de gestion consiste à configurer la solution de sauvegarde et à en contrôler le bon fonctionnement. Note : la solution de sauvegarde doit être souscrite séparément, par exemple GCBDR (Backup & DR GCP). La restauration du service sur incident peut impliquer la restauration des données à partir d'une sauvegarde.
- Corrections et anti-virus du système d'exploitation : le maintien du système d'exploitation à jour et sans virus est un service managé par Managed Virtual Machine / Managed OS. **Référez-vous à la description détaillée**.

- Spécificités : certains services natifs du cloud peuvent avoir des tâches de configuration ou de gestion spécifiques.
- Spécificités de l'application métier : par défaut, les alertes standard sont surveillées. La configuration d'alertes, de logs sur un service cloud natif qui sont spécifiques à une application métier fait l'objet d'un périmètre de travail spécifique.

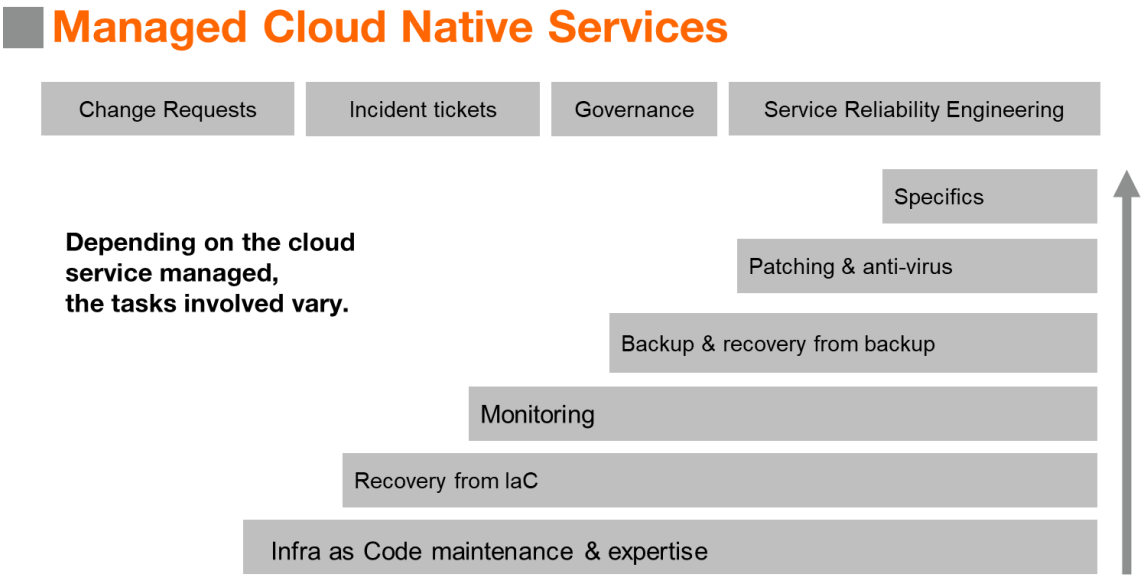

Tâches liées à la gestion des services Cloud natifs

Selon le service cloud natif managé, plus ou moins de tâches de management sont nécessaires et incluses dans le service managé. Cela détermine la complexité du service managé.

Les tâches impliquées dépendent généralement de la catégorie du service cloud natif, qu'il s'agisse du plan utilisateur, du plan de données dont dépend l'application métier ou d'autres services dont l'application métier ne dépend pas.

	Modèle de facturation	Services du plan utilisateur	Services du plan données	Autres services
Objectif		Utilisé pour soutenir l'application client	Utilisé pour soutenir l'application du client	Utilisé pour faire fonctionner le plan utilisateur ou le plan de données
Build	Charge unique basée sur l'étendue des travaux	IaC dans Git, poussé via CI / CD	IaC dans Git, poussé via CI / CD	IaC dans Git, poussé via CI / CD
Maintien du IaC sans changement	Frais mensuels récurrents	Oui	Oui	Oui
Surveillance et alertes	Frais mensuels récurrents	Oui	Oui	
Restauration de la configuration sur incident	Inclus dans le MRC	Oui, par export du IaC	Oui, par export de l'IaC	Yes, du IaC quand c'est applicable
Sauvegarde et restauration	Inclus dans le MRC		Oui	

des données sur incident				
Gestion des réseaux et de la sécurité	Basé sur l'étendue des travaux	Optionnel : Basé sur l'étendue des travaux	Optionnel : Basé sur l'étendue des travaux	
Service Desk	Par ticket d'incident ou pourcentage	Oui	Oui	Oui
Gestion des changements	Par changement, en jetons par rapport à la complexité	Via IaC dans Git, poussé via CI / CD	Via IaC dans Git, poussé via CI / CD	Via IaC dans Git, poussé via CI / CD
Reprise après sinistre	Conception et devis spécifiques	Optionnel : Basée sur l'étendue des travaux	Optionnel : Basée sur l'étendue des travaux	

2.4 Liste non-exhaustive des tâches incluses dans la gestion d'un service "Cloud Natif"

GCP service	Type	Configuration	Surveillance et alertes configurées dans Google Cloud Monitoring	Sauvegarde configurée dans GCBDR	Procédure de récupération	Gestion des correctifs	Antivirus management	Spécificités
Prérequis en cas de		**Classe 2, classe 4 lorsque GCBDR n'est pas disponible pour le service.**	**Classe 2**	**Classe 2**	**Classe 2 Si différent d'une restauration, alors Classe 4, Classe 5**	**Classe 2**	**n/a**	
Cloud Load Balancing	Managé	Terraform ou Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	A partir du IaC	n/a	n/a	
Cloud DNS	Managé	Terraform ou Google Cloud Deploy	Google Cloud Monitoring	n/a	A partir du IaC	n/a	n/a	Le DNS privé doit être inclus dans le tenant managé.

		ment Manager						
Content Delivery Network (CDN)	Managé	Terraform ou Google Cloud Deployment Manager	Google Cloud Monitoring		A partir du IaC	n/a	n/a	HA par conception
Cloud NAT	Managé	Terraform ou Google Cloud Deployment Manager	Google Cloud Monitoring		A partir du IaC	n/a	n/a	Nativement redondant
Cloud Router	Managé	Terraform ou Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	A partir du IaC	n/a	n/a	Complexe Récupération du Iac ou par des actions de l'équipe opération
Cloud VPN	Managé	Terraform ou Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	A partir du IaC	n/a	n/a	HA en option par conception
Cloud SQL	Managé	Terraform ou Google Cloud Deployment Manager	Google Cloud Monitoring	GCBDR	A partir du IaC	n/a	n/a	
Cloud Storage	Managé	Terraform ou Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	A partir du IaC	n/a	n/a	
Storage Transfer Service	Managé	Terraform ou Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	A partir du IaC	n/a	n/a	

Google Kubernetes Engine (Std)	Managé	Terraform ou Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	A partir du IaC	n/a	n/a	La gestion des correctifs est incluse dans le service
Google Kubernetes Engine (Autopilot)	Managé	Terraform ou Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	A partir du IaC	n/a	n/a	La gestion des correctifs est incluse dans le service
Compute Engine	Managé	Terraform ou Google Cloud Deployment Manager	Google Cloud Monitoring	GCBDR	A partir du IaC	Google VM Manager	OBS Sophos	Uniquement les versions d'OS supportées
Virtual Private Cloud	Managé	Terraform ou Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	A partir du IaC	n/a	n/a	
Persistent Disk	Managé	Terraform ou Google Cloud Deployment Manager	Google Cloud Monitoring	n/a	A partir du IaC	n/a	n/a	

Liste non-exhaustive des tâches incluses dans la gestion d'un service "Cloud Natif" (Extrait des services)

2.5 Outils utilisés pour les services managés cloud natifs

Les outils GCP et les outils d'OBS pour les opérations backend sont utilisés pour fournir les services managés. Si le client souhaite utiliser un autre outil, la faisabilité doit être confirmée avec OBS et le RACI et les unités d'œuvre peuvent être révisés.

Processus	Outil utilisé par le Delivery OBS MA
Configuration de l'infrastructure	Script Terraform ou Google Cloud Deployment Manager40 Référentiel GIT CI / CD
Solution de supervision	Surveillance du cloud avec connecteur à la supervision OBS
Sauvegarde	GCBDR (snapshots inclus)
Solution de patchs pour OS	VM Manager pour la gestion des correctifs de l'OS Outil de patching OBS MA (BRAC) OBS OS factory
Solution d'antivirus	Outil OBS MA Sophos
Solution de journalisation	Journalisation de Google Cloud
Rétablissment du service	À partir de la sauvegarde lorsqu'elle existe A partir d'un script Terraform dans GIT, si nécessaire Idéalement, à partir d'une infrastructure à jour en tant que code avec CI/CD.
Connexion de l'admin	VPN vers OBS Bastion Zone - Connexion IaaS via CyberArk ou ORUN / Autre : Internet via API
Portail d'accès aux contrats MA, aux incidents et aux changements ITSM	OBS CloudStore (ECCS)

2.6 Prérequis généraux pour le RUN des services manages

Les prérequis suivants sont nécessaires pour tous les services managés :

- Le Client doit avoir défini une architecture valide. (OBS peut éventuellement fournir des services professionnels pour la définition de l'architecture).
- Le Client doit avoir un **abonnement valide à GCP, y compris un abonnement au plan de support GCP, et se procurer les ressources GCP et le plan de support GCP. OBS peut éventuellement fournir cet abonnement incluant le support GCP (réf. à l'offre Multi-Cloud Ready pour GCP), cependant, l'abonnement, les ressources IaaS, le support GCP ne font pas partie des Services Managés. Les services managés s'appuieront sur ce contrat de support pour faire remonter les incidents au CSP GCP.**
- La plateforme GCP pour le Client doit être urbanisée selon les meilleures pratiques de la landing zone GCP ou doit offrir des services comparables.
- OBS propose un RACI par défaut en fonction de la classe de transition et de la ressource gérée. En tant que prérequis au projet, OBS et le Client devront s'accorder sur le RACI.
- Accord sur l'outillage utilisé pour GIT, la chaîne CI / CD, la solution de Monitoring, Logging et Alerting.
- Des prérequis supplémentaires sont nécessaires lorsque la transition n'est pas entièrement sous la responsabilité d'OBS (pas de Full Build, voir le chapitre Build du document).

- Dans le cas d'un service managé complet, OBS utilise sa propre solution de Git, de chaîne CI / CD, de Monitoring, de Logging et d'Alertes.
- Dans le cas d'un service managé, OBS et le Client se mettent d'accord sur la solution Git, chaîne CI / CD, Monitoring, Logging et Alerting à utiliser. Par défaut, l'outillage est
    - Soit basé sur les outils GCP i.e. Google Cloud Deployment Manager, Google Cloud Monitoring
    - Soit basé sur l'outillage générique multicloud proposé par OBS, par exemple CaasCad (Prometheus, Grafana, …).

Cet outillage n'est pas inclus dans les unités d'œuvre Managed Applications et peut être acheté séparément dans le cadre de l'abonnement GCP ou en tant que proposition d'outillage multicloud faite par OBS.

2.7 Critères pour le RUN d'un service Cloud Natif

Les critères doivent être respectés et approuvés par le niveau 2 avant de transformer un composant cloud natif en un service managé actif (c'est-à-dire Run) par les opérations de niveau 2 / niveau 1. Le propriétaire du Build et du support de niveau 3 a la responsabilité de s'assurer que les critères sont respectés :

- L'architecture et le déploiement du service doivent être définis.
- Le service doit être déployé grâce à l'Infrastructure-as-Code et testé avant d'être transmis à l'équipe des opérations. Typiquement, le test réussi en Préproduction, avec un environnement de préproduction iso-production. Remarque : le IaC est nécessaire pour récupérer les services en cas de panne majeure.
- L'utilisation du service doit être expliquée à l'équipe d'exploitation.
- Les politiques de sécurité et le contrôle d'accès doivent avoir été configurés.
- L'accès doit avoir été configuré pour permettre aux équipes de niveau 2 d'OBS d'y accéder.
- Le service doit exporter les métriques nécessaires vers **Google Cloud Monitoring**.
- La sauvegarde des données doit être configurée dans **GCBDR** lorsque la sauvegarde est applicable.
- La reprise après sinistre doit être configurée le cas échéant.
- Les procédures de dépannage et de restauration du service doivent être fournies au niveau 2.
- Si une procédure nécessite des journaux ou un tableau de bord, ceux-ci doivent avoir été développés et déployés avant de passer à la phase d'exécution.
- Une procédure de dépannage sur incident ne doit pas durer plus de 15 minutes. Au-delà de cette durée, l'effort sera facturé sur la base du temps passé.

3 Le build des services & services managés sur GCP

3.1 Critères pour être qualifié comme modèle de "backend Build" pour une ressource

Le modèle de Build "Backend Build" pour une ressource est utilisé pour :

- Une ressource/service dans le périmètre de travail des services managés pour lequel l'infrastructure est déjà construite et déployée par le client en tirant parti de l'Infrastructure-as-Code.
- Et, pour lequel l'outillage GCP est entièrement configuré et opérationnel avant la transition sous la responsabilité du client. L'outillage utilisé sera :
  - Cloud Monitoring et Cloud logging pour la supervision avec des alertes appropriées définies.
  - GCBDR correctement configuré et fonctionnel
  - VM Manager et Patch Management configurés pour l'application de correctifs aux machines virtuelles.
  - Les procédures de remédiation et de dépannage pour les incidents connus sont définies et fournies.
  - Les procédures de récupération à utiliser sont définies et fournies par le client.

Et le client fournit la documentation, c'est-à-dire le schéma, le HLD et le DAT/LLD, l'architecture expliquant comment la disponibilité et le HA, la surveillance, les politiques de sécurité et le contrôle d'accès, la sauvegarde, la reprise après sinistre, la sécurité de base, le SLA sont réalisés.

L'effort de Build fourni par OBS dans le "backend Build" comprend l'intégration des alarmes de GCP Monitoring aux systèmes backend d'OBS, la capture des guides de procédures fournis par le client dans le référentiel de connaissances des opérations d'OBS, et la préparation des

opérations. Il s'agit également de préparer le backend administratif, l'ITSM d'OBS, le portail et la facturation pour les opérations.

3.2 Critères pour être qualifié comme modèle de " operations Build " pour une ressource

Le modèle "operations Build" pour une ressource est utilisé pour :

- Une ressource/service dans le cadre d'un service managé pour lequel l'infrastructure est déjà construite et déployée par le client en tirant parti de l'Infrastructure-as-Code.
- Le client fournit la documentation, c'est-à-dire le schéma, le HLD et le DAT/LLD, l'architecture expliquant comment la disponibilité et le HA, la surveillance, la sauvegarde, la reprise après sinistre, la sécurité de base et le SLA sont réalisés.
- Enfin, un accord est conclu entre le client et OBS pour utiliser les outils backend de GCP et d'OBS.

L'effort de Build fourni par OBS dans le "operations Build" comprend celui du "backend Build" plus la configuration et le déploiement de l'outillage de GCP grâce à l'Infrastructure as Code et du backend OBS, à savoir

- Cloud Monitoring et Cloud logging pour la supervision avec alertes
- Configuration et déploiement de GCBDR
- Configuration de Update Manager pour le Patching des VM
- Configuration de l'anti-virus pour les VM
- Utilisation des procédures standards de remédiation et de dépannage en cas d'incident connu pour le service natif du cloud.
- Utilisation de procédures de récupération standard pour le service natif du cloud.

Pour plus de détails sur les opérations par service, Référez-vous au **chapitre 11 : description détaillée par service cloud**.

3.3 Critères pour être qualifié comme modèle de "Full Build " pour une ressource

Le modèle "Full Build" pour une ressource est utilisé pour :

- Une ressource/service dans le cadre d'un service managé qui n'est pas encore construit et déployé.
- Le client fournit la documentation, c'est-à-dire le schéma, le HLD et le DAT/LLD, l'architecture expliquant comment la disponibilité et le HA, la surveillance, la sauvegarde, la reprise après sinistre, la sécurité de base et le SLA sont réalisés.
- Enfin, un accord est conclu entre le client et OBS pour utiliser les outils backend de GCP et d'OBS.

L'effort de Build fourni par OBS dans le "Full Build" comprend celui du "Backend Build" plus celui du "Operations Build" plus

- La configuration de la Landing Zone et de l'infrastructure de la ressource en utilisant Infrastructure as Code.

Pour plus de détails sur les opérations par service, référez-vous au **chapitre 11 : description détaillée par service cloud.**

Pour plus de détails sur l'Infrastructure as Code pour le modèle "Full Build", **Référez-vous au chapitre Méthodologie de l'Infrastructure as Code.**

3.4 Mesures à prendre en cas de non-respect des prérequis ou des critères

L'évaluation peut révéler que les critères ne sont pas remplis pour être qualifié comme un modèle de Build donné. Trois options sont alors possibles :

- L'étendue des travaux doit être réexaminée avec un modèle de Build plus approprié. Cela peut affecter la durée du projet et les efforts.
- Le client peut remédier aux critères manquants. Cela peut affecter la durée du projet et les efforts de gestion et de coordination du projet.
- Le client et OBS peuvent accepter de continuer avec certaines limitations dans les capacités et les responsabilités de gestion dues aux critères manquants.

Si le projet devait être retardé et si OBS devait dépenser plus de ressources que prévu en raison du non-respect des prérequis et des critères relevant de la responsabilité du client, OBS serait en droit de facturer le dépassement sur la base du temps et du matériel.

3.5 Modèle de facturation du Build

Service	Unité d'oeuvre
Gestion de projet	Temps et matériel
Service Implementation Coordination	Temps et matériel
Service Reliability Engineer	Temps et matériel
Architecte technique	Temps et matériel (si nécessaire pour la documentation)
Modèle Full Build – Unité 1ére ressource*	One Time Charge par ressource
Modèle Full Build – Unité ressource suivante du même type*	One Time Charge par ressource
Modèle Operations build model – Unité 1ére ressource*	One Time Charge par ressource
Modèle Operations Build - Unité ressource suivante du même type*	One Time Charge par ressource
Modèle Backend build - Unité 1ére ressource*	One Time Charge par ressource
Modèle Backend Build - Unité ressource suivante du même type*	One Time Charge par ressource

Unité de ressource*: Référez-vous au Chapitre 11 : description détaillée par service Cloud pour la définition de l'unité de ressource par service Cloud.

4 Directive sur la sécurité

Dans le cadre d'une offre gérée par Orange Business Services, nous considérons deux cas :

1. Le déploiement de charges de travail dans l'organisation GCP du client.
2. Le déploiement de workloads dans l'organisation GCP d'OBS

Nous avons alors 3 points à prendre en compte lors du déploiement de GCP Security Native Services :

1. Comment gérer l'identification et l'authentification pour OBS et pour le Client ?
   a. Identité de l'utilisateur,
   b. Identité du compte de service.
2. Comment donner des droits d'accès sur ces identités (identités de groupe, d'utilisateur et de compte de service) ?

3. Comment s'assurer que nous avons un audit périodique sur les 2 points précédents ? Effectuer une revue d'accès sur une base périodique, et donner un compte au client, ajoute de la valeur à l'offre.

Dans le cadre du projet NGOT (Next Generation Operation Tools), une infrastructure pérenne et automatique de fournisseur d'identité est actuellement développée, afin d'assurer l'attribution des droits aux utilisateurs de Managed Services.

Cette infrastructure IAM s'appuie sur les processus RH d'Orange Business Services pour gérer les utilisateurs OBS et leurs droits. Les identités et les droits des clients peuvent également être gérés à l'aide de l'infrastructure IAM d'OBS. Sinon, les identités et les droits cogérés peuvent être fédérés à partir de l'infrastructure IAM d'OBS, d'une part, et du fournisseur d'identité du client, d'autre part.

4.1 Hiérarchie des ressources

La zone d'atterrissage des clients comprend la hiérarchie des ressources dans l'organisation du client et dans l'organisation gcp.orange-business.com afin de pouvoir gérer le client séparément par OCB et par le client en cas de besoin. Les ressources de gcp.orange-business.com sont également utilisées pour automatiser le déploiement de l'infrastructure du client via IaC.

La hiérarchie des ressources restera simple et synthétique, nous donnant un bon aperçu des contrôles par types d'environnement, puis au dernier niveau, les projets eux-mêmes. **Les 4 niveaux en gras** ci-dessous sont la conception standard, par rapport à la norme de sécurité des groupes d'Orange. En cas de niveaux supplémentaires, les règles s'appliquent à tous les niveaux. Les entités peuvent avoir plus de 4 niveaux.

1. Organisation
2. **Dossier Client (dans le cas d'une gestion complète) (Dossier de 1er niveau)**
3. **Dossier de gestion de la zone d'atterrissage (Dossier de 2ème niveau)**
4. **Dossiers d'environnement et dossiers communs (Dossier de 3ème niveau)**
5. **Dossier Entité (Dossier de 4ème niveau)**
6. Projets (Dossier de 5ème niveau)

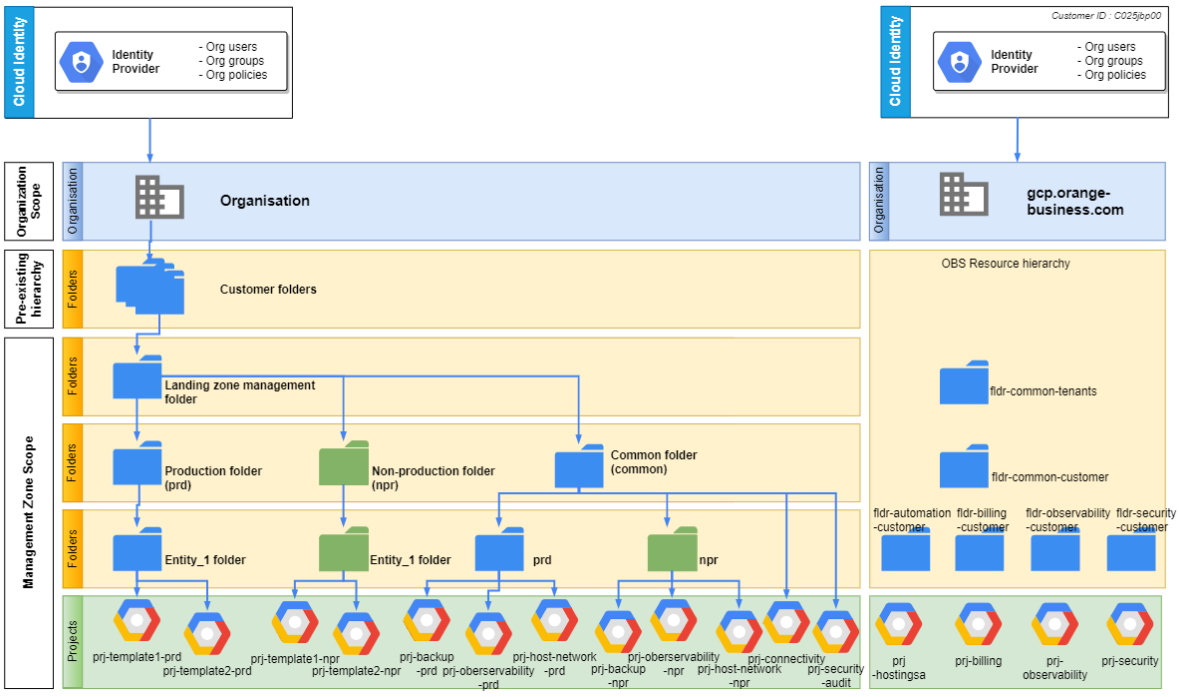

4.2 IAM policy

La gestion des identités et des accès (IAM) **couvre les produits, les processus et les politiques utilisés pour gérer les identités des utilisateurs et réguler leur accès au sein d'une organisation**.

La mise en œuvre de la politique IAM suit les trois principes ci-dessous :

1. Le moindre privilège : s'assurer qu'un utilisateur dispose des permissions nécessaires et suffisantes en fonction de son rôle.
2. Isolation : s'assurer qu'il n'y a pas d'accès ou même de visibilité possible d'un fournisseur d'identité à un autre.
3. Autonomie : limiter la perte de temps liée à l'acquisition de permissions

4.2.1 Groupes (niveau organisation)

- ✓ Chaque règle de politique IAM est prise en charge par des groupes définis au niveau auquel la politique est pertinente. Le responsable de chaque groupe est chargé de déterminer et de remplir les membres.
- ✓ Le nom du groupe est préfixé en fonction du niveau auquel appartient la règle de politique IAM.

Les types de groupes suivants sont actuellement définis :

4 groupes privilégiés

- admins (ressources admins) : qui peuvent créer des projets et des dossiers
- security-admins : qui peut donner accès en donnant des rôles IAM
- billing-account-users : qui peut attribuer un compte de facturation à un projet
- administrateurs de réseau

groupes les moins privilégiés

- viewers : ceux qui peuvent voir les ressources
- dev-ops : qui peut déployer en tant que code. Principalement qui peuvent se faire passer pour des comptes de services dédiés
- dev-apps : utilisateurs avec des droits limités
- security-reviewers : qui peuvent revoir la sécurité du cloud, principalement voir les paramètres des rôles IAM

Groupes IAM créés par OBS IaC pour gérer efficacement les privilèges

Groupes privilégiés	Groupes Client		Groupes OBS
Oui	grp-gcp-iam-ocb-<tenant-name>-**tenant-resources-admins**@<customer-domain>		grp-gcp-iam-ocb-<tenant-name>-**tenant-resources-admins**@gcp.orange-business.com
Oui	grp-gcp-iam-ocb-<tenant-name>-**tenant-security-admins**@<customer-domain>		grp-gcp-iam-ocb-<tenant-name>-**tenant-security-admins**@gcp.orange-business.com
Oui	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

Groupes privilégiés	Groupes Client		Groupes OBS
Oui	grp-gcp-iam-ocb-<tenant-name>-**tenant-dev-ops**@<customer-domain>		grp-gcp-iam-ocb-<tenant-name>-**tenant-dev-ops**@gcp.orange-business.com
Non	grp-gcp-iam-ocb-<tenant-name>-**tenant-viewers**@<customer-domain>		grp-gcp-iam-ocb-<tenant-name>-**tenant-viewers**@gcp.orange-business.com
Non	grp-gcp-iam-ocb-<tenant-name>-**tenant-dev-apps**@<customer-domain>		grp-gcp-iam-ocb-<tenant-name>-**tenant-dev-apps**@gcp.orange-business.com
Non	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 des groupes

Permissions des groupes IAM privilégiés et des comptes de service

ressources limitées <customer-domain> (customer organisation)	membre <ocb> is here the entity reseller	définition du rôle	rôle	état de mise en œuvre
organisation	grp-gcp-iam-ocb-<tenant-name>-**tenant-admins**@<customer-domain>	Les membres sont responsables de l'organisation de la structure des ressources utilisées par l'organisation.	roles/cloudsupport.admin roles/resourcemanager.folderCreator roles/resourcemanager.projectCreator roles/securitycenter.admin	à jour
organisation	sa-node-creator-<CUSTOMER>-prd	Ce compte de service peut créer des nœuds (dossiers ou projets) dans l'organisation <CUSTOMER>.	roles/resourcemanager.folderCreator roles/resourcemanager.projectCreator	à jour
organisation	grp-gcp-iam-ocb-<tenant-name>-**tenant-security-admins**@<customer-domain>	Les membres sont chargés d'établir et de gérer les politiques de sécurité pour l'ensemble de	roles/compute.viewer roles/container.viewer roles/resourcemanager.organizationAdmin roles/iam.organizationRoleViewer roles/iam.securityReviewer roles/logging.configWriter	à jour

		l'organisation, y compris la gestion des accès.	roles/logging.privateLogViewer roles/orgpolicy.policyAdmin roles/orgpolicy.policyViewer roles/resourcemanager.folderIamAdmin roles/resourcemanager.tagAdmin roles/resourcemanager.tagUser roles/securitycenter.admin	
organisation	grp-gcp-iam-ocb-<tenant-name>-**tenant-billing-account-users**@<customer-domain>	Les membres peuvent lier le compte de sous-facturation du client à un projet dans l'organisation du client.	roles/billing.projectManager	à jour
<CUSTOMER> compte de sous-facturation	grp-gcp-iam-ocb-<tenant-name>-**tenant-billing-account-users**@<customer-domain>	Les membres peuvent utiliser le compte de sous-facturation Client	roles/billing.user	à jour
<CUSTOMER> compte de sous-facturation	sa-node-creator-<tenant-name>-prd	Ce compte de service peut utiliser le compte de sous-facturation <CUSTOMER>.	roles/billing.user	à jour
organisation	grp-gcp-iam-ocb-<tenant-name>-tenant-viewers@<customer-domain>	Peut visualiser l'organisation	roles/resourcemanager.organizationViewer	à jour
organisation		Peut visualiser le dossier <TENANT> et ses sous-dossiers	roles/resourcemanager.folderViewer	à jour
<CUSTOMER> compte de sous-facturation	grp-gcp-iam-ocb-<tenant-name>-tenant-viewers@<customer-domain>	Les membres peuvent voir toutes les sous-factures de <CUSTOMER>.	roles/billing.viewer	à ajouter
organisation	grp-gcp-iam-ocb-<tenant-name>-tenant-dev-apps@<customer-domain>	à définir	à définir	non appliqué

organisation	grp-gcp-iam-ocb-<tenant-name>-tenant-security-reviewers@<customer-domain>	Les membres font partie de l'équipe de sécurité responsable de l'examen de la sécurité du cloud dans le(s) dossier(s) OCB.	roles/iam.securityReviewer roles/cloudasset.viewer	à jour

4.3 Gestion des secrets pour le déploiement

Pour des raisons de fiabilité, de performance et de capitalisation, le déploiement est automatisé et mis en œuvre via Terraform et Ansible.

Le déploiement automatisé nécessite l'utilisation de comptes de service dédiés, dotés de privilèges particuliers. L'accès à ces comptes est protégé par des secrets spécifiques, dont l'utilisation doit être limitée aux agents composant les outils de déploiement automatisé.

4.3.1 Standard OBS deployment tools

Les outils suivants sont utilisés en complément de Terraform et Ansible pour le déploiement de l'automatisation :

Outils	Description	Règles	Référence externe
GitLab	Gestion du code Exécution du processus CI/CD (runners GitLab)	Meilleures pratiques, Gitflow et convention de nommage : https://ocbconfluence.equant.com/display/practicegcp/GIT	https://sourcehub.orange-business.com/
Hashicorp Vault	Gestion des secrets	Documentation sur les meilleures pratiques	https://www.hashicorp.com/products/vault
Artifactory	Stockage des artefacts	Documentation sur les meilleures pratiques	https://multirepo.orangeapplicationsforbusiness.com/ui/

Comptes de services privilégiés requis

ressources limitées <customer-domain> (customer organisation)	Member <ocb> is here the entity reseller	définition du rôle	rôles
Organisation	sa-tenant-deployer-groups-mgmt@deployment-core-3c64.iam.gserviceaccount.com		Visualisateur d'organisation

ressources limitées <customer-domain> (customer organisation)	Member <ocb> is here the entity reseller	définition du rôle	rôles
identité cloud client	sa-tenant-deployer-groups-mgmt@deployment-core-3c64.iam.gserviceaccount.com	Ce compte de service peut gérer les groupes de l'organisation <CUSTOMER>.	Groupes d'administrateur
Organisation	sa-tenant-deployer-ocb-dev@deployment-ocb-dev-3ea1.iam.gserviceaccount.com	Ce compte de service peut • - modifier l'IAM sur l'organisation entière • - créer et supprimer des projets	Administrateur de l'organisation Créateur de projet Suppression du projet

4.3.2 Outils de déploiement définis par le client

Selon la classe de service achetée par le client, les scripts de déploiement automatisé peuvent être adaptés aux besoins du client. Ces solutions sont conçues et mises en œuvre dans le cadre d'ateliers dédiés avec Orange Business Services.

La conception des processus et des scripts de déploiement automatisé englobe le sujet de la gestion des secrets protégeant les comptes de services dédiés.

4.4 Examen des comptes et des accès

La revue des comptes et des droits d'accès est une composante essentielle de notre politique IAM, étroitement liée à la gestion du cycle de vie des identités et au provisionnement des comptes et des droits. Elle consiste à s'assurer que les droits d'accès des utilisateurs du système d'information sont conformes à ce qu'ils devraient être, et à les certifier, ou - si nécessaire - à effectuer des opérations de remédiation en cas de non-conformité à la politique d'autorisation du client.

Cette activité s'inscrit donc dans une logique de gouvernance et de contrôle des autorisations, afin d'apporter les garanties de conformité attendues.

La revue des comptes et des accès est réalisée tous les six mois, sous la responsabilité du responsable de la sécurité des entreprises. Si le service vendu ne comprend pas de service BSO, alors le Service Delivery Manager est responsable de la revue des comptes et des accès.

4.5 Gestion des correctifs

Les charges utiles standard s'exécutant sur les VM traditionnelles sont basées sur des images d'OS renforcées. Ils font l'objet d'une gestion centralisée des correctifs. Lorsque la charge utile est conteneurisée, l'OBS peut effectuer l'équivalent de la gestion des correctifs, à condition qu'il ait la responsabilité et le contrôle de la construction de l'image de base.

5 Responsabilités et obligations détaillées

Les tableaux suivants décrivent les responsabilités standards par défaut entre OBS et le client en fonction **du modèle de Build**.

Les tableaux suivants décrivent les responsabilités standard par défaut entre OBS et le client en fonction des classes de service. Elles peuvent être modifiées avec le consentement du mutuel en fonction du projet.

- R signifie Responsable (réalisateur)
- A pour Accountable (autorités ou responsable)
- C signifie Contributeur
- I signifie Informé

5.1 RACI pour Managed OS

Implementation du Service	OBS	Client	OBS	Client	OBS	Client
Mise en œuvre de l'infrastructure de l'OS du serveur	Full build		Operations build		Backend build	
Déploiement de l'infrastructure	R, A	I	I	R, A	I	R, A
Déploiement des composants du réseau local (LAN)	R, A	I	I	R, A	I	R, A
Déploiement des services DNS et NTP	R, A	I	R, A	I	I	R, A
Outils de sauvegarde pour les opérations (GCBDR & GCP Snapshots)	R, A	I	R, A	I	I	R, A
Déploiement de la solution de Patching de l'OS (Google VM Manager) or OBS Patch management	R, A	I	R, A	I	I	R, A
Déploiement de la solution Antivirus	R, A	I	R, A	I	SoW	SoW
Déploiement de la solution de supervision (Google Cloud Monitoring)	R, A	I	R, A	I	I	R, A
Déploiement de la solution de supervision (Google Cloud Monitoring)	R, A	I	R, A	I	I	R, A
Déploiement des groupes de sécurité et des règles de pare-feu	R, A	I	SoW	SoW	I	R, A
Procédure de récupération (Infra as Code, restauration, autre…)	R, A	I	I	R, A	I	R, A
Test et validation de l'implémentation de l'infrastructure	R	A	I	R, A	I	R, A
Test et validation de l'implémentation des outils GCP et de la gestion du cycle de vie.	R	A	R	A	I	R, A
Mise en œuvre de l'OS du serveur						
Évaluation ou déploiement du système d'exploitation	R, A	I	R, A	I	I	R, A
Déploiement de nouveaux paquets	R, A	I	R, A	I	R, A	I
Test et validation de l'implémentation du système d'exploitation pour les nouveaux packages	R, A	I	R, A	I	R, A	I
Documentation sur la mise en œuvre des services						
Conception, architecture et low-level design pour l'infrastructure	I	R, A	I	R, A	I	R, A
Documentation de mise en œuvre et d'exploitation de l'infrastructure	R, A	I	I	R, A	I	R, A
Conception et low-level design pour l'outillage (GCP)	R, A	I	R, A	I	I	R, A
Documentation de mise en œuvre et d'exploitation de l'outillage (GCP)	R, A	I	R, A	I	I	R, A

5.2 RACI pour Database as a Service

Implémentation du service	OBS	Client	OBS	Client	OBS	Client
Database aaS services conception et implémentation	Full Build		Operations Build		Backend Build	
Maintenance du référentiel d'architecture de l'infrastructure	R, A	I	I	R, A	I	R, A
Maintenance du référentiel de configuration de l'outillage	R, A	I	R, A	I	I	R, A
Déploiement de l'infrastructure	R, A	I	I	R, A	I	R, A
Déploiement de la solution de supervision (Google Cloud Monitoring)	R, A	I	R, A	I	I	R, A
Déploiement de la solution de Logging (Google Cloud Logging) (optionnel)	R, A	I	R, A	I	I	R, A
Déploiement de la solution de sauvegarde (GCBDR, Snapshot)	R, A	C, I	R, A	C, I	I	R, A
Procédure de récupération de l'infrastructure à partir du référentiel (Infra as code, restauration à partir de la sauvegarde, autre...)	R, A	C, I	I	R, A	I	R, A
Procédure de récupération de l'outillage à partir du référentiel (Infra as code, restauration, autre...)	R, A	C, I	R, A	C, I	I	R, A
Test et validation de l'implémentation de l'infrastructure	R, A	I	I	R, A	I	R, A
Test et validation de l'implémentation des outils et de la gestion du cycle de vie.	R, A	I	R, A	C, I	I	R, A
Exécution de scripts fournis par le client sur l'instance de la base de données	R	A, I	R	A, I	R	A, I
Exécution du script OBS sur l'instance de la base de données	R, A	C, I	R, A	C, I	R, A	C, I
Documentation sur la mise en œuvre des services						
Conception, architecture et low-level design pour l'infrastructure	C, I	R, A	I	R, A	I	R, A
Documentation de mise en œuvre et d'exploitation de l'infrastructure	R, A	C, I	I	R, A	I	R, A
Conception et low-level design pour l'outillage (GCP)	R, A	C, I	R, A	C, I	I	R, A
Documentation de mise en œuvre et d'exploitation de l'outillage (GCP)	R, A	C, I	R, A	C, I	I	R, A

Opérations des services	OBS	Client	OBS	Client	OBS	Client
Opérations Database aaS services	Full Build		Operations Build		Backend Build	
Monitoring via Google Cloud Monitoring	R	I	R	I	R*	I
Investigation via Google Cloud Monitoring & Google Cloud Logging	R, A	C,I	R, A	C,I	R*	A
Restauration à partir d'Infra comme Code et sauvegarde	R, A	C,I	R, A	C,I	R*	A
Modification de la capacité de l'instance de la base de données	R, A	C,I	C, I	R, A	C, I	R, A
Opérations ITSM						
Gestion du changement	R	A	R	A	R	A
Gestion des incidents	R, A	R**,I	R, A	R**,I	R, A	R**,I
Gestion des événements	R, A	I	R, A	I	R, A	I
Gestion de la sécurité de base	R	A	SoW	SoW	SoW	SoW
Gestion des configurations	R, A	C, I	R	A	R	A
Gestion des rapports via le service SDM	R, A	C, I	R, A	C, I	R, A	C, I
Gestion de la facturation	R, A	I	R, A	I	R, A	I

R* : dans les limites de l'outillage fourni par le client
R** : dans le modèle de co-management, le client peut avoir des responsabilités conjointes liées à l'activité et aux incidents.

5.3 RACI pour les Services Native managés

Implémentation du service	**OBS**	**Client**	**OBS**	**Client**	**OBS**	**Client**
Implémentation des services natifs	Full Build		Operations Build		Backend Build	
Déploiement de l'infrastructure	R, A	I	I	R, A	I	R, A
Outils de sauvegarde pour les opérations (GCBDR) (1)	R, A	I	R, A	I	I	R, A
Déploiement de la solution de supervision (Google Cloud Monitoring) (1)	R, A	I	R, A	I	I	R, A
Déploiement de la solution de journalisation (Google Cloud Logging) optional (1)	R, A	I	R, A	I	I	R, A
Déploiement des groupes de sécurité et des règles de pare-feu.	R, A	I	SoW	SoW	I	R, A
Procédure de récupération (Infra as Code, restauration, autre...)	R, A	I	I	R, A	I	R, A
Test et validation de l'implémentation de l'infrastructure	R	A	I	R, A	I	R, A
Test et validation de l'implémentation des outils GCP	R	A	R	A	I	R, A
Packages						
Deployment of new packages (1)	I	R, A	I	R, A	I	R, A
Documentation de l'implémentation des services						
Conception, architecture et low-level design pour l'infrastructure	C, I	R, A	I	R, A	I	R, A
Documentation de mise en œuvre et d'exploitation de l'infrastructure	R, A	I	I	R, A	I	R, A
Conception et low-level design pour l'outillage (GCP)	R, A	I	R, A	I	I	R, A
Documentation de mise en œuvre et d'exploitation de l'outillage (GCP)	R, A	I	R, A	I	I	R, A

Opérations des services	**OBS**	**Client**	**OBS**	**Client**	**OBS**	**Client**
Opérations des services natifs	Full Build		Operations build		Backend build	
Monitoring (1)	R, A	I	R, A	I	R*	A
Sauvegarde (1)	R	A	R	A	R*	A
Restauration de l'Infra comme Code et sauvegarde (1)	R, A	C, I	R, A	C, I	R*	A
Groupes de sécurité, paramétrage des règles du pare-feu	R	A	SoW	SoW	I	R, A
Opérations ITSM						
Gestion du changement	R	A	R	A	R*	A
Gestion des incidents	R, A	R**, I	R, A	R**, I	R*, A	R**, I
Gestion des événements	R, A	I	R, A	I	R*	A
Gestion de la sécurité de base	R	A	SoW	SoW	SoW	SoW
Gestion des rapports via le service SDM	R, A	I	R, A	I	R, A	I
Gestion de la facturation	R, A	I	R, A	I	R, A	I

R* : dans la limite de l'outillage fourni par le client.

R** : dans le modèle de co-management, le client peut avoir des responsabilités conjointes liées à l'activité et à l'incident.
(1) Si applicable selon la description détaillée par 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

Example : For a Customer we have configured :

- Monitoring CloudSQL Postgres :
    - Utilization CPU :
        - Warning at 80% : p4
        - Critical at 90% : p2
    - Utilization Memory :
        - Warning at 80% : p4
        - Critical at 90% : p2
    - Utilization Bandwidth Disk :
        - Warning at 80% : p4
        - Critical at 90% : p2
    - Different State of « RUNNING »: Critical for severity P1

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.

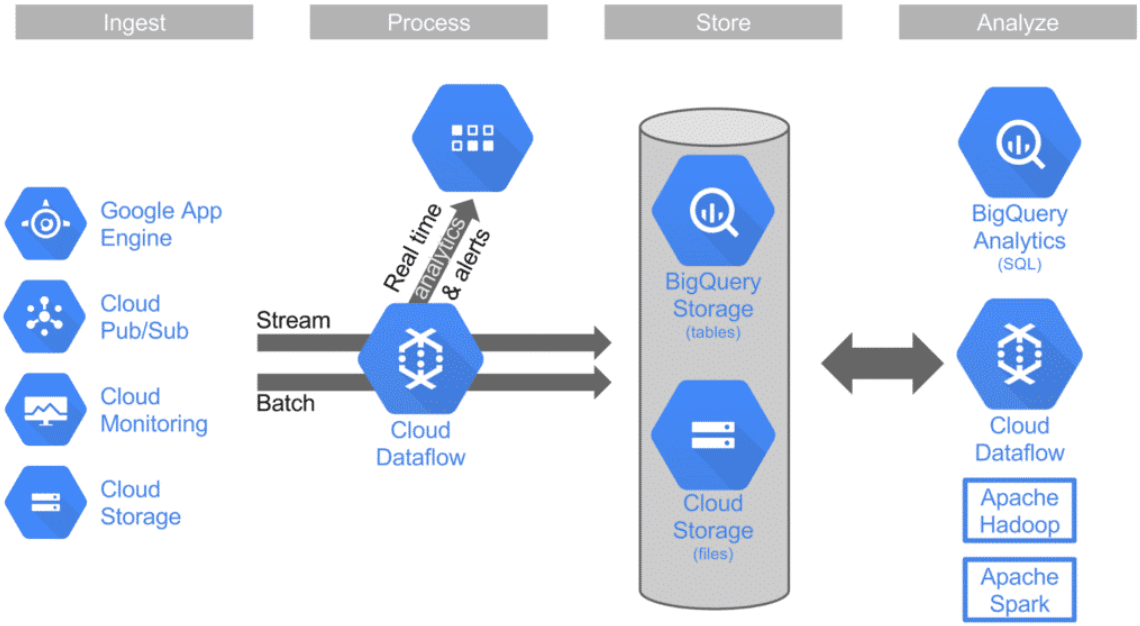

6.16.2 Build to run service included in the OTC

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

Alerts observed:

Toute erreur dans les logs entraîne une alarme « Critical », de sévérité p2 [Tickets P2]

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