Business

Service Description
Flexible Engine

Table of contents

1 Definitions

Complementary to the definitions as per General Terms and Cloud Specific Terms, the following specific definitions shall apply with respect to this Service Description.

Availability Zone refers to a separate data center sufficiently distant from the others, if any, in the same Region to allow the implementation of a local resilience. Availability Zones in each Region are listed in the Service Description.

Downtime refers to the period(s) during which an incident causes a significant malfunction of the Service or Feature concerned, affecting all Users. Calculating the duration of the unavailability obeys specific criteria for each Service or Feature.

Flexible Engine Console refers to the web interface allowing to administrate the Flexible Engine Services.

Multi-AZ Region refers to a Region offering several Availability Zones.

Region refers to a geographical area where the Service is available on one or several Availability Zone(s). The Regions are listed in the Service Description.

Service refers to the service “Flexible Engine” provided for one Tenant. Each Tenant constitutes a separate Service.

Tenant refers to a virtual private pool of resources on “Flexible Engine” cloud, only accessible to Users which are authenticated by login and password. Creation, deletion, modification and listing of these resources and associated Features may be performed by those Users only.

Virtual Machine (VM) refers to a software computer which, like a physical computer, runs an operating system and applications. The virtual machine consists of a set of specification and configuration files. It is supported by the physical resources of a host. Each virtual machine has virtual devices which provide the same functions as hardware.

2 Purpose of the document

The purpose of this service description is to define the conditions under which the Provider provides the "Flexible Engine" service (hereafter the "Service") to the Customer.
The present description is attached to the Cloud Specific Terms.

3 Overview of the Service

3.1 Overall description

Flexible Engine is a global Infrastructure as a Service solution. Flexible Engine offers a rich portfolio of cloud services available as a services represented in the following diagram. The roadmap is highly evolutive.

3.2 Geographical footprint

Flexible Engine services spans across different Regions of the world enabling users to deploy workloads globally. Flexible Engine Console allows for central control and deployment of the workloads. When setting up his Service, the Customer determines the Region(s) in which his data will be processed and stored. The Provider does not move the Customer's data to a Region other than the one(s) chosen by the Customer.

The solution offers Regions operated by the Provider and Regions operated by the Provider partners. The regions operated by the Provider are as follows:

- Paris (eu-west-0) which is made of 3 datacenters in Paris’s Area in Western Europe
- Amsterdam (eu-west-1) which is made of 3 datacenters in Amsterdam in Western Europe
- SAP Hana Paris (eu-westxvp-28) which includes a datacenter in the Paris area

Partner Regions are visible on the Customer Space.

The availability of some Features may be different according to the Regions and is indicated in the Price List and/or the Flexible Engine Console and/or on partner websites. The SAP Hana Paris Region is only accessible for SAP Hana projects in managed mode (Managed Applications offer)

3.2.1 Regions and Availability Zones (AZ)

Flexible Engine services are based on an architecture designed for resilience.

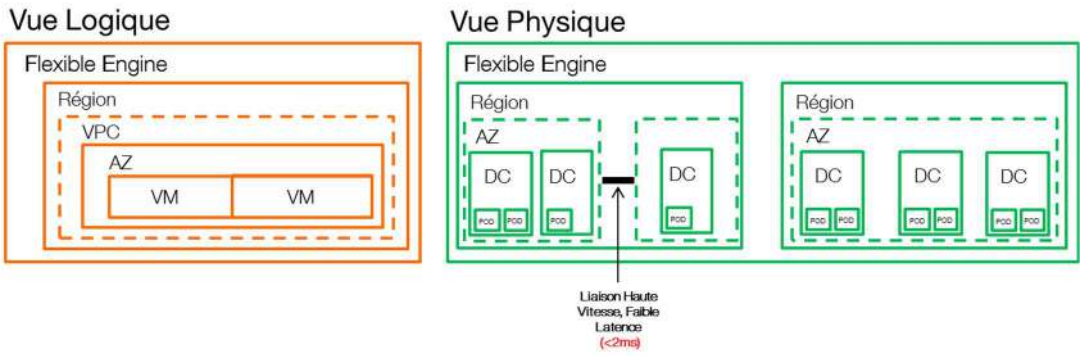

Figure n° 1: *Resilience and protection against major disasters*

Main concepts and principles:

- At the world level: separate Regions allow us to address the distributed use cases or to set up a global resilience.
- At a Region level: separate Areas of Availability (AZ). Each Availability Zone is implemented in a separate and distant Data Center to implement a local resilience. The Data Centers are close enough to implement load balances (less than 2 ms latency between 2 Azs in a Region).
- At the cloud server instance level, anti-affinity can be used to ensure that instances are deployed on different hypervisors. At the level of an AZ: different service areas isolated from each other, implementing all redundant infrastructure components:
  - Local internet access by different providers
  - Redundant security and interconnection equipment
  - IaaS infrastructure based on equipment offering High Availability

Orange has designed the resilience of its services by drastically dissociating:

- The traffic collection and security management layers via infrastructure-independent points of presence (POPs), located in each zone of availability
- The IaaS catalogue service delivery layer via the Data Center
- The interconnection between the two layers is done locally in each of the Availability Zone.

3.3 HDS certification

Flexible Engine used in conjunction with Business support solution or higher is certified Hébergement de Données de Santé (HDS) for healthcare service providers in France to manage sensitive health data of patients.

4 Terms of use

4.1 Prices

The prices for the Service are revisable as per General Terms, except for reserved instances whose prices are firm and settled for the duration of the commitment at the time of the reservation.

The prices of the Service are subject to review under the conditions set forth in the General Conditions and the Cloud Specific Terms and may be updated monthly. The new prices apply to current Contracts. The Customer will be informed of the new prices by publication on the User Interfaces or by any other means, no later than the date on which the new prices come into force. In the event of a price increase for an existing Functionality, the Customer will be informed by e-mail or by any other means no later than 30 days before the new prices come into force. The prices in force on the Activation Date may differ from those communicated at the time of subscription.

The prices are defined per Region. The Customer identifies in the Order a primary Region which shall bear the fees that are not connected to a specific Region.

For partner Regions, only the catalogue prices in US dollars on the partner website are valid.

4.2 Licenses

The Customer undertakes to use Software, including operating systems, in compliance with the article “Intellectual property” of the General Terms.

4.2.1 Microsoft products

The Customer may either subscribe to Microsoft Software licenses from the Provider in rental mode or bring licenses subscribed by him directly to Microsoft or a third-party reseller in mobility mode, according to the terms of use applicable to each Software, available at the following address:

https://www.microsoft.com/en-us/licensing/product-licensing/products

The Customer's use of Microsoft Software must comply with the terms of use associated with Microsoft's SPLA (Service Provider License Agreement). The Customer is the sole responsible of non-conformities in using Microsoft software and services and he is liable of the consequences decided by Microsoft in such cases.

4.2.1.1 Rental mode

The Microsoft licenses offered by the Provider are in rental mode, the Customer must not use the corresponding licenses for any purpose other than the Service subscribed to from the Provider.

4.2.1.2 Mobility mode

Microsoft license mobility, for previously acquired software, is possible in accordance with the "License Mobility" or "Qualified Multitenant Hoster" (QMTH) add-on of the SPLA contract, depending on the Software concerned.

Among other conditions, the Customer is responsible for the following operations:

- to have subscribed with Microsoft, when required by Microsoft, the "Software Assurance" (SA) which is an additional license to allow its mobility;
- for Mobility License, declare mobility to Microsoft, indicating ORANGE's references as a mobility partner, via a specific form published by Microsoft and provided to the Customer by the Provider at the Customer's request;
- for QMTH, declare to the Provider the number of Users for each Software concerned.

Microsoft licenses from a Customer’s SPLA or SPLA Academics contract are eligible to mobility. In order to enable this capability, an “Outsourcing Company Agreement” contract have to be signed between the Customer and the Provider as stated in §7 of the Customer’s SPLA contract.

4.3 Reversibility

Within the context of the Flexible Engine solution, the Customer may perform backups of its virtual ECS machines and its EVS volumes. To the extent that it is a usage online platform, the Customer is free to retrieve by using the APIs of the Flexible Engine nominally through the connection that it uses, for example the Internet or a Direct Connection, backups thus made as well as the data stored on its OBS Storage Object. The Provider is not involved.

Optionally, the performance of Services can be envisaged subject to a quotation.

4.4 HDS Data Retrieval and Deletion

HDS data from Backups of ECS virtual machines and EBS volumes can be recovered by the Customer in autonomous mode or accompanied mode.

Autonomous mode:

Insofar as it is an online platform for use, the Customer is autonomous in retrieving, nominally using the APIs of Flexible Engine HDS through the connection he uses, for example the Internet or a Direct Connection, the backups thus made as well as the data stored on its Object Storage. The Provider does not intervene. The data belonging to the Customer is automatically deleted from the Datacenters within a maximum period of two (2) months from the contractual end date of the Service. The Customer may decide to manually delete all of its data before the end of the service from the Flexible Engine HDS Console.

Accompanied mode:

Optionally, services can be considered on quotation.

Once transferred, the data belonging to the Customer, as well as all of their backups, are deleted from the Datacenter within a maximum period of two (2) months from the signing of the acceptance report for the reversibility plan and in any condition. no later than two (2) months from the contractual end date of the Service. At the Customer's request, the Supplier may certify in a written document that the Customer's data has been deleted.

5 Access to the Service

5.1 Portals

Once the Order accepted by the Provider, the Customer will receive a confirmation email of its registration and request for initialization of its password to access the Flexible Engine portal called Flexible Engine's Cloud Customer Space, to the Flexible Engine Console and to administer the services in the Tenant created for it.

The Customer can invite Users with the rights to use Services into its Tenants. The Users to whom the Customer has given the corresponding rights may themselves invite other Users.

5.1.1 The Provider cloud Orange Business portal

The Provider cloud Orange Business portal is the live website where our services offers are presented. It includes as well events, references and partner pages. URL http://cloud.orange-business.com/

5.1.2 Cloud Customer Space

Flexible Engine's Cloud Customer Space is a secured environment from which the user can manage is account. It allows for

- **Dashboard**: this section allows to view general account information, access the console
- **Requests**: this section allows to view the history of requests made on the Cloud Customer Space, and their status
- **Catalog**: this section allows to access a catalogue of Flexible Engine services and to order them: order Reserved Instances, reserve dedicated resources, rename a contract, activate Partner Regions, etc
- **Rights**: this section allows to manage the rights of Flexible Engine users on the Cloud Customer Space and to give them access to the Technical Console
- **Invoices**: Allows to consult all online invoices and other related files: detailed invoices, invoice simulation during the month, etc. This section also offers budget management (alerts)
- **Subscriptions**: from the dashboard, this section allows to consult the list of Reserved Instances and Flexible Reserved Instances with their status and end date of commitment
- **Need help?**: from the dashboard this section allows to access all the online help and to create support requests/tickets
- **Partner**: this section can only be accessed with a "partner login" issued to the Client by the Provider at its request, it allows :
    - the Customer to create additional Tenants - the uses of these Tenants are the Customer's uses, they are aggregated to his invoice
    - the Customer to manage its various Tenants on the Cloud Customer Space from a single login - access to the various sections mentioned above,

  This section does not allow the management of Flexible Engine user rights or access to the Flexible Engine Technical Console.

The Cloud Customer Space also provides an API service based on the rights of the Cloud Customer Space users to the different Flexible Engine tenants.

This API allows the Customer, if the Cloud Customer Space login used has sufficient rights to :

- access the billing files of the Tenants,
- create additional Flexible Engine Tenants,
- give access to the Flexible Engine Tenants to existing or new users

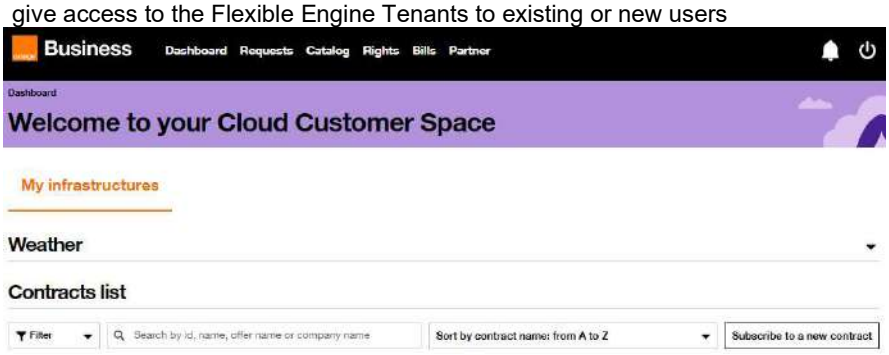

Figure n° 2: *Cloud Customer Space*

5.1.3 Flexible Engine Console

The Flexible Engine Console is the main web portal to manage all the User's Flexible Engine services.

Using the Flexible Engine Console, the user can create, configure, monitor and managed all on-line services of Flexible Engine. It is accessible from: https://console.prod-cloud-ocb.orange-business.com/console/#/home.The Flexible Engine Console connects to all international Regions of Flexible Engine global cloud.

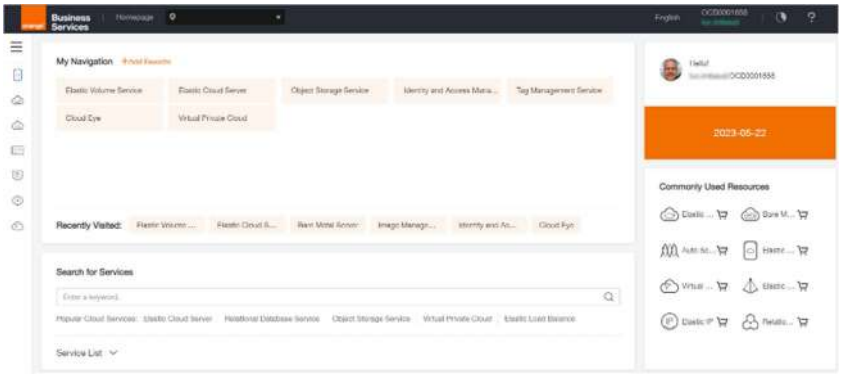

Figure n° 3: *Flexible Engine Console*

5.1.4 Partner websites

For partner Regions, the Customer has access to the partner site, by means of a specific login/password for functionalities requiring a connection, giving him access to:

- technical information on the solution available in the partner Regions
- the public prices applicable in the context of the Service
- the partner customer space, only for usage reports (usage reports are not invoices)
- the partner technical console, for the configuration of the Service

The other elements present on the partner site are not applicable to the Service, in particular:

- the contractual documents submitted by the partner.
- the Order, renewal or termination journey offered by the partner (outside the technical console). The use of this journey has no impact on the Service and does not engage the Provider.
- the "prepaid" invoicing method proposed by the partner.
- partner promotions.
- access to partner support. The terms and conditions for access to the Service's support are exclusively those described in Chapter 7 of this document.

5.2 Network

Connection to the Service is carried out via the Internet network.

The Flexible Engine services do not include the network interconnection between the Customer's company and the Provider' infrastructure. Depending on the Customer's needs, network interconnection solutions can be subscribed to separately from the Provider.

6 Content of the Service

6.1 Flexible Engine Compute services

6.1.1 Elastic Cloud Server

Elastic Cloud Servers (ECS) are Virtual Machines (VM). This service, available in all Availability Zones, offers self-service virtual machines. The User can start, stop, resize VMs using the Flexible Engine Console or using the ECS API.

Cloud Servers use a system disk based on the Elastic Volume Service (EVS), and for some flavors a local disk is also included. The ECS is invoiced:

- when it is turned on for flavors that do not include a local disk,
- including when it is turned off and until it is completely deleted for flavors including a local disk.

The technical specifications of each instance are available in the Flexible Engine Console.

The invoicing principles are specified in the Price List.

6.1.2 Baremetals servers

A BMS (Bare Metal Server) is a physical server that can be subscribed and launched from the Flexible Engine console. This server is entirely dedicated to the Client in order to install key applications or to optimize Oracle database licenses for example. The client can combine this physical server with other Flexible Engine services such as a virtual private cloud (VPC), OS images (IMS), storage or network. With a dual network card (Multi-NIC), the BMS can connect to two networks.

6.1.3 Reserved instances

"Reserved Instances" Feature allows the User to benefit from specific rates for the ECS service, depending on the duration of its commitment (1, 2, 3 or 5 years) and whether or not one-time upfront fees are included. With this Feature, ECS, Workspace, BMS, or CSS are no longer billed on a per use basis but on a monthly subscription basis.

"Reserved Instances" are a pricing mechanism that allows you to benefit from improved rates than those applied on a pay-per-use basis. They apply to resources provisioned via the user console or APIs. The subscription to "Reserved Instances" is done via the Cloud Customer Space. They do not guarantee the availability of associated resources at the time of subscription. Before subscribing to a Reserved Instance, it is essential to check the availability of the cloud resource in the Flexible Engine Console that will physically allocate them.

Reserved Instances are available in two models:

- No upfront fee: the user pays a set amount every month, regardless of usage.
- With upfront fee: the user will be charged an upfront fee (which is a portion of the total cost of the reserved instance during the subscription period) in addition to a defined monthly amount, regardless of usage.

The Customer may have both standard ECSs and reserved instances within the same tenant. The reserved instances are subscribed for a given flavor and region. Each month, the ECS corresponding to the reserved flavor with the highest consumption in the applicable region will be covered by the subscription, while the others will be billed on a per-use basis. The reserved instance applies from the first day of the calendar month following the subscription. Thus, the first month of use will be billed on a pay-per-use basis.

At the end of the commitment period, the pay-per-use invoicing resumes.

The subscription journey of Reserved Instances is described in the following document: https://cloud.orange-business.com/wp-content/uploads/2021/05/RI-explanation-page.pdf

6.1.4 Flexible reserved instances

The "Flexible Reserved Instances" allow the user to benefit from a specific pricing structure compared to pay-per-use, and offer the user the possibility to modify his subscription during the subscription period. The "Flexible Reserved Instances" are subject to the same conditions of access and use as the "Reserved Instances". The different possibilities for modifying existing "Flexible Reserved Instances" are described in the following document : https://cloud.orange-business.com/wp-content/uploads/2020/11/FRI-explanation-page.pdf

6.1.5 Cancellation of Reserved Instances or Flexible Reserved Instance

It is possible to cancel a Reserved Instance or a Flexible Reserved Instance before the end of the commitment. In return, a cancellation fee of 12% of the amount due for the initial subscription period will be applied. The modalities of cancellation of the Reserved Instances and Flexible Reserved Instances are described in a specific document accessible here: https://cloud.orange-business.com/wp-content/uploads/2021/03/Early-Termination-Penalty-Explanation-Page.pdf

6.1.6 VM Auto-Recovery

In the event of a server failure, the Virtual Machines which have been configured with the VM auto-recovery Feature by the User, will be automatically migrated to another computing server host. The new VM will be a clone of the failed VM.

Activation of this Feature is done using the Cloud Eye Service monitoring console, and is limited to compatible flavors.

6.1.7 Auto-Scaling

Auto Scaling (AS) uses preset AS policies to automatically scale service resources up and down based on service requirements. The User can configure scheduled and periodic scaling tasks, monitoring policies, and AS group capacity thresholds to enable AS to automatically increase or decrease the number of Elastic Cloud Server (ECS) instances.

A flexible web-based self-service management console is provided for the User to manage and control the AS service.

In addition, AS can work with Elastic Load Balance (ELB) to automatically scale load balancers members.

If the User need to deploy a distributed application system on a cloud platform, the User can use AS to follow the demand curve for the User's system closely by planning scaling activities and configuring automatic resource adjustment based on monitoring data.

6.1.8 Image Management Service

An image is used to create ECSs and consists of a preinstalled public or private operating system and, if any, applications. IMS allows the User to create, edit, upload, and delete images in self-service mode.

The User can use the Flexible Engine console to provision ECSs using images, either one by one or in batches.

IMS allows the User to:

- Create ECSs using public images available in a Region with preinstalled Software.
- Create a private image using an existing ECS.
- Query details about a private image.
- Delete an existing private image.
- Upload an image file and register it as a private image.
- Export a private image in a specified format.
- Share a private image with other users.

The list of public images is available on the Flexible Engine Console and is subject to evolution.

In case the application images are provided by a provider different from the Provider, only the provider is accountable for the applications Features. Support on these applications is ensured solely by the provider, a specific support contract must be bought by the Customer from the provider. Some images are provided in BYOL (Bring Your Own License) mode in which case the Customer should take care of the licensing. The providers and licensing modes for each image are indicated on the Flexible Engine Console.

The User can import private images towards the Public Cloud. Verifying the proper operation of private images is under the responsibility of the Customer.

The private images that can be exported include those that the User has uploaded to the system or established from the ECSs created from the free public images. Exported images may be used as backups.

6.1.9 Cloud Container Engine (CCE)

The Cloud Container Engine (CCE) service is a container service that features high availability and elastic scalability. With CCE service, users can create, run, and stop Docker containers conveniently. The CCE service also provides a graphical application orchestration tool for users to create and deploy applications efficiently.

The CCE service supports only stateless Docker applications at present.

The CCE service uses Kubernetes to deploy and manage Docker applications, and provides a unified interface for users to manage applications.

Features:

- **Application management**: This feature enables users to create, update, delete and query Docker container applications. It also supports management of application templates and component templates.
- **Graphical orchestration**: This feature provides a graphical orchestration tool for users to define topology structures by dragging components and to deploy applications.
- **Private image management**: This feature enables users to manage private images, such as uploading, updating, or deleting images.
- **Cluster management:** This feature enables users to manage container clusters, such as creating, updating, or deleting a container cluster.
- **Application Elastic scaling:** This feature enables users to scale required resources based on load conditions of applications to flexibly respond to Internet traffic changes.
- **Monitoring and log query:** This feature enables users to monitor applications' CPU usage and memory usage with a graphical display. It also supports collection and download of logs.
- **ELB to Application:** This feature enables user to apply ELB to Application.

For each Docker container, the User can configure the memory size and CPU specifications

The total number of nodes that can be created on clusters of each tenant is also restricted by the resource quota (ECS, VPC, etc…) of the tenant.

6.1.10 Dedicated Cloud (DEC)

The "Dedicated Cloud" service makes it possible to provision a pool of isolated hypervisors in the public Cloud. In this way, the Customer benefits within its Tenant from dedicated physical servers to build its own virtual resource groups. The Client can connect its dedicated cloud to virtual networks, dedicated storage resources or distributed storage resources (EV, OBS) and use other Flexible Engine services to create ECS, load public or private OS images (IMS)..., establish backups (VBS)...

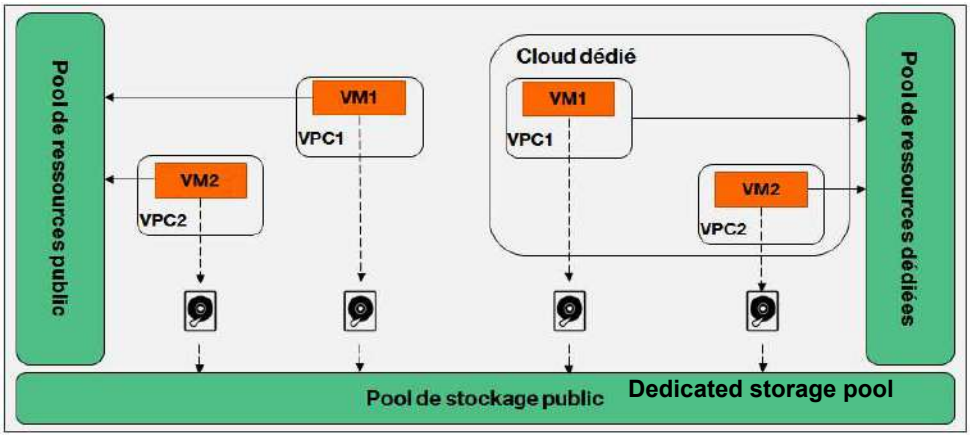

Figure No. 4: Schematic diagram of the Dedicated Cloud Service

The available servers and the invoicing principle are specified in the Price List.

6.1.11 Flexible Engine Stack

Using Flexible Engine Stack, the Customer can use Flexible Engine services on a private infrastructure for the organization's needs. The Customer will be able to use the same user interface as Flexible Engine to control the Customer's cloud, public and private, or automated APIs.

6.1.12 Dedicated Host

Dedicated Host Service (DeH) is a service that provides dedicated physical servers for the Customer's use, on which the Customer can create or migrate Elastic Cloud Server (ECS) virtual machines. DeH allows the Customer to improve the isolation, security and performance of the Customer's ECS. When the Customer migrates servers to a DeH, the Customer can continue to use the software licenses of his old server.

6.1.13 FunctionGraph

FunctionGraph is an IT service that provides serverless functionality. It automatically changes to accommodate fluctuations in resource demand during peak activity periods, while requiring no server reservations or reserved capacity. Flexible Engine users can write their own code and specify conditions.

FunctionGraph is compatible with both Log Tank Service (LTS), which allows users to view function execution logs without having to configure them, and Cloud Eye, which allows users to view graphical function monitoring data without having to configure them.

6.1.14 Server Migration Service (SMS)

Server Migration Service (SMS) offers Private to Virtual (P2V) and Virtual to Virtual (V2V) migration services to help the Customer move data and applications from local physical x86 servers or virtual machines on private clouds (Flexible Engine Stack, Dedicated Cloud) or public clouds to ECS virtual machines on shared (Elastic Cloud Server) or dedicated (Dedicated Host) resource pools.

6.2 Flexible Engine storage services

6.2.1 Elastic Volume Service

Elastic Volume Service (EVS) is a scalable virtual block storage service based on the distributed architecture. The method for using an EVS disk is the same as that for using hard disks on traditional servers.

The EVS disk provides high data reliability and I/O throughput and is easy to use. Therefore, it can be used by file systems, databases, and other system software or applications that require block storage devices.

Volumes are highly available and are used as partition of servers starts or also as storage devices of additional data. Block volumes are available in two performance ranges:

- Standard range using SATA disks
- High I/O range using SSD disks

6.2.1.1 Description

EVS provides the User with high-performance, persistent block storage. The User creates EVS disks and attach them to Elastic Cloud Servers (ECSs) so that the ECSs can access and use the disks.

EVS provides the following features:

- Supports different types of EVS disks, including common I/O ultra-high I/O EVS disks.
- Allows the User to expand the EVS disk capacity elastically to meet the increasing requirements for storage capacity.
- Works with Volume Backup Service (VBS) to provide the backup service..
- Provides a system disk with a capacity of 1 GB to 32 TB and a data disk with a capacity of 10 GB to 32 TB.

6.2.1.2 Specifications

Item	Common I/O	Ultra-High I/O
Maximum capacity of a single disk	32 TB	32 TB
Maximum IOPS per EVS disk	1000	20,000
Maximum throughput per EVS disk	40 MB/s	320 to 350 MB/s
Average response time	10 ms to 15 ms	1 ms to 3 ms

Charging of EVS service is based on usage.

6.2.2 Object Storage Service

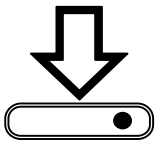

Serving as a cross-platform storage architecture featuring high reliability and secure data sharing, Object Storage Service (OBS) provides customers with secure and reliable data storage at an affordable price. OBS delivers powerful capabilities, including bucket creation, modification, and deletion as well as object upload, download, replication, modification, and deletion. It can store any type of files and is suitable for common users, websites, enterprises, and developers.

As an Internet-oriented service, OBS provides web service interfaces (WSIs) over Hypertext Transfer Protocol (HTTP) and Hypertext Transfer Protocol Service (HTTPS). Users can use OBS Console or clients to access and manage data stored on OBS from any computer connected to the Internet anytime, anywhere. Additionally, OBS is compatible with most Amazon Simple Storage Service (S3) application platform interfaces (APIs). Users can invoke the Representational State Transfer (REST) APIs of OBS and software development kits (SDKs) to develop software adapting to upper-layer applications or connect to the Amazon S3 storage. This allows users to focus on service applications instead of the underlying storage implementation technologies.

OBS offers three storage classes: Standard, Warm and Cold. OBS Standard features low access latency and high throughput. OBS Warm is suitable for storing data that is infrequently accessed but requires fast access response. OBS Cold is oriented to data archiving and long-term backup with rare data access.

Charging of OBS service is based on on storage used and requests executed, depending on the storage classes.

6.2.3 Local Storage to Cloud servers

Local Storage of Cloud Servers optimized for Big Data (‘d’ i) is intended for high performance intensive uses of Big Data.

Described as "ephemeral" disk because it has the characteristic of being located on the internal disks of the hypervisor where the Client creates the server, it is destroyed when the VM is destroyed. This behavior must therefore be managed at the application level.

It is particularly suitable for Big Data clusters and No SQL databases whose applications take full advantage of its reduced access time, the possibility of parallelization and its large bandwidth. Up to 24 1.8 TB local volumes can be configured for distributed Big Data clusters.

6.2.4 Volume Backup Service

Volume Backup Service (VBS) provides snapshot-based protection for Elastic Volume Service (EVS) disks in the public cloud.

VBS provides online one-click backup and restoration for EVS disks, such as system and data disks of ECSs, allowing the User to leverage another layer of security. If an Elastic Volume Service (EVS) disk of an ECS is faulty or logic errors occur on data, the User can use the backups to quickly restore data.

VBS provides disk backup services. A web-based management console is provided for the User to back up the User's EVS disks.

VBS provides the following functions:

- EVS disk-level complete or incremental backup service
- Manual backup service or automated backup policy
- Backup task status query
- Efficient EVS disk creation or reversion to the original state
- Storage type is SATA supported by OBS.
- Multiple copies in different AZs supported by OBS.
- Cross AZ restoration of EVS disks

Charging of VBS service is based on usage.

1. Specifications

- Maximum 360 backup sets for each Tenant.
- Each EVS disk of a Tenant supports up to 20 backup sets.
- Total 200TB of capacity for each Tenant;
- Maximum 5 concurrent VBS backup operations executing at one time, including create backup, delete backup and rollback. More operation will be queued.

6.2.5 Cloud Server Backup Service

Cloud Server Backup Service (CSBS) offers the backup protection service for Elastic Cloud Servers (ECSs) towards Object Storage Service (OBS). It works based on the consistent snapshot technology for Elastic Volume Service (EVS) disks. Backups of all the EVS disks on an ECS are generated at the same point in time.

By default, only the first backup is full and subsequent ones are incremental. CSBS performs the following functions: manual backup, automatic backup and restoration.

The CSBS service is charged based on OBS usage plus one fixed monthly fee for each backed up VM.

2. Limitations

- Applications and file systems on the ECS are not suspended before backup, and memory data is not backed up.
- Each ECS can be associated with only one backup policy.
- A maximum of five EVS disk backup creation and/or deletion jobs can be executed concurrently for each Tenant.
- Backup creation or deletion jobs are applied to whole ECS, including all their EVS disks.

6.2.6 Storage Disaster Recovery Service (sDRS)

Storage Disaster Recovery Service (sDRS) allows the Customer to restart his IT activity on another Flexible Engine AZ. Thus, sDRS allows the Customer to set up a DRP (Disaster Recovery Plan) adapted to failures or disasters affecting his applications or the nominal infrastructures on which these applications run.

The Client is autonomous and solely responsible for maintaining in operational conditions and activating the protection of his activity. sDRS allows him to select the VMs to be protected ; to test the restart of his activity on the recovery site ; to switch his activity to the recovery site ; to restore his activity on the nominal site.

sDRS is based on VM replication, with associated applications and data. This replication is done between two AZs of the same Flexible Engine Region.

The following table presents the costs incurred by the Client in setting up a DRP :

Mode : Costs incurred	Protection	Test	Recovery	
Number of protected VMs	X	X	X	Costs specific to sDRS
Volume of data transferred between the nominal AZ and the recovery AZ	X	X	X	
Storage, on the recovery AZ, of protection data	X	X		Additional costs on Flexible Engine to be taken into account
CPU/RAM/Storage activity, at the recovery site, of test or production VMs		X	X	

3. *Limitations*

- sDRS doesn't constitute a DRP (Disaster Recovery Plan), but only a solution for the Client to set one up.
- The nominal site and the recovery site must be 2 AZs from the same Flexible Engine Region.
- In the Paris Region, the recovery AZ must be EU_West-0a (PA3).
- The Provider makes no commitment to data freshness (Recovery Point Objective) and recovery speed (Recovery Time Objective).

6.2.7 Scalable File Service

Scalable File Service (SFS) provides an on-demand, scalable, and high-performance shared file system accessible to all Elastic Cloud Servers (ECSs) of a given Virtual Private Cloud (VPC) across AZs within a Region.

SFS is charged based on volume of storage used.

4. *Limitations*

- Scalable File Service supports only NFSv3 protocol.
- SFS does not allow modifying the name, AZ and VPC of existing file systems.

6.2.8 Scalable File Service Turbo (SFS Turbo)

SFS Turbo addresses NAS-like service scenarios, providing file services with low latency and high IOPS.

5. *Feature Description:*

- Flexible Interconnection: supports ECS and EVS interconnection, to be deployed as needed.
- Isolation: Each SFS Turbo Service is dedicated to its domain and is not shared with other ECS or EVS.
- High performance: supports 1-2 ms latency and a maximum IOPS of 20,000.
- Support for standard NFS protocol, as well as encryption, backup, and recovery of file data.
- The SFS Turbo instance provides high availability within an AZ

6.2.9 Dedicated Distributed Storage Service (DSS)

Dedicated Distributed Storage Service (DSS) provides dedicated, physical storage resources. It can interconnect with various computing services, such as ECS, BMS, and DCC. It supports disk sharing, disk encryption, disk backup, and snapshots.

6.2.10 File & Application Backup

File & Application Backup is a BaaS (backup as a Service) that enables you to backup your files and applications in Orange Cloud. You can backup resources located on-premise or on any other Cloud infrastructure as illustrated in the below figure.

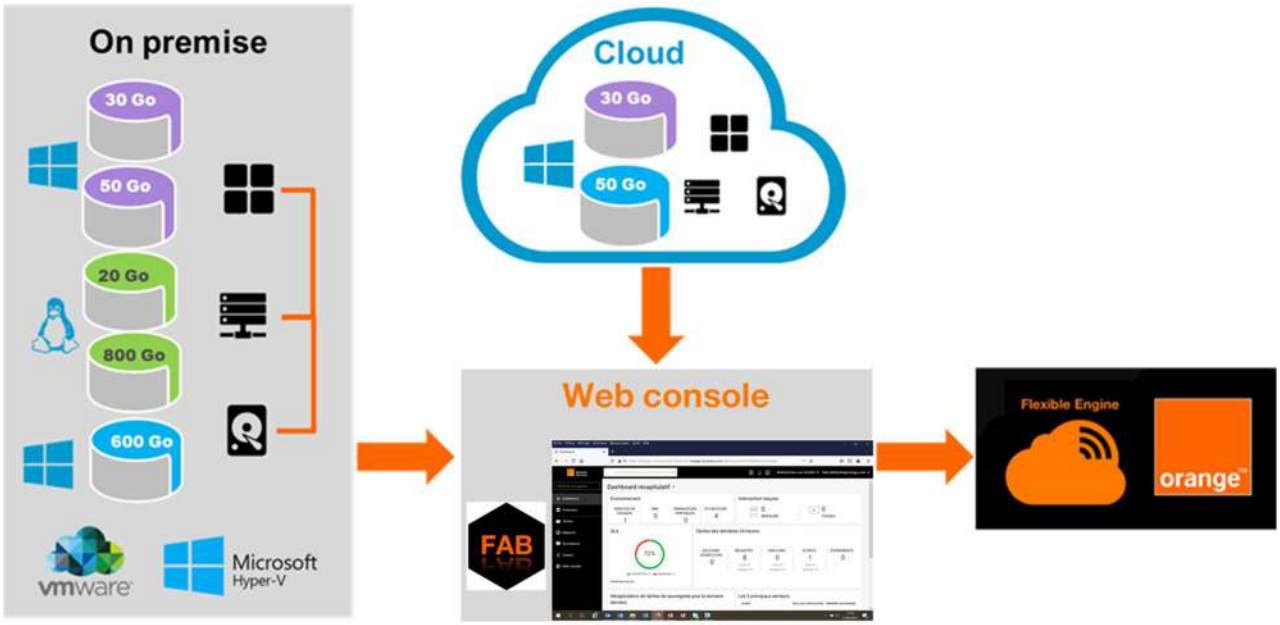

6. Features

Backup features are added progressively according to customer needs. Available to date:

- Granular file backup for Windows and Linux servers
  - You can backup by default all data on a file server, or you can choose drives, folders or even single files,
  - You can filter to exclude or include files based on name or type.
- Database backup for SQL Server, Oracle and MySql
  You can have a consistent backup for your database and a transaction based restore depending on your database.
- Exchange infrastructure
- Office 365 (only Azure based directory)
  You can backup and restore your applications (Exchange Online, Sharepoint, OneDrive, Teams) in a granular way.

7. Characteristics

For your backup and restore FAB has the following characteristics:

- Secured web interface to use your backup solution
- Secured data with encrypted transport tunnels from end to end
- Bandwidth optimization based on protected data compression and deduplication
- Secured and redundant Cloud storage for your backup (provided by Flexible Engine)
- Predefined backup policies (daily backup between 8pm and 8am + 1 weekly) with a choice list for retention periods
- On demand backup at any time according to your needs
- Point-in-time restore based on available backup copies
- Restore on initial destination or any other chosen one
- Granular management for backup / restore
- Access management tool to enable users access and authorization
- Data Replication option on another region for higher security (based on Flexible Engine)
- Dashboard and reports to follow up backup and restore tasks.
- the name, AZ and VPC of existing file systems.

6.2.11 Cloud Backup and Recovery

Cloud Backup and Recovery (CBR) allows users to back up their ECS instances and EVS and SFS Turbo drives. The retention policy (number and duration of backups) can be configured by the user.

8. Cloud Disk Backup:

CBR allows to backup one or more EVS of the domain of the Customer (system or data disk)

9. Cloud Server Backup:

CBR allows the Customer to backup an entire server based on snapshots for ECS and BMS. It is recommended to use Cloud Server Backup in scenarios that require high data consistency, such as RAID clusters.

10. File System Backup:

CBR allows to back up SFS Turbo file systems and to use the backups to create new SFS Turbo file systems so as to avoid the loss of important data.

6.2.12 Data Express Service (DES)

Data Express Service (DES) is a terabyte-scale data transmission service. It uses physical storage media (Teleport devices) to transmit large amounts of data from enterprise data centers to the cloud. It helps resolve problems associated with massive data transfer, such as high network costs and long transfer times.

The service consists in a secure NAS server sent to the Customer's data center and client software (DES Client) to be deployed on a Linux operating system. DESClient allows to select, migrate and encrypt data from the Customer's site to the NAS server (called teleport). The teleport is then sent to the Flexible Engine data center, where the data is migrated to the Customer's selected Object Storage Service (OBS) bucket.

6.3 Flexible Engine Network Services

6.3.1 Virtual Private Cloud

Virtual Private Cloud (VPC) enables the User to provision a logically isolated, configurable, and manageable virtual network environment, improving security of resources of a Region and simplifying network deployment.

The VPC service enables Tenant Users to have complete control over their virtual network environments, including network creation and DHCP configuration. Tenant Users can use security groups to improve security of their network environments. Additionally, they can assign Elastic IP Addresses (EIP) for their VPCs to connect the VPCs to the public network. Tenant Users can also connect VPCs to their physical data centers using a virtual private network (VPN) or using a direct connection,

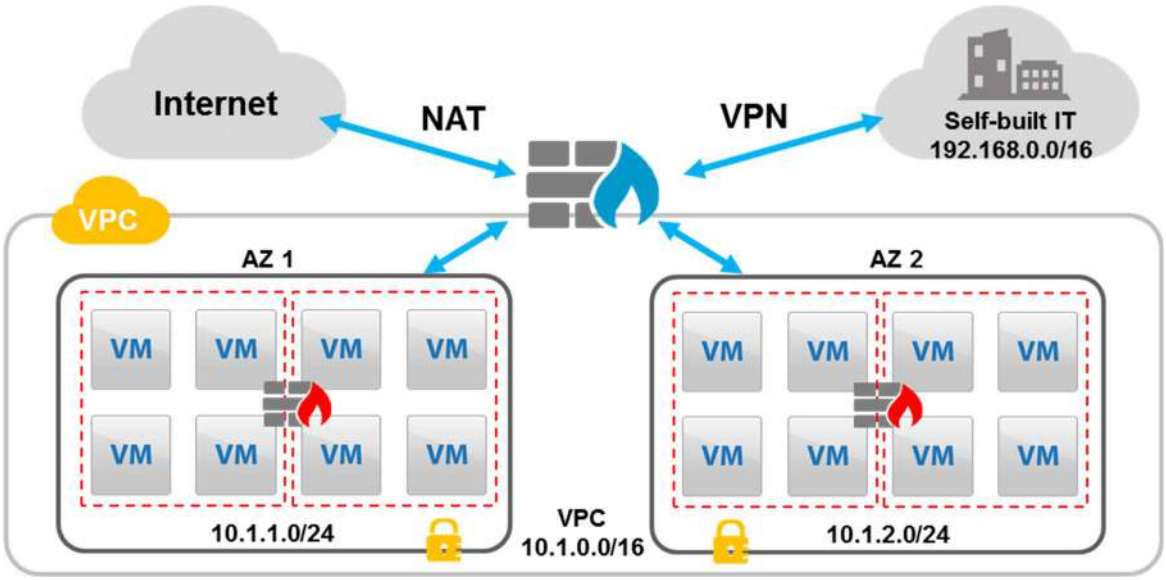

Figure n° 4: *VPC in multi-AZ Region*

11. VPC limitations

- IP address range (RFC1918): 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16
- Bandwidth range for EIPs: 1 Mbit/s to 1Gbps
- VPN subnets: 1000

Technical quotas per Tenant are available on the Flexible Engine Console.

6.3.2 Elastic IP Public Adresses (EIP)

An Elastic IP address (EIP) is a static IPv4 address reachable from the Internet and designed for dynamic computing. An EIP is associated to your Flexible Engine Tenant and you are responsible for binding it to an ECS to enable communication with the Internet.

EIP states are:

- Allocated: Reserved for a Tenant
- Bound: Allocated and bound to an ECS

When an ECS is deleted, bound EIPs remain allocated to the Tenant and may be bound to another ECS.

Allocated public IP addresses are billed per hour (PAYG model).

6.3.3 VPN IPSec aaS

Virtual Private Network as a service is a function of the VPC which allows the creation of a secured IPSec tunnel from the VPC to another IPsec endpoint (such as another VPC within the Flexible Engine cloud) or over the Internet to an external infrastructure.

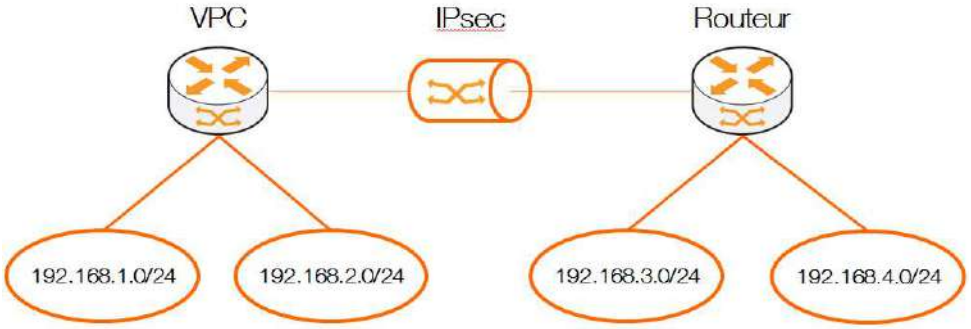

Figure n° 5: *VPN IPSec as a Service*

The VPN as a Service is charged based on usage.

6.3.4 Security Groups

A security group acts as a virtual firewall for your ECS to control inbound and outbound traffic. You can attach to an ECS several security groups. Security groups act at the instance level and not at the subnet level. Therefore, each ECS in a subnet of your VPC could be assigned to a different set of security groups. If you do not specify any security group at the launch of an ECS, the ECS is automatically assigned the default security group of the VPC.

Each security group allows to create and edit rules specifying the source and destination addresses, sources and destination port numbers and protocols.

Security groups are not charged.

6.3.5 Elastic Load Balancer

Elastic Load Balance (ELB) is a service that automatically distributes access traffic to multiple Elastic Cloud Servers (ECSs) to balance their service load.

ELB provides the following functions:

- Traffic distribution across availability zones (AZs), improving reliability and maintainability
- Elastic automatic scaling based on traffic demands
- Linear capacity expansion, eliminating SPOFs
- Support for public network load balancers, which receive requests from clients over the Internet and route the requests to the User's ECSs.
- Support for private network load balancers, which receive requests from clients in the User's VPC and route the requests to the User's ECSs in the subnets.
- Layer 4 (TCP) and layer 7 (Http/Https) load balancing
- Support for ELB monitoring metrics, such as incoming and outgoing traffic, new requests, concurrent requests, incoming and outgoing data packets, active connections, inactive connections
- Working with Cloud Eye (CES) to display monitoring metrics and to allow alarm thresholds to be configured
- Working with AS to implement automatic scaling based on service workload
- Support the load balancer connection draining for http and https protocol
- Support for access logs that capture detailed information about requests sent to the User's load balancer.

The ELB service is charged based on usage.

6.3.6 The private Elastic Load Balancer Service

This Feature automatically distributes access traffic inside a Virtual Private Cloud, without using internet access. The private Elastic Load Balancer Service distributes traffic to multiple ECSs in the VPC within an Availability Zone or across Availability Zones in a given Region.

Private ELB (e.g. for database servers) can be combined with internet-facing ELB (for web servers).

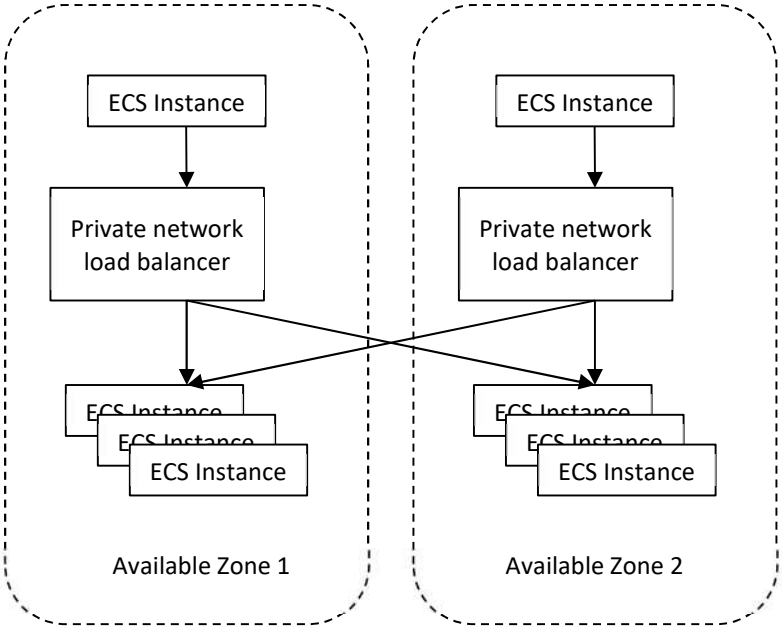

Figure n° 6: *Private network load balancer*

12. Limitations

ECS can only access the private network load balancer in the same available zone.

The private ELB service is charged based on usage.

6.3.7 Internet access

Flexible Engine services are available over the Internet as a default access. Flexible Engine Console, Cloud Customer Space, Object Storage Service are accessed via authenticated login / password.

The User can configure VPC and EIP and Security Groups such that ECS can access to the Internet.

Outband traffic to public addresses is charged on volume.

6.3.8 Direct Connect

Flexible Engine Direct Connect is a solution for directly connecting customer network to FE Virtual Private Cloud (VPC) without using Internet.

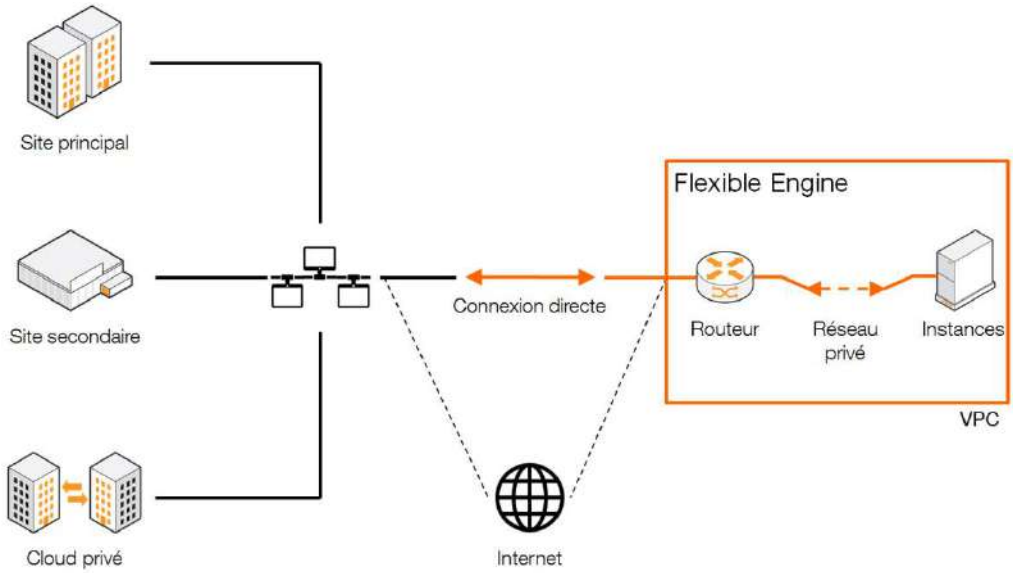

Customer resources (ECS) are contained within a Virtual Private Cloud (VPC) and externalized back to the enterprise over a direct connection which may be through a cross connect called FE Dedicated Port or through a partner network. Available capacity for each solution is given in the table below.

Monthly Direct Connect fee is based on the subscribed bandwidth and do not include transport fees for partner or any other charge to deploy and connect customer routers to FE Direct Connect.

Direct Connect Capacity	FE Dedicated Port	Partner Networks	
		Business VPN Galerie	Equinix Cloud Exchange
2 Mbps		X	
5 Mbps		X	X
10 Mbps		X	X
30 Mbps		X	X
40 Mbps		X	X
50 Mbps		X	X
100 Mbps		X	X
200 Mbps		X	X
300 Mbps		X	X
500 Mbps		X	X
1 Gbps	X	X	X
2 Gbps			X
3 Gbps			X
4 Gpbs			X
5 Gbps			X
6 Gbps			X
7 Gbps			X
8 Gbps			X
9 Gbps			X
10 Gbps	X		X

6.3.8.1 Direct Connect through Flexible Engine Dedicated Port

Customer can directly access 1Gbps or 10Gbps port on FE routers.

For dedicated port (1Gbps and 10Gbps), FE customer is accountable to rent colocation space to deploy its own routers in the Flexible Engine PoP (Point of Presence), and to interconnect these routers to its internal network and to purchase the circuits to connect these routers to FE dedicated port.

6.3.8.2 Direct Connect through a partner network

Current partners are: the Provider (Business VPN Galerie) and Equinix (Equinix Cloud Exchange).

FE Direct Connect through VPN Galerie provides secured MPLS private cloud connectivity between Flexible Engine and customer's existing Business VPN, allowing the End-to-End Private Enterprise Network connectivity to be extended to a Flexible Engine VPC, which will then be seen as an additional site.

Customer can use Equinix Cloud Exchange to connect its private network to Flexible Engine. In that case, the customer should subscribe to Equinix Cloud Exchange.

In order to benefit from the end-to-end MPLS network solution, the customer on one hand needs to activate the FE Direct Connect billable option through a change request and purchase on the other hand the necessary service by the chosen partner.

6.3.9 Domain Name Service

Flexible Engine Domain Name Service (DNS) provides a way for Users and developers to translate a domain name (such as www.example.com) into an IP address (such as 192.0.2.1) so that computers can access applications. With this service, Flexible Engine users can configure the DNS on the FE technical console or through API. DNS service can be used for public and private zones.

Customer is charged according to the number of hosted zones and the number of DNS queries.

6.3.10 VPC Endpoint (VPCEP)

The VPC Endpoint Service (VPCEP) provides secure, private channels to connect the VPC of the Customer to VPC Endpoints (cloud services on your current platform or on your private services), offering a flexible network configuration without the need for an EIP.

6.3.11 NAT Gateway (NAT)

The NAT Gateway service offers the Network Address Translation (NAT) function for Elastic Cloud Servers (ECSs) in a Virtual Private Cloud (VPC), allowing these ECSs to access the Internet using elastic IP addresses (EIPs) or to provide services for external networks.

6.4 Flexible Engine Security and Identity Services

The different security mechanisms generate also events and alerts, consolidated in real-time in a "security events" zone, not accessible to Users. Flexible Engine relies on SOC services for the recurring operation 24/7/365 of these events. SOC offers a specific follow-up of failed VPN connections on the administration network.
Concerning the traces of security equipments, Flexible Engine has access logs to APIs services, administration console and customer dashboard. These data are intended to be communicated to legal authorities.

6.4.1 Confidentiality, Integrity & Proof

Flexible Engine provides services that allow a user to create a virtualized infrastructure over a shared physical infrastructure for all users. The virtualization mechanisms implemented ensure a strong logical partitioning of the client's virtualized resources (one per client). The access to the resources of a Tenant is done through the OpenStack APIs implementing a strong (login / password / token) and secure (in SSL via https) authentication.

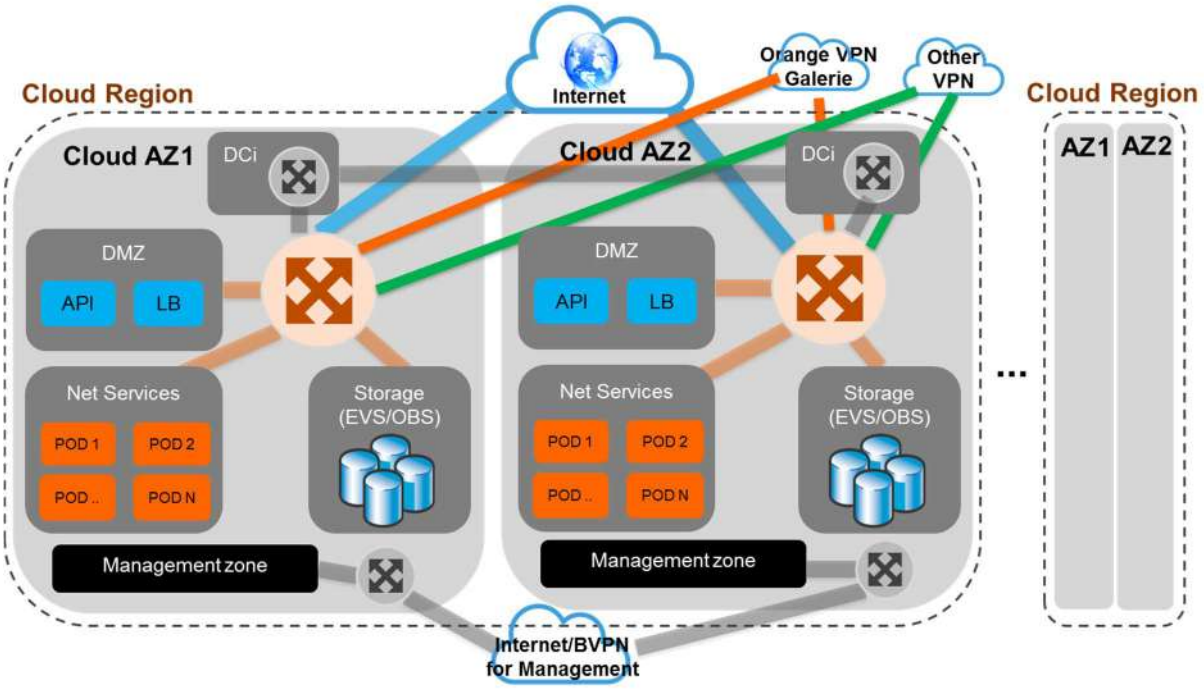

Figure n° 7: *Security and AZ segregation*

13. System Virtualization

The Flexible Engine virtualization platform is based on the OpenSource XEN virtualization engine. In addition to the standard security of this solution, the hypervisor has been hardened to strengthen its partitioning:

- **Processing**: Different virtual server processors have no visibility on each other
- **Memory**: Tests conducted by Orange show that memory remnants from a previously allocated VM cannot be recovered
- **Persistent data**: no local storage on the hypervisor (excluding BigDisk type specific template). The access to the virtual servers deployed in the holding is established by means of a secure connection SSH or RDP.

14. Storage Virtualization:

Segregated access to the stored data (block and object) is ensured by an application layer (Openstack Cinder for storage block, compatible with AWS S3 for object storage) which allows data access only to data owners or the corresponding storage. In addition, data written to the infrastructure is not recoverable once it is deleted by the client or when the corresponding virtual infrastructure is terminated by the customer. These mechanisms are regularly tested thanks to intrusion tests carried out by Orange Cyber Defense and its partners, whose skills are recognized. Physical disks needing to be replaced are destroyed by grinding through a traceable and certified process.

15. Network Virtualization:

VPC features, carried by the Neutron Openstack component, provide a logical logical partitioning of communications on the user network. Any form of network traffic that is not naturally authorized on the customers tenant is not processed by the devices supporting the client's virtual network, preventing any use of spoofing technologies.

6.4.2 Anti-DDoS

The anti-distributed denial of service (Anti-DDoS) aims to provide precise capabilities for defending DDoS attacks, such as challenge collapsar (CC) attacks, SYN flood, and User Datagram Protocol (UDP) flood, for tenants by encapsulating professional anti-DDoS device functions. Tenants can configure anti-DDoS thresholds based on leased bandwidth and service models. The system promptly notifies tenants of the defense status of websites after detecting attacks.

The Anti-DDoS service can defend tenants' public IP addresses against traffic attacks and application-layer CC attacks (mail).

The Anti-DDoS service provides the following functions:

- Defends against traffic attacks and application-layer CC attacks.
- Allows tenants to customize Anti-DDoS policies.
- Allows tenants to select public IP addresses to be defended.

- Provides real-time monitoring reports.
- Provides weekly security reports.

6.4.3 Identity & Access Management (IAM)

IAM centrally controls Users security certificates and Users access policies (which include an access control list). All APIs used for Flexible Engine Services (as well as access to the Flexible Engine Console) are protected by authentication and authorization controls by the IAM Feature.

The IAM service also includes a multi-factor authentication capability (MFA-Multiple Factor Authentication) and a temporary access creation Feature (STS-Security Token Service).

6.4.4 Key Management Service (KMS)

Key Management Service (KMS) is a service that helps Users centrally manage and safeguard their Customer Master Keys (CMKs). KMS uses hardware security modules (HSMs) to protect CMKs.

A Customer Master Key (CMK) is a Key Encryption Key (KEK) created by a user with KMS. It is used to encrypt and protect Data Encryption Keys (DEKs). One CMK can be used to encrypt one or multiple DEKs.

6.4.5 Web Application Firewall (WAF)

The Web Application Firewall (WAF) targets to increase the stability and security of web services as well as the management of related risks. It examines all HTTP and HTTPS requests to detect and block the following attacks: Structured Query Language (SQL) injection, Cross Site Scripting (XSS), web shells, command and code injections, file inclusion, access to sensitive files, exploitation of third party vulnerabilities, Challenge Collapsar (CC) attacks, malicious crawlers, and Cross Site Request Forgery (CSRF).

The WAF allows to protect domain names or IP addresses.

6.4.6 Host Security Service (HSS)

Host Security Service (HSS) is a security manager for servers. It improves overall host security and provides features such as vulnerability management, asset management, baseline inspection and intrusion detection to discover intrusions more rapidly in order to meet compliance requirements.

To use HSS, an agent must be installed on the Elastic Cloud Servers (ECS) to be protected, which allows O&M (operations & maintenance) staff to centrally manage host security through the Security Management Center.

16. Feature List:

- Asset Management: Manages and analyzes security asset information, such as account, open port, process, web directory, and software.
- Vulnerability Detection: Detects vulnerabilities in systems and software (such as Secure Shell, OpenSSL, Apache HTTP Server and MySQL) and provides remediation suggestions.
- Basic Inspection: Discover weak passwords and common system configurations to identify and prevent risks.
- Intrusion Detection: Detects and protects accounts from brute force attacks, webshells, cryptocurrency miners, ransomware and Trojans.

6.5 Flexible Engine Analytics services

6.5.1 Elastic Big Data Service (Map Reduce Service)

The MapReduce Service (MRS) enables the deployment of secure clusters providing computing and storage resources for massive data analysis or real-time processing.

The resources used for calculation and storage can be created and deleted according to the necessary processes in order to optimize costs.

MRS delivers the following functions:

- Hadoop: distributed platform using MapReduce to perform parallel processing on large volumes of data and HDFS for storage
- Spark: distributed processing framework capable of reducing the latency of processing large amounts of data through its in-memory analysis functionalities. It supports Scala, Java and Python languages. It integrates into MRS, Spark SQL to request and analyze data via the standard SQL language.

- HBase (Hadoop Database): distributed non-relational database management system, written in Java, with structured storage for large tables. It provides a reliable, powerful and scalable solution to complete relational databases in massive data processing.
- Hive Apache: data warehouse infrastructure integrated with Hadoop allowing analysis, query via a syntactically close language to SQL as well as data synthesis.
- HDFS (massive data storage): distributed file system that gives high-performance access to data distributed in Hadoop clusters. Like other Hadoop related technologies, HDFS has become a key tool for managing Big Data pools and supporting analytical applications. After being processed and analyzed, the data is encrypted via SSL and stored in object storage (OBS) or HDFS.
- Kerberos: MRS uses KrbServer to provide Kerberos authentication on all components, securing authentication mechanisms.
- Hue: provides a graphical interface (WebUI) for MRS applications, enabling HDFS, MapReduce and database management, HQL and SparQL editing.
- CarbonData: column data format here associated with Spark and which allows to accelerate the requests of an ordeur of magnitude.
- Kafka: distributed streaming platform. It uses the concepts of "publisher" and "subscriber" and allows the collection and consumption of messages in real time.
- Storm: real-time processing system. During Hadoop for real time, Storm allows large scale data flows in real time.
- Loader: Sqoop implementation, allowing to transfer data from Hadoop to structured datastores and to use multiple datasources and exchanges between HDFS, HASE, RDBMS, NFS, SFTP.
- Apache Flume: software for collecting and analyzing log files. The tool is designed to work within a distributed computing architecture to support peak loads.

The types of ECS supported by MRS are listed in the Price List.

17. Usage Restrictions

The following restrictions must be noted during MRS usage: If files are uploaded through the web, the file size cannot be larger than 50 MB. If data is dumped from HDFS to OBS, the maximum data size is 5 GB. The maximum network bandwidth is 5 Gbit/s. For details, see the specification limitations of ECS, VPC, EVS, and OBS.

MRS Service pricing is based on the choice of ECS machines used within the MRS Cluster and is in addition to the price of the ECS service.

6.5.2 Cloud Stream Service (CS)

Cloud Stream Service is a real-time big data stream analysis service running on the public cloud. Computing clusters are fully managed by Cloud Stream Service, enabling the Customer to focus on Stream SQL services. Cloud Stream Service is compatible with Apache Flink APIs, and Cloud Stream Service jobs run in real time.

Powered on Flink and Spark Streaming, Cloud Stream Service integrates enhanced features and security, and supports both stream processing and batch processing methods. It provides mandatory Stream SQL features for data processing. CS is going to be end of life and evolved to Data Lake Insight (DLI).

6.5.3 Data Ingestion Service (DIS)

The Data Ingestion Service (DIS) is a scalable real-time streaming service which can captur and process streaming data. Data send to DIS can be stored for offline processing and analytics. DIS can be useful to scenarios such as capturing IoT sensor data, website clickstreams, stock transactions, social feeds, mobile app gaming telemetries or sensors from autonomous vehicles.

DIS is a high-throughput, distributed message Publish-Subscribe managed web service system. User can create, delete, describe stream throughput through a web console.

6.5.4 Data Pipeline Service (DPS)

Data Pipeline Service (DPS) is a web service running on the public cloud. It enables the Customer to automate the movement and transformation of data between different services.With DPS, the Customer can define a pipeline to describe data processing tasks, task execution sequence, and task scheduling plan. DPS then schedules and controls the execution of tasks based on the pre-defined scheduling plan and relationship, to achieve inter-service data processing and movement. DPS will be end of life and evolved to Data Lake Governance Center (DGC) which is going to be launched soon.

6.5.5 Data Warehouse Service (DWS)

Data Warehouse Service (DWS) is an online database based on the MPP (massively parallel processing) cloud. DWS also provides some tools for database operation and maintenance, including backup and restore, monitoring and database connection.

6.5.6 Machine Learning Service (MLS)

Machine Learning Service (MLS) is an analysis platform service that helps the Customer find patterns in data to construct a machine learning model. The Customer can use this model to process new data and make predictions about service application. MLS is going to be end of life and will evolve to Modelarts.

6.5.7 ModelArts

It allows the Customer to develop AI models for multiple use cases: image, video, voice, object detection, scoring, recommendations and exception detection.

On ModelArts, integrate, build, train and deploy AI models from end to end. Integrate and label the Customer's own datasets, or through external vendors, train models using GPU VM templates, and apply models via inference. Then deploy models in the cloud or on the edge.

6.5.8 Cloud Search Service (CSS)

Cloud Search Service (CSS) is a managed, distributed search service. It is compatible with open-source Elasticsearch and OpenSearch and provides users with structured and unstructured data search, statistics, and report capabilities. CSS works in a similar way as a database.

CSS can be automatically deployed. It provides search engine optimization practices and requires no O&M. Additionally, it has a robust monitoring system that provides key metrics, including systems, clusters, and query performance.

6.5.9 Data Lake Insight (DLI)

Data Lake Insight (DLI) is a serverless data processing and analysis service compatible with standard SQL, Spark and Flink SQL. It also supports multiple access modes, and is compatible with mainstream data formats.

6.5.10 Data Lake Governance Center (DGC)

Data Lake Governance Center (DGC) is a one-stop full-lifecycle data development and operation platform. It provides functions such as data integration and development, supports intelligent construction of the industry's knowledge databases, and incorporates data foundations such as big data storage, computing, and analysis engines, helping enterprise customers build data operation capabilities.

6.5.11 Graph Engine Service (GES)

Graph Engine Service (GES) facilitates querying and analysis of graph-structure data based on various relationships. It is specifically suited for scenarios requiring analysis of rich relationship data, including social relationship analysis, marketing recommendation, public opinions and social listening, information communication, and anti-fraud.

6.5.12 HiLens

HiLens is a multimodal AI development platform that enables device-cloud synergy. It provides a framework, an environment, a cloud-based management console and AI skill market. It also provides Atlas AI edge device with camera that allows the Customer to install AI skills in it.

6.6 Flexible Engine Database Services

6.6.1 Flexible Engine Relational Database Service (RDS)

The relational database service (RDS) allows the deployment of MySQL, PostgreSQL or Microsoft SQL Server databases, with a deployment in simple mode or in active-passive mode.

The installation and deployment of databases is done automatically. The service also offers operation and maintenance tools: PRA, backup and restore, monitoring, migration. The service reduces complexity and maintenance costs, allowing the customer to focus on the application and the business.

The RDS templates and systems as well as their pricing are presented in the Price List.

6.6.2 Distributed Cache Service (DCS)

Distributed Cache Service (DCS) is an in-memory database service compatible with Redis and IMDG. Based on an HA architecture, DCS supports three instance types: single-node, master/standby, and cluster. DCS ensures high read/write performance and fast data access.

6.6.3 Document Database Service (DDS)

Document Database Service (DDS) is a MongoDB-compatible database service that is secure, highly available, reliable, scalable, and easy to use. It provides a variety of functions including DB instance creation, scaling, redundancy, backup, restoration, monitoring, and alarm reporting.

6.6.4 Data Replication Service (DRS)

DRS is a cloud service for database online migration and synchronization in real time which simplifies data migration processes. The Customer can use DRS to transmit data between databases in various scenarios.

6.6.5 Data Admin Service (DAS)

Data Admin Service (DAS) is a one-stop cloud database management platform that allows users to manage databases on a web console. It offers database development, Operations and Maintenance (O&M), intelligent diagnosis, and enterprise-level DevOps, making it easy to use and maintain databases.

6.7 Enterprise Applications

6.7.1 WorkSpace [End of Life]

WorkSpace is a Desktop-as-a-Service (DaaS) solution that allows the Customer to provision Users virtual, cloud-based Microsoft Windows desktops including vCPU, disks and operating systems. This way, Users are able to access them from supported devices.

The list of instances is available on the console and is subject to roadmap evolution.

WorkSpace can be purchased as a flat monthly subscription with unlimited usage rights or as a pay-per-use hourly fee, with or without a monthly subscription.

6.7.2 Remote Desktop Services (RDS/SAL)

RDS allows a User to remotely connect to an enterprise application hosted on a Windows server. By default, two connections are included into Windows server licenses provider by the Provider.

In order to use this service with more than two concurrent accesses, the Customer has to bring RSD/SAL licenses held by him in mobility mode, under the conditions described in the "Licenses / Microsoft Products" section. The Customer must subscribe an RDS/SAL (Subscriber Access License) for each User who may have access to the relevant enterprise application. Machines cannot be licensed.

6.7.3 Office

Office is an office automation software suite. Flexible Engine does not offer this type of licenses.

Each Office license (Standard or Professional Plus) has to be subscribed for a single User, a natural person. These licenses are not eligible for mobility.

However, "Office 365 Professional Plus" licenses can be provided by the Customer and used on Flexible Engine, provided they are declared to the Provider.

Each Office (Standard or Professional Plus) or Office 365 Professional Plus license must be associated with a "Remote Desktop Services" license.

6.7.4 oneclick™

oneclick™ is a Virtual Desktop Infrastructure (VDI) service that allows the Customer to provision desktop environment for Users with the dedicated ECS flavors, disks and operating systems. This way Users are able to access the usual office applications remotely from supported devices in an efficient way.

The list of dedicated ECS flavors instances is available on the console and is subject to roadmap evolution.

oneclick™ can be purchased as a monthly or annual subscription according to license models defined in the Provider website https://cloud.orange-business.com/en/offers/infrastructure-iaas/public-cloud/appliance-catalog/oneclick/ .

Support for the oneclick™ software is provided by oneclick AG themselves, oneclick AG should be the first point of contact. See oneclick website https://help.oneclick-cloud.com/en/ for contact details and instructions how to open a ticket. Specific hours of operation, availability target and response times for tickets are given in the oneclick SLA https://oneclick-cloud.com/en/general-service-level-agreement/ -Gold level.

6.7.5 Distributed Message Service (DMS)

Distributed Message Service (DMS) is a scalable message queuing service that is hosted on the cloud computing platform. Through the use of distributed cluster technology. This service decouples the components of a cloud application.

DMS provides a web-based console for managing message queues and application programming interfaces (APIs) for accessing messages. Using the DMS console, the Customer can create queues and perform message production and consumption tests. User applications can then directly call RESTful APIs, making the DMS service immediately available for applications.

An all-round monitoring and maintenance system has been launched to ensure reliable running of DMS.

All messages stored in isolated message service area are secured against unauthorized access.

6.7.6 Distributed Message Service for Kafka

Distributed Message Service (DMS) for Kafka is a message queuing service based on Apache Kafka and more specifically on Kafka premium instances. The compute, storage and bandwidth resources used by an instance are exclusively occupied by the user.

Apache Kafka is a distributed message middleware that offers high throughput, data persistence, horizontal scalability, and continuous data processing. It adopts the publish-and-subscribe model and is widely used for log collection, data streaming, online/offline systems analysis, and real-time monitoring.

6.7.7 Distributed Message Service for RocketMQ

Distributed Message Service (DMS) for RocketMQ is message-oriented middleware that delivers low latency, high flexibility, high throughput, dynamic expansion, easy management, and abundant messaging functions.

DMS for RocketMQ has the following features:

- Compatibility with open-source RocketMQ clients.
- Abundant messaging functions, including ordered message delivery, delayed messages, scheduled messages, message retry, dead letter messages, and transactional messages, which meet diverse needs in e-commerce and finance scenarios.
- Monitoring and analysis functions, including message tracing, message tracking, trace analysis, dead letter message export, monitoring and alarms, which allow the Customer to monitor services and keep them up and running.

6.7.8 Simple Message Notification (SMN)

Simple Message Notification (SMN) is a hosted simple message notification service that is flexible and large-scale. SMN allows the Customer to send messages to email addresses and HTTP/HTTPS.

6.8 Developer tools and APIs

6.8.1 Flexible Engine Open APIs

The APIs made available by Flexible Engine are RESTful APIs based on OpenStack technology and documented in the Help Center.

6.8.2 Orchestration: Resource Template Service (RTS)

With the Heat RTS (Resource Template Service) orchestrator made available via APIs, the Customer can automatically and configurably deploy a virtualized infrastructure (servers, routers, networks, volumes, etc.) using the various APIs of the Openstack modules.

It is thus possible to create HOT (Heat Orchestration Template) templates that allow to specify the configuration, description and relationships of all resources to automate and facilitate the deployment of the platform.

6.8.3 API Gateway

API Gateway allows the developers to create, publish, secure and monitor the APIs of their applications.

The Feature is charged based on the number of API calls and the outgoing traffic.

6.9 Flexible Engine Management Tools and Portals

6.9.1 Cloud Eye Service

The Cloud Eye Service (CES) is an open monitoring service that allows you to set up monitoring, alerting and supervision for your resources in real time.

It allows metrics to be monitored directly on computing instances (ECS), storage volumes (EVS), Virtual Private Clouds (VPC), load balancers (ELB), autoscaling groups (AS) and aaS relational databases (RDS).

It is thus possible for the Client to configure alerting rules and notification policies based on its metrics to track the status and performance of monitored objects over time.

The features are as follows:

- **Automatic monitoring**: The system automatically starts monitoring based on the resources obtained. Apart from certain ECS metrics, the User do not need to install any plug-ins to monitor service metrics.
- **Flexible alarm function:** The User can flexibly set alarm rules on any of the monitoring metrics, configure alarm thresholds, as well as enable or disable the alarm function.
- **Real-time notification:** The User can configure alarm notification to receive short messages or emails when alarms are generated.
- **Following metrics:** On the Dashboard page of the CES console, the User can follow a metric of a monitored object or delete a followed metric. After following a metric, the User can view monitoring data concerning it each time the User log in to the CES console.

6.9.2 Cloud Trace Service

Cloud Trace Service (CTS) provides operation log on cloud service resources. With this service, the User can query, audit and backtrack operation log, and store traces in OBS buckets with high reliability. Cloud Trace Service records all traces that are triggered by open APIs and Console from every cloud service that's integrated with this feature. The User can create only one tracker for each Region in each Tenant. This feature is not charged.

6.9.3 Simple Message Notification

Simple Message Notification (SMN) is a message notification service. It enables users to send messages through emails, SMS or HTTP/HTTPS to a group of subscribers in batches.
SMN can be integrated with other Features to receive event notifications from them.

SMN is charged based on number of API calls, number of notifications, number of SMS and their destination, and volume of Internet traffic used.

18. Limitations

- The subscription takes effect only after the subscriber confirms the subscription. Subscribers must be invited and confirm their subscription to receive messages.
- The maximum message size is restricted to 256 kB.
- Messages are reserved for 7 days and the system automatically clears them afterwards.

- Upon a message pushing failure, the system tries to send the message for another 6 times. If the pushing still fails, the system abandons the message.

6.9.4 Tag Management Service (TMS)

Tag Management Service (TMS) is a service for tagging and categorizing cloud services. Users can use tags to classify and search cloud resources by purpose, dimension, project, environment... Supported resources are: ECS, OBS, VPC, VBS, EVS, AS, IMS.

6.9.5 Application Operations Management (AOM)

Application Operations Management (AOM) is a multi-dimensional O&M management platform for cloud applications. It monitors applications and related cloud resources in real time, collects and associates resource metrics, logs, and events to analyze application health status. It also provides flexible alarm reporting and abundant data visualization, thus facilitating the detection of errors.

Specifically, AOM comprehensively monitors and uniformly manages servers, storage devices, networks, web containers, and applications hosted in Docker and Kubernetes, so as to prevent problems, facilitate fault locating and reduce O&M costs.

6.9.6 Log Tanks Service (LTS)

Log Tank Service (LTS) provides the Customer with a solution to collect log data from various cloud services, servers, and network devices. It helps the Customer quickly find specific data from mass data sets and locate problems. In addition, it supports long-term storage and prevents loss of log data.

6.10 Container

6.10.1 Application Performance Management (APM)

Application Performance Management (APM) is a cloud service that monitors and manages cloud application performance and faults in real time. It provides professional distributed application performance analysis capabilities to help O&M personnel quickly resolve problems such as error assessments and removal of performance bottlenecks in the distributed architecture.

6.10.2 Application Orchestration Service (AOS)

With Application Orchestration Service (AOS), users can deploy applications in the cloud by writing templates (declarations of resources that make up stacks), and create stacks from the templates. AOS also provides application lifecycle management features, such as starting, changing, and deleting.

6.10.3 Application Service Mesh (ASM)

Application Service Mesh (ASM) is a non-intrusive solution for the Customer to manage microservice lifecycle and traffic. It is compatible with the Kubernetes and Istio ecosystems and hosts a wide range of features such as load balancing, circuit breaking, and fault injection. In addition, it provides diversified built-in grayscale releases, including canary release and blue-green deployment, enabling one-stop automatic release management.

6.10.4 Intelligent EdgeFabric (IEF)

Intelligent EdgeFabric (IEF) provides users with a complete edge computing solution where cloud applications are extended to the edge. By leveraging edge-cloud synergy, users can manage edge nodes and applications remotely while still processing data nearby. In addition, they can perform Operations & Maintenance (O&M) in the cloud, including edge node monitoring, edge application monitoring, and log collection.

6.10.5 Multi-cloud Container Platform (MCP)

Multi-Cloud Container Platform (MCP) provides multi-cloud and hybrid-cloud containerized solutions for unified cluster management across clouds, unified deployment and traffic distribution of multi-cluster application. It is designed to resolve multi-cloud disaster recovery and plays an important role in various scenarios including traffic sharing and separation of services and data, development and production, and computing and services.

6.10.6 Software Repository for Container (SWR)

Software Repository for Container (SWR) provides management of Docker container images throughout their lifecycles, featuring image push, pull, and deleting.

Private image repository and fine-grained permission management allow users to grant different access permissions, namely, read, write, and edit, to different users. Every time an image is updated, the application deployed in Cloud Container Engine (CCE) with this image will be automatically updated. The trigger option enables automatic application update.

Users can push, pull, and manage Docker images by using the SWR console, SWR APIs, or Docker Command Line Interface (CLI).

6.11 Flexible Engine HDS Certification

Flexible Engine HDS certification guarantees to healthcare companies in France that, as customers of Flexible Engine HDS, the infrastructure services they use are compliant with French Public Health Code (Article L.1111-8), which states that any healthcare organization (hospitals, pharmaceutical companies, laboratories) in France that manages personal medical data must use an HDS certified service provider.

Flexible Engine should be used in conjunction with Business support solution or higher to be certified HDS.

Flexible Engine HDS customers have access to a French speaking Helpdesk who's staff are trained in HDS certification.

Flexible Engine HDS covers the Levels 1, 2 and 4 of the HDS cerfiication. The Provider only intervenes in these levels. It is the Customer's responsibility to ensure the applicability, and compliance where applicable, of the regulations relating to the hosting of health data concerning the other Levels.

Level 1 = Datacenter

Level 2 = Hardware for Servers, Storage, Network

Level 4 = Virtual infrastructures

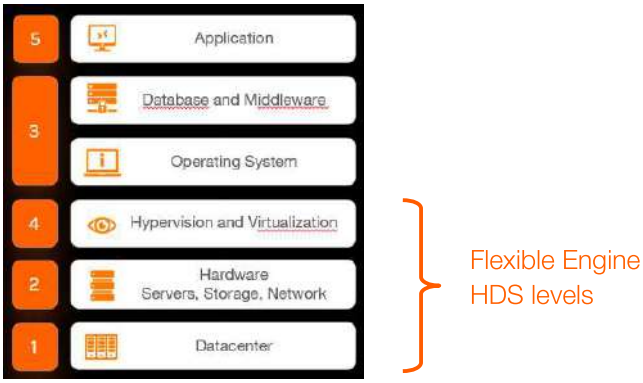

6.11.1 Audit

The Provider can communicate certification audit reports to Customers who request them.

The Provider regularly commissions internal audits in terms of information security. In addition, it may host third-party audits on behalf of the Customer in compliance with the conditions set out in the article Protection of personal data of the General Conditions.

As part of the Flexible Engine HDS service, it is specified that no Personal Health Data is transferred outside the European Union and/or to a country that does not have legislation on the processing of personal data deemed adequate by decision of the European Commission.

7 Support

The purpose of this chapter is to describe the support services provided by the Provider within the context of the Services of the Flexible Engine solution, their organization and the models of the associated processes.

This chapter gives details of:

- The support solutions offered to the Customer;
- The organization of communications between the Provider and the Customer;
- The organization and scope of the support activities provided by the Provider;
- The prerequisites required for the provision of the support by the Provider;
- How an incident should be reported or a request should be made to Technical Support;
- How Technical Support records and deals with an incident or a request;

7.1 Scope of application

The Provider does not take any support commitment for Beta Features.

7.2 Definitions

- **The Catalogue** means the catalogue of Services as described in the Description of Flexible Engine Services annex
- **An Incident** means an unplanned break in the Services or a reduction in the quality of the Services.
- **A Ticket** means a record in the Provider ticketing tool for any request or reporting of an Incident on the Services. A Ticket is used to exchange information between the Technical Support team and the contacts named by the Customer when managing and processing a Request or an Incident regarding the Services.
- **A Problem** means an Incident regarding the Services which is recurrent or major which has been resolved but whose origin (cause) has not yet been identified.
- **A Known Error** means an Incident regarding the Services that has already occurred (recurrent) and whose origin is known.
- **A Request** means a request from the Customer to obtain an item of information of advice or a standard change or to have access to one of the Services.
- **A Loss of Service** means a total break in the Services caused by the shutdown of one or more elements of the Provider production platform (critical or non-critical) and which makes the provision of the service concerned impossible.
- **A Scheduled Shutdown** means a shutdown period of the platform of Services that was planned and previously announced by the Provider. During this shutdown period, all or some of the Services are unavailable.
- **A Unscheduled Shutdown** means an Incident that occurs outside of the scheduled shutdown period and which it results in a loss of Services.
- **The Help Center** means the documentary pages of Flexible Engine accessible from the Flexible Engine Console.
- **Working Hours** are between 9.00 and 19.00 (French time), Monday to Friday, excluding French public holidays.
- **The Support Services** are the support services of the Services delivered by the Provider.
- **The Support Plans** are the products to which the Customer may subscribe, giving an entitlement to different levels of commitment of Support Services by the Provider.
- **The Customer Service** of the Provider is the Customer Service for the Flexible Engine solution Services.
- **The Website** means the cloud Orange Business portal website of the Provider whose URL is http://cloud.orange-business.com/
- **The Technical Support** of the Provider is the technical support for the Services.
- **The Single Point of Contact (SPOC)** is the whole Technical Support team of the Provider for the Services of the Catalogue which is the dedicated contact for the Customer in the event of a Request or an Incident.
- **A Request for Change (RFC)** is a formal request for the improvement or modification of the Services, issued by the Provider, and which gives details of the required change. The description of an RFC is indicated in the change management procedure as indicated in the “Change management” paragraph of this document.
- **A Named Contact** is a functional expert designated by the Customer with good knowledge of the Services. Only a Named Contact is entitled to issue Requests or report Incidents on the Services to the Provider technical support.
- **Response Time** refers to the time elapsed between a ticket’s opening and notification to the Customer of the Provider’ having taken it into account, minus the periods during which the Provider’ engagements do not apply.
- **Response Time Objective** refers to the Response Time within which the Provider undertakes to respond to an incident ticket, in accordance with the level of support to which the Customer has subscribed. Failures to

reach Response Time Objectives do not entitle the Customer to claim Service Credits. Only a repeated and prolonged failure to comply with Response Time Objectives may be considered a breach to the Provider commitments.

- **Fault Repair Time** refers to the time elapsed between an Incident ticket's opening and its resolution, minus the periods during which the Provider' engagements do not apply.
- **Fault Repair Time Objective** refers to the Fault Repair Time within which the Provider commits in the event of Incident in production environment, in accordance with the level of support to which the Customer has subscribed. Failures to reach Fault Repair Time Objectives do not entitle the Customer to claim Service Credits. Only a repeated and prolonged failure to comply with Fault Repair Time Objectives may be considered a breach to the Provider commitments.
- **Priority** refers to the following levels used by the Provider to classify Incident tickets:
    - **Priority 1** (or **P1**): complete loss of Service for more than 50% of Users, or Incident with a critical impact on the Customer's activities
    - **Priority 2** (or **P2**): Services deteriorated. Users are able to access the Services, but experience difficulties or must deal with significant delays, or complete loss of Service for 5 to 50% of Users.
    - **Priority 3** (or **P3**): Services provided with delay or minor difficulties, or complete loss of Service for less than 5% of Users. The Customer's activity is not significantly impeded.
    - **Priority 4 (**or **P4):** these tickets are not related to Incidents, and quality of service commitments by the Provider are applicable only for Platinum support plan.

7.3 Organisation of the Support Services

The organisation of the Support Services is structured around two entities:

- Customer Service;
- Technical Support (International or European, based on the support level).

Once the Customer has subscribed to a Support Services solution, it can contact the Provider to:

- Make service Requests;
- Report an Incident.

The purpose of the Provider support is to manage Requests and Incidents, by carrying out the following actions:

- Take charge of the Requests and Incidents, and processing and resolving them by following procedures for managing Incidents and Requests as defined in this service agreement;
- Communicate the appropriate, up-to-date information to the Customer's Named Contact regarding the processing of the Incidents and Requests which have been duly reported;
- Improve and update the technical support procedures.

7.3.1 Support Plans for Flexible Engine

The Provider offers its Customers the following Support solutions:

- **Basic**: documentation, FAQ online on the Website and on the Help Center for independent viewing for all Customers of the Services.
- **Standard**: support solution designed for Customers whose use of the Services is intended for application developments excluding production applications.
- **Business**: support solution designed for Customers whose use of the Services is intended for production applications.
- **Business Europe**: support solution designed for Customers who prefer their data not to be transferred to or accessible from a country outside Europe.
- **Premium**: support solution designed for Customers whose use of the Services is intended for demanding production applications.
- **Platinum**: support solution designed for Customers whose use of the Services is intended for highly business critical and large infrastructure.

Support services are subscribed for a minimum of 6 months. The Customer may only upgrade its Order during the minimum period. The minimum period is then extended for 6 month on the new level subscribed.

Changes of support level take effect at the beginning of a calendar month.

Support Plans for Flexible Engine Services	BASIC	STANDARD	BUSINESS	PREMIUM	BUSINESS EUROPE

Customer Service					
Flexible Engine documentation via Help Center	Included	Included	Included	Included	Included
API documentation via Help Center					
Questions relating to the account, subscriptions, invoicing	Working hours	Working hours	Working hours	Working hours	Working hours
Support location	International	International	International	International	Europe
Technical Support					
Supervision of data-centres 24/7	Included	Included	Included	Included	Included
Means of Access					
Ticket via Flexible Engine Customer Space	Ticket received 24/7	Ticket received 24/7	Ticket received 24/7	Ticket received 24/7	Ticket received 24/7
eMail	N/A	N/A	N/A	Yes	Yes
Phone	N/A	N/A	Phone with prior ticket	Phone with prior ticket	Phone with prior ticket
Ticket processing hours	Working day	Working day	P1, P2: 24/7 P3: Working hours	24/7	P1, P2: 24/7 P3: Working hours
Response Time Objectives					
for P4 incident tickets	N/A	N/A	N/A	24 hours / 5 working days	N/A
for P3 incident tickets	N/A	1 working day	1 working day	8 hours	12 working hours
for P2 incident tickets	N/A	1 working day	2 hours	1 hour	2 hours
for P1 incident tickets	N/A	12 working hours	1 hour	15 mins	1 hour
Fault Repair Time Objetives					
for P4 incident tickets	N/A	N/A	N/A	N/A	N/A
for P3 incident tickets	N/A	N/A	N/A	N/A	N/A
for P2 incident tickets	N/A	N/A	API incident only 4 hours	API incident only 4 hours	API incident only 4 hours
for P1 incident tickets	N/A	N/A	API incident only 4 hours	API incident only 4 hours	API incident only 4 hours
Assistance on best practices	Optional: Cloud Coach Service				
Change management	Optional: Managed Application Offers				
OS and application monitoring and supervision	Optional: Managed Application Offers				
Names Contacts					
Availability of single-tenant support services	N/A	Standard single-tenant support offer	Single-tenant business support offer	Single-tenant premium support offer	Single-tenant business support offer
Number of named contacts for a single-tenant service	N/A	2	5	Unlimited	5
Availability of multi-tenant support services	N/A	N/A	Multi-tenant business support offer	Multi-tenant premium support offer	Multi-tenant business support offer
Number of named contacts for a multi-tenant service	N/A	N/A	Unlimited	Unlimited	Unlimited
Priority refers to the following levels used by the Provider to classify Incident tickets: o Priority 1 (or P1): A complete breakdown/outage of the Service Loss of Service, or Service unavailability for 50% and more of customers, or High impact (key) or corporate customer(s) affected. o Priority 2 (or P2): Services deteriorated. Users are able to access the Services, but experience difficulties or must deal with significant delays, or complete loss of Service for 5 to 50% of Users. o Priority 3 (or P3): Services provided with delay or minor difficulties, or complete loss of Service for less than 5% of Users. The Customer's activity is not significantly impeded. o Priority 4 (or P4): General development question, request a feature, These tickets are not related to Incidents and quality of service commitments.					

Tableau n° 1: *Support Plans Table*

7.3.2 Self-service for Flexible Engine Support

Various documentary supports and contents are available to the Customer on the Help Center to assist with the use of the Services:

API documentation: the Customer will find the functions of the Provider platform here and find out more details of its uses.

This content is regularly updated by the Provider.

7.3.3 Technical Support

The Technical Support team deals with requests from the Customer for the Services and takes into account the Incidents reported during the time slots and within the timeframes laid down according to the support solution to which the Customer has subscribed.

For the STANDARD, BUSINESS AND PREMIUM Support Service solutions the Provider undertakes to:

- Take charge of the Customer's Requests and Incidents according to the timeframes stipulated in the support solution concerned;
- Respond to all requests for information regarding the Customer account, its Named Contacts and its invoices;
- Respond to all requests for information regarding the functions of the Services;
- Make its best efforts to resolve any Incident linked to the proper operation of the Services;
- Make its best efforts to resolve any Incident regarding the APIs of the Flexible Engine Services;
- Make its best efforts to a Customer having subscribed to a BUSINESS or PREMIUM Support solution to provide technical assistance related to the use of the operating systems provided by the Provider in the Catalogue. In any event, it is stated that the Provider support team does not provide direct administration of systems as part of Flexible Engine support.

In addition to the above, for the BUSINESS EUROPE Support Service solution:

- The L2 and L3 support teams are located in the European Union

Technical Support of the Provider does not cover:

- Requests and Incidents related to operating systems / third party software other than that contained in the list of available images provided by the Provider as described in the Catalogue.
- Requests and Incidents related to operating systems / software not provided by the Provider;
- Requests and Incidents related to the Customer's architecture or software;
- Requests and Incidents related to system administration tasks of the Customer's virtual machines;
- The development of code or scripts.

7.4 The Customer's competencies and responsibilities

The Customer's Named Contact entitled to contact the Technical Support is a person who is deemed to be trained and competent in the use of the Cloud and the Services' APIs.

The Customer undertakes to establish good practice in the use of the cloud and all possible means to process the Requests and resolve the Incidents reported on its service before reporting an Incident to the Technical Support. It has, among other things, APIs of the Services made available to it to take action by itself. Should the Customer be unable to process the Request or resolve the Incident detected on the Services due to a malfunction of the Services, it will then be able to send the Request or the Incident to the Technical Support by following the procedure for managing incidents described in this document.

7.5 The interfaces and ways of contacting the Customer Support

The Customer must provide the Customer Service with the list of Named Contacts who alone are entitled to make a Request or report an Incident with the Technical Support

To do so, it will communicate to the Provider, a list designating the names of the Named Contacts including, at least, the following information:

- Last name;
- First name;
- Position;
- Email address;
- Phone;
- Availability time-slot.

These Named Contacts will be the only points of contacts used to inform the Customer of any Incident detected by the Provider on the Services used by the Customer. The number of Named Contacts entitled to contact the support is defined in the **Support Plans Table** in this document.

7.6 Description of the support model

The support model defines the organization and the exchanges between the Customer's support and that of the Provider.

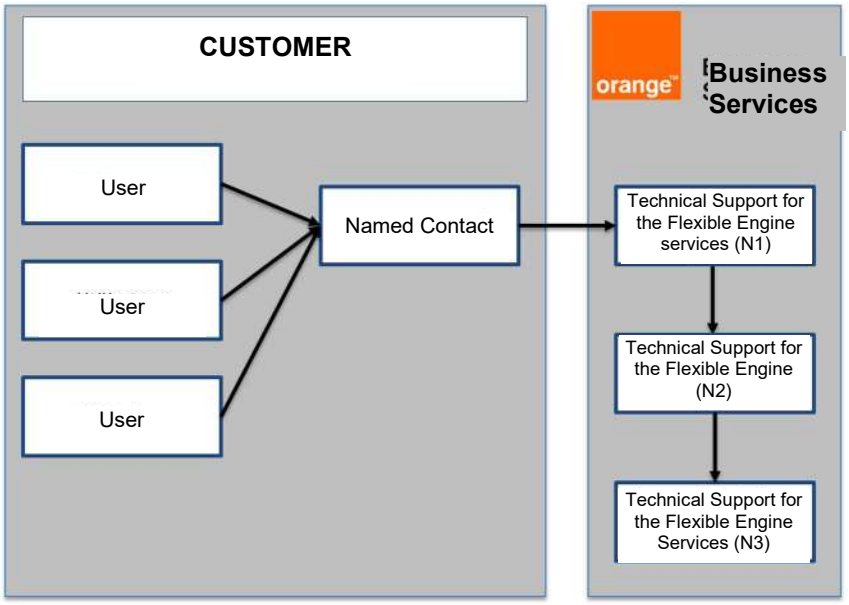

Diagram: Support model for the Customer

7.6.1 RACI Chart - support for the Services

For Support Services, here is the table of responsibility

Task	Customer	The Provider
Supervision of virtual infrastructure instantiated by the Customer, software and services of the Customer	R,A	
Implementation of good practice for the use of the cloud: automation of deployments, restarting / redeployment procedures, making highly available, security rules, backups.	R,A	
Resolution of Customer incidents of services or applications possibly making use of the APIs of the Services	R,A	
Resolution of incidents related to the Operating Systems and the images of the Customer or to the Operating Systems modified by the Customer.	R,A	
Management of Requests and Incidents related to the Services provided by the Provider	I	R,A
Supervision of the Provider physical infrastructure and technical services	I	R,A
Management of Requests for services (cf. Catalogue of Services)	I	R,A

R = Responsible / A = Accountable / C = Consulted / I = Informed

7.6.2 Monitoring of the virtual infrastructure

The Provider monitors its platform and associated Services. The monitoring scope of the Provider covers all of the physical components (routers, firewalls, servers, etc.), software (OpenStack nodes, hypervisors, etc.) and technical services (IAM, LDAP, etc.), with the exception of the virtual infrastructure instantiated by the Customer (cloud server instances, OS, storage).

7.7 Process catalogue

The provision of Services model covers the main processes:

- Incident and Problem management;
- Request management
  - Standard Request
  - Non-standard Request;
- Production release management ;

- Change management

7.7.1 Incident management

This section describes the steps which must be respected by the Customer for managing Incidents it reports to the Provider.

7.7.2 Incident Report

7.7.2.1 Incident reported by the Customer to Technical Support

Technical Support implements the Incident management process only when one of the following cases occurs:

- The Customer has subscribed to a support solution
- The Customer creates a Ticket from its Cloud Customer Space, section “Need Help?”
- The Customer sends an e-mail to the Provider if it cannot access its personal space (PREMIUM solution).
- The Customer calls Technical Support to report a critical or major Incident under a BUSINESS or PREMIUM support plan only;

The prerequisites for opening an Incident Ticket are as follows:

- The Incident regards the Services only; for example, an Incident does not therefore relate to the architectures, software, the production applications or the services designed and operated by the Customer on the Services.
- The Customer has implemented the means to resolve the Incident made available to it, in particular through the use of the Services (among other things the APIs of the Services) and has not been able to resolve the Incident (cf Chapter 5 The Customer’s competencies and responsibilities)
- Only the Named Contacts designated by the Customer may contact the Provider’s technical support;
- The Customer has collected all information relating to the Incident;
- The Customer completes the mandatory fields in its Ticket in the Ticket tool of the Cloud Customer Space following he Provider instructions.

7.7.2.2 Specific process in the event of an Incident reported by phone or by email

Information reported by email or phone must be complete and accurate.

In order to improve the reporting process and allow the Provider to process the incident, the Customer undertakes to confirm the Incident to Technical Support using the Incident reporting model a model of which is contained in the Annex to this document. Failing this, the Provider will be unable to process the Customer’s request.

In order to be processed by the Provider, the Tickets or, in the case of PREMIUM SUPPORT, emails sent by the Customer’s support team must contain the following information:

- Email address of the administrator of the Customer account
- Subject: “[Ticket number issued by the Customer’s ticketing software] Description of the Incident”.
- Body: the required information is described below in the “Ticket reporting model”

Ticket reporting model		
	Mandatory	Description
> Customer information		
# Information about the reporting party		
Customer company name	Yes	
Name of the reporting party and position	Yes	The person reporting the Incident (e.g. a Named Contact)
Reporting party’s email address	Yes	
Reporting party’s phone number	Yes	
# Information about the user		
User's last name / first name	Yes	The person experiencing the incident (e.g. an end Customer, a user of the Customer, etc.)
User’s email address	Yes	
User’s phone number	No	
> Information about the Ticket		
# Context		
Ticket Category: request, incident, problem, information	Yes	Allows the support team to better assess the purpose of the ticket

Criticality perceived by the Customer: critical, major, minor	Yes	Allows the support team to better assess the urgency of processing the ticket
Tenant ID	Yes	Important key allowing the Support team to investigate the incident.
ID of the resources concerned by the incident	Yes	For example: InstanceID, routerID, etc.
Orange Service Cloud for Business concerned	Yes	Compute or Storage Object Services
Impacted component	Yes	Compute Services: Authentication, network, storage block Storage Object: authentication, file upload/download
API Login	Yes	
API URL	Yes	URL used to access the API
Source IP address accessing the API	Yes	The IP address which accesses the API
# Description		
Incident description	Yes	Explanation of the incident context
Date and time of the incident	Yes	
Error message	If displayed	
Error code	If displayed	
File(s) attached		
Screenshot of the incident	If relevant	
Logs relating to the incident	If relevant	

Once the Ticket has been created by the Customer, the Provider support teams inform the Customer's Named Contact by issuing the Ticket. The exchange between the Customer and Technical Support then continues through the Ticket management tool of the Cloud Customer Space.

In the event of an exchange by email or by phone, a Ticket will be created and the follow-up will be through this Ticket.

The Technical Support team will define the appropriate priority level for the Incident created, based on the latter's criticality. In this respect, Technical Support reserves the right to change the priority of an Incident reported by the Customer.

7.7.3 Processing of Incidents

Once the Ticket has been created, the Customer is informed by email of the processing status of the Incident.

The Provider undertakes to record it and report back initially to the Customer within the time frames that depend on the criticality of the Incident (according to the "Table of priority of Incidents" in this document) and the Support Solution subscribed to by the Customer. The Mean Time to Investigate (MTTI) commitments are described in the **Support Plans Table** of this document.

It is understood that any Ticket not falling within the support scope provided by the Provider will be automatically closed by it.

In the event of a disagreement about the priority level of an Incident, the Provider will apply the re-prioritisation procedure described in the "Change to the severity of an Incident Ticket" paragraph.

The Customer may request the Provider to provide it with support through the Flexible Engine Console or the APIs of its tenant. In this case, it is the responsibility of the Customer to invite the person from the Provider support temporarily into its tenant.

The person from the Provider support invited by the Customer will temporarily have the same rights as the Customer over the configuration of its tenant.

When the Provider has completed its intervention, it will be the Customer's responsibility to delete the invitation of the person from the Provider support.

7.7.3.1 Incident Resolution

Once the solution to resolve the Incident has been found (workaround, patch, etc.), the Technical Support team will change the status of the Ticket from "In progress" to "Resolved" and the Customer will be notified of this in the Ticketing tool.

- If the resolution of the Incident requires the release of an application code: the Incident Ticket will be considered to have been resolved once the code is available in the pre-production environment, pending deployment to the production environment. In the event of an incident so requiring it, the urgent release process will be applied (see "Urgent Change" paragraph).
- If the resolution does not require the release of an application code, the Incident will be considered to have been resolved once the solution has been applied in the production environment.

7.7.3.2 *Incident closure*

When the resolution is confirmed by the Customer, the Ticket will be closed by the Technical Support team.

Failing confirmation by the Customer within five (5) working days from the Incident resolution, the Ticket will be automatically closed.

7.7.3.3 *Table of priority of Incidents*

The Table of priority of Incidents defines each priority level for any Incident on the Services.

Incident Priority	Definition
P1 - CRITICAL	A Priority 1 Incident means that one of the following APIs is totally and permanently unavailable: - ECS (Elastic Compute Service) API - IMS (Images Service) API - EVS (Elastic Volume Service) API - VPC (Virtual Private Cloud) API - IAM (Identity Authentication Management) API - OBS (Object Storage Service) API - CCE (Cloud Container Engine) API - MRS (MAP Reduce Service) API - RDS (Relational Database Service) API Or that **all the instances and disk volumes attached to the** instances of the tenant concerned are unavailable, and that it is impossible to create new instances and new volumes, in the tenant using the API of the Services. The Provider commits itself to a Guaranteed Restoration Time (GTR) of 4 hours for Customers which have subscribed to a BUSINESS or PREMIUM Support Solution, for the following APIS: - ECS (Elastic Compute Service) API - IMS (Images Service) API - EVS (Elastic Volume Service) API - VPC (Virtual Private Cloud) API - IAM (Identity Authentication Management) API - OBS (Object Storage Service) API - CCE (Cloud Container Engine) API - MRS (MAP Reduce Service) API - RDS (Relational Database Service) API
P2 - MAJOR	A priority 2 Incident means that the Storage Object Service (OBS) and/or the Compute Service (ECS) is available but that it has malfunctions that have a strong impact on its use. This case occurs when one or more elements of the Provider platform (critical or non-critical) significantly deteriorate the service concerned. P2 Incidents are: • Partial unavailability of the service concerned. - For example: authentication failures with the IAM API which occur intermittently but recurrently. • Total unavailability of one or more functions significantly deteriorating the service concerned.

	- For example, Compute Service (ECS): inability to start new instances, but existing instances work. - For example, Storage Object Service (OBS): inability to upload a file in the container, but the already uploaded files are available. - For example: Intermittent loss of network connectivity or demonstrated highly deteriorated quality of the network connections.
P3 - STANDARD	A priority 3 Incident means that the Storage Object Service (OBS) and/or the Compute Service (ECS) is available but that it has malfunctions that have a limited impact on the said service. This case occurs when one or more elements of the Provider platform (critical or non-critical) significantly deteriorate the service concerned, with the essential functions of the service concerned remaining operational. P3 Incidents are: • An application bug that modifies the display of information relating to the use of Services. - For example: Problem of style sheets on the Flexible Engine Console • Partial unavailability of the service concerned caused by a one-off and temporary event and for which full availability of the service concerned can be restored. - For example: restart of a hypervisor or of a firewall. • Partial unavailability of one or several functions that have a limited impact on the use of the service concerned. - For example, Compute Service (ECS): inability to rename a virtual machine. - For example Storage Object (OBS): inability to list files present in a container.
Not covered	The following situations in particular are not covered: • Any Incident that has no impact on the production platform • Any Incident related to the statistical tools • Any Incident related to the ticketing tool • Any platform shutdown caused by a scheduled activity, such as maintenance, or activities or events over which the Provider has no control, such as damage due to a fire or a failure of the network due to the sectioning of a cable. In all cases, the Provider will remain responsible for the actions of its subcontractors.

7.7.3.4 Change to the priority level of an Incident Ticket

The priority level of an Incident Ticket may be changed as follows:

a) Reduction in the priority level of Incident Tickets

When a P1 critical or P2 major Incident; reported by the Customer, does not comply with the definitions of the incident priority table, the Technical Support team may lower the level of priority of the Ticket.

The Customer can provide additional information and details about the incident, its impacts and explain the choice of the priority of the Incident to the Technical Support team (during this period, the status of the ticket will be changed from “In progress” to “On hold”):

- If the Incident matches the level of priority defined by the Customer, the defined level will remain unchanged;
- Failing this, the priority level of the ticket will be lowered.

In the event of a disagreement, the Provider will ultimately define the priority level of an Incident.

b) Increase in the priority level of Incident Tickets

When a P3 standard Incident or a P2 major Incident reported by the Customer, does not comply with the definitions of the incident priority table, the Technical Support team may raise the level of priority of the Ticket.

The Customer can provide additional information and details about the incident, its impacts and explain the choice of the priority of the Incident to the Technical Support team (during this period, the status of the ticket will be changed from “In progress” to “On hold”):

- If the Incident matches the level of priority defined by the Customer, the defined level will remain unchanged;
- Failing this, the priority level of the Ticket will be raised.

In the event of a disagreement, the Provider will ultimately define the priority level of an Incident.

7.7.4 Management of Problems

The management of Problems process is a process that is internal to the Provider and is carried out by the Provider Problems manager. As such, the Provider has no commitment with respect to the Customer over this scope.

As part of the establishment of this process for managing Problems, the Provider makes its best-efforts to:

- Seek the main cause(s) which create the Incident(s);
- Identify the solution that will allow the Incidents to be resolved definitively;
- Make available to the Incident Management Process recommendations and good practice to ensure resolution of the Incidents under the best conditions.

The Provider will be able to initiate the process for managing Problems in the event of P1 critical Incidents, recurring Incidents or Incidents whose root cause is not identified.

- If a Problem can be resolved by a new version of a code, or of an application, etc., it will be processed as part of the change and/or release process.

7.7.5 Release management

7.7.5.1 Release of major versions

Definition of a major version (major code release): Major change to the Services providing a set of new functions that have a strong impact on the Services or change the application architecture in a structural way.

The release schedule is managed by the Provider. Major versions which are released are versions of the Services.

The Customer is informed by the Provider of the release of a major version with notice of ten (10) working days.

For a change of version of an API of the Services that involves an incompatibility with the previous version of the API, the Provider undertakes to inform the Customer with minimum notice of sixty (60) days.

In this event, the Customer will make the necessary updates to ensure compatibility of its services with the Services. If the update is not made, the Provider will not be able to provide the support services.

7.7.5.2 Release of minor versions

Definition of a minor version (minor code release): Any change not described as major. A minor version generally includes application correction elements (bug fixes) and/or minor developments of application components already deployed and/or new functions that have a limited impact on the service concerned.

The Customer is informed by the Provider of the release of a minor version with notice of 24 hours.

7.7.6 Request management

7.7.6.1 Catalogue of service requests

Only the Customer’s Named Contacts will be able to make service Requests to the Provider via the Customer’s personal space or by email, by specifying in the subject of the Request, ”Support request”.

Only the following Requests are possible for Flexible Engine Services:

Request Name	Type of Request	Description	Pre-requisites	Processing time (from recording of the Ticket by the Customer Service)
Account creation request	Standard	Creation of a tenant for an End Customer at the Customer's request	Opening of a request Ticket. All the information necessary	2 working days

			for opening an account must be provided in the request	
Account deletion request	Standard	Deletion of the tenant for an End Customer at the Customer's request	Opening of a request Ticket. The Customer will have deleted all data and stopped all services on the date of the request	10 working days
Request for quota change	Standard	Change to quotas for the Services	Prior authorisation from the Provider The Customer will provide reasons for its request by providing its needs and expected use forecasts	5 working days

7.7.6.2 Standard Request

A standard Request does not require any specific development. It is an operational measure carried out by the Provider support team at the Customer's request

The standard Request categories are integrated into the Provider's catalogue of service requests as indicated in the "Catalogue of service requests" article of this document.

This type of Request is dealt with by the Technical Support team as part of the Requests management process, via Tickets of the "support request" type.

7.7.6.3 Non-Standard Request

This is a Request which is not contained in the Catalogue of service requests as indicated in the "Catalogue of service requests" article of this document. A request may be made by the Customer to the Provider Customer Service but the Provider makes no commitment with respect to processing it.

7.7.7 Change management

7.7.7.1 Request for Change (RFC)

Requests for change will be processed by the Change management process. This process that is internal to the Provider guarantees the proper management of the various changes made on the Provider production platform. As such, the Provider has no commitment with respect to the Customer over this scope.

7.7.7.2 Standard change

A standard change is a pre-approved change which presents a low risk, somewhat current, and which is made according to a procedure or work instruction. A Request for Change (RFC) is not necessary for making a standard change and its approval by the change authorisation committee of the Provider for the Flexible Engine Services is not required.

7.7.7.3 Urgent change

The Provider Services can implement urgent changes (release of an application code, maintenance operation on the infrastructure, etc.) for critical security or maintenance in operational state reasons (to resolve a critical or major Incident),

Urgent changes are reserved for operations intended to correct critical Incidents or to correct major malfunctions which have an immediate or imminent impact on the production services.

These urgent changes are subject to the approval of the urgent change authorisation committee of the Provider (which meets when the situation so requires).

7.7.7.4 Normal change

Any change not classified as standard or urgent, is considered to be normal. Normal changes are submitted to the Provider's change authorisation committee.

7.7.7.5 Notification of change

The three types of change, detailed above, can be carried out by the Provider through:

- The release of application code;

- A change related to the infrastructure;
- A change related to maintenance.

Information regarding the notice period, the duration of the service interruption and notifications are listed below, divided according to the priority level of these changes.

TYPE OF CHANGE	NATURE OF THE CHANGE	NOTICE PERIOD	SERVICE INTERRUPTION
URGENT	Urgent code release	The Provider inform the CUSTOMER as soon as possible	Maximum 3 hours per month per service concerned
	Urgent corrective maintenance		
NORMAL	Code Release	24 hours	
	Major Code Release	10 working days	
	- Non-urgent corrective maintenance	24 hours	
STANDARD	- Service request included in the service catalogue - Technical tasks with no impact on production	N/A	N/A

8 Service limitations

8.1 Resource quota

In order to guard against misuse or uncontrolled use, the Customer is informed that the Provider sets a maximum quota of resources for use of the Services. This quota may be adjusted at the Customer's request subject to acceptance by the Provider. This quota can be viewed in the Flexible Engine Console.

8.2 Backups

It is the Customer’s responsibility to perform backups of its virtual machines and its data. Unless a specific service has been agreed, the Flexible Engine Services do not include systematic backups by the Provider. Therefore, the Provider may not be held responsible for reconstructing data as part of the Flexible Engine solution.

9 APPENDIX 1: HDS Responsibilities

The purpose of this Appendix is to define the responsibilities of the Customer and the Provider to be compliant with HDS regulations. It is intended to be used in conjunction with the Flexible Engine Service Description which defines Flexible Engine HDS.

9.1 Matrix of responsibilities of the Customer and of the Provider

	Rule	Detail of the rule	Responsibility	Customer Responsibility	Provider Responsibility
Respect for the rights of individuals	Presentation of the services to the people concerned	To present the modalities allowing to ensure the absence of opposition for a legitimate reason of the interested party to the hosting of the data concerning him	Customer	The Customer must obtain or ensure the absence of opposition for a legitimate reason from the natural persons ("patients") whose personal health data are being processed. These persons must be clearly informed of the following - That the hosting of their data is subcontracted to the Provider. - On the modalities of access and transmission of these data, being reminded that the Provider will not ensure these obligations.	N/A
	Individuals' consent to access data	To present the methods used to ensure that access to personal health data and their possible transmission only take place with the agreement of the persons concerned and by the persons designated by them.	Customer	The Customer is solely responsible for access to business applications and personal health data that it may give to its staff or health professionals or directly to patients.	N/A
	Rights of individuals under the GDPR	To make available the procedures and means to respond to requests to exercise the rights of data subjects defined by Articles 15 to 22 of the RGPD (access, rectification, erasure, limitation of processing, portability, opposition)	Shared	The Customer is responsible for the process of taking into account requests to exercise the rights of data subjects as defined by Articles 15 to 22 of the GDPR	The Provider undertakes to make available the procedures and means to enable its Customers to respond to requests to exercise the rights of the persons concerned. The rights covered are those defined by Articles 15 to 22 of the RGPD (access, rectification, erasure, limitation of processing, portability, opposition).
	Notification in case of disclosure of personal data	A procedure must define the modalities for notifying the client of a transmission following judicial seizures, unless such notification is prohibited	Shared	The Customer is responsible for the notification process in case of disclosure of personal data	The Provider undertakes to set up a procedure for notifying the customer in the event of the disclosure of personal data in the context of a judicial seizure, unless such notification is prohibited.

	Rule	Detail of the rule	Responsibility	Customer Responsibility	Provider Responsibility
	Purpose	The Host processes personal data only on the basis of documented instructions from the customer and must not deviate from the purposes specified in the instructions	Shared	The Customer agrees not to process personal data provided by patients for any purpose other than those stated in the consent	The Provider undertakes not to process the personal data provided by its clients for purposes other than those provided for the performance of the client's contract.
	Personal information management	The host must have defined and formalized a policy for the provision and return of personal data to its customers, as well as their destruction	The Provider	N/A	The Provider has a reversibility procedure for the return of personal health data in the event of termination of the contract or withdrawal of certification.
Access control	Identifiers and authorizations	Access to personal data or to the systems used for their processing must be made through nominative accounts	Shared	The Customer is responsible for the implementation of technical means intended to ensure identification, authentication and access control to health data for institutions and health professionals, for the persons concerned by the hosted data, and for other actors (e.g. publishers or integrators) at the level of the hosted Customer applications.	Within the framework of an IaaS offer, The Provider teams do not have access to health data. With regard to the administration of the virtualization infrastructure, The Provider undertakes to maintain an updated register of authorized persons. Used credentials should not be reassigned once they have been deactivated or expired.
	Stakeholder access to systems	Technical means implemented to ensure the identification and authentication and the means of access control of the participants on the systems.	Shared	The Customer is responsible for the access of its stakeholders to the operating systems, software and applications that the Customer manages, as well as for access control on the workstations of these stakeholders.	The Provider shall provide an authentication device to Customer's operators that it has nominally authorized, with the rights defined within the Flexible Engine service.
Telecommunications	Encryption of personal data transmitted over public networks	Personal data must be encrypted before being transmitted over public networks	Customer	The Customer undertakes to implement the technical means necessary to ensure state-of-the-art encryption when technically possible for any transmission of personal health data over public networks.	N/A
	Traceability and integrity of communications	The host must log the transmission of Personal Data to third parties and ensure that the data is received by the target system.	Customer	The Customer is responsible for the integrity of the personal health data during transfers under its responsibility	N/A

	Rule	Detail of the rule	Responsibility	Customer Responsibility	Provider Responsibility
Traceability	Trace management	The hosting company must implement the means to ensure the traceability of user actions, failures and events related to information security	Shared	The Customer is responsible for the traceability of all activity both at the system and application level and within its virtualized environment. Customer is responsible for implementing the Cloud Trace Service (6.9.2) to store activities related to the management of Flexible Engine services.	The Provider collects the traces of the activities performed on the Cloud Store client interface, as well as the traces of the administration activities of the Provider Infrastructure teams.
	Communication of administration traces to clients	Technical and organizational means must be implemented in order to communicate to the client the traces of the administrators	Provider	N/A	The Provider may provide the Customer with traces of its activity on the Cloud Store interface, as well as traces of infrastructure administration activities if requested.
	Entitlements required for access to application traces	Management of authorizations for access to application traces	Customer	The Customer is responsible for the nominative authorization of each operator authorized by him to access the backup and archiving of the traces under his responsibility.	N/A
	Perimeter related to health data access monitoring	Technical perimeter in terms of access to health data	Customer	Customer shall monitor the software and applications it manages and the security, administration and operational components of such software and applications. The traceability of actions on the software installed by the Customer is the responsibility of the Customer.	N/A
Incident management	Alert and escalation procedures	Alerting and escalation procedures to respond to detected security incidents.	Shared	The Customer defines and implements alert and escalation procedures relating to security incidents detected on the virtual infrastructures under its management.	The Provider shall only handle incidents related to the platform and associated Services, with the exception of virtual infrastructures instantiated by the Customer (cloud server instances, OS, storage).
	Data breach notification	The host notifies its customer of any personal data breach as soon as possible after becoming aware of it	Shared	The Customer undertakes to notify any personal data breach in accordance with the requirements set out in the GDPR	The Provider undertakes to notify the Customer of any breach of personal data as soon as possible after becoming aware of such breach.
	Incident classification	Principles of incident classification	Shared	The Customer defines a classification of incidents detected on the virtualized environment under its management.	The Provider defines a priority matrix for incidents detected within its scope of responsibility.
Backup	Backup security	The security of health data backups must be guaranteed, regardless of the medium.	Customer	Backups of health data are controlled during transfer and storage	N/A

	Rule	Detail of the rule	Responsibility	Customer Responsibility	Provider Responsibility
	Backup of traces	Definition and implementation of a trace backup policy	Shared	The Customer is responsible for the backup and archiving policy of the access traces on its virtualized environment.	The Provider shall be responsible for the policy of saving and archiving access traces on the perimeter of the virtualization infrastructure, with the exception of the virtualized environment instantiated by the Customer
Continuity of service	Continuity and availability of services	Service continuity and availability requirements are identified and approved jointly by Customer and Hosting Company. The Hosting Company defines and implements a service continuity plan to meet the defined and approved requirements. The service continuity plan is tested regularly (at least once a year) and the test results must be recorded Capacity and performance are monitored regularly.	Shared	The Customer defines and implements a Business Continuity Plan on the perimeter of its virtualized environment that it instantiates (cloud server instances, OS, storage).	The Provider shall define and implement a Business Continuity Plan for the perimeter of the platforms and associated Services, with the exception of the virtualized environment instantiated by the Customer (cloud server instances, OS, storage), and shall perform recovery tests in accordance with a defined annual plan. The Provider maintains a capacity plan for the platforms to meet the expected evolutions.
Evolution management	Planning new or modified services	Formalizes a planning procedure for the implementation or modification of services to meet service requirements.	Provider	N/A	The Provider maintains an ISO/IEC 20000-1:2011 certification on The scope of the platforms covers all activities related to the hosting of personal health data.
	Change management process for applications	Description of the change management process related to system evolutions for application components	Customer	The Customer shall install, update and maintain the virtualized environment under its management. The Customer defines and implements the change management processes on this perimeter.	N/A
Compliance	ISO 27001 certification	ISO/IEC 27001:2013 certification on the scope of personal health data hosting activities.	Provider	N/A	The Provider maintains an ISO/IEC 27001:2013 certification on the perimeter of the platforms covering all activities of hosting personal health data.
	WSIS scope	The scope of the ISMS must cover all of the host's personal health data hosting activities	Provider	N/A	The scope of the ISMS selected for ISO 27001 certification covers all activities related to the hosting of personal health data.
	WSIS Applicability Statement	The ISMS LoA (Statement of Applicability) must include all the requirements of the HDS certification standard	Provider	N/A	The ISMS LoA (Statement of Applicability) selected for ISO 27001 certification includes all the requirements of the HDS certification framework.

	Rule	Detail of the rule	Responsibility	Customer Responsibility	Provider Responsibility
	Application audit	The hosting company must allow its customers to perform audits on the applications put into production	Provider	N/A	The Provider undertakes to implement a procedure for the performance of audits requested by the Customer on the applications put into production.
	Subcontracting	The host must not use a subcontractor without the prior information of the customer	Provider	N/A	The Provider undertakes to inform the customer in advance in the event that a subcontractor is used.
	Confidentiality agreements	The employment contracts of the host's employees must include a confidentiality clause	Provider	N/A	The Provider undertakes to ensure that its staff and subcontractors are bound by a commitment of confidentiality regarding the dissemination of personal data.
	PGSSI-S compliance	The hosting company must inform its customers that they are required to comply with the PGSSI-S and must put in place a means to collect the commitment of this compliance.	Shared	The Customer agrees to comply with the PGSSI-S (General Security Policy for Health Information Systems)	The Provider informs its Customer that it is required to comply with the PGSSI-S. The signature of this contract is the collection of the customer's commitment
	Places of accommodation	The hosting company must specify the list of all countries in which the customer's data is or can be hosted	Provider	N/A	The Provider agrees to provide the Customer with a list of all the countries in which the data is or may be hosted, and to allow the Customer to choose the country(ies) in which its health data will be hosted.
	Certification audit report	The Hosting Company shall provide certification audit reports to Customers upon request.	Provider	N/A	The Provider undertakes to communicate the certification audit reports to the Customer upon request.
Security policy	Monitoring Policy	Definition and implementation of a policy for monitoring user access to the application	Customer	The Customer shall implement a policy of monitoring accesses by users of the application (including healthcare professionals and institutions, and even directly to patients). This monitoring is a functional matter and the definition and implementation of measures for the preservation, access and use of said traces is outside the scope of The Provider's responsibility.	N/A
	Business application access authorization policy	Definition and implementation of a policy for access to business applications	Customer	The Customer is solely responsible for the access authorization policy to business applications and personal health data that it may give to its staff (operators or others) or to health professionals or even directly to patients.	N/A

	Rule	Detail of the rule	Responsibility	Customer Responsibility	Provider Responsibility
Working post	Workstation safety	Definition and implementation of a security policy for workstations	Shared	The Customer is responsible for the security policy of its operators' workstations (e.g.: management of passwords, tokens, automatic locking)	The Provider is responsible for the security policy of its operators' workstations

It is the responsibility of the Customer to ensure that all of its responsibilities are met.

Where applicable, the Provider may request proof from the Customer of strict compliance with all or part of these provisions if its liability were to be incurred.

9.2 Compliance with the interoperability and security standards of the ANS

The Customer guarantees to have implemented the technical means to ensure the identification, authentication and access control means of the professionals and persons concerned by the health data, for any operation provided for by the hosting service under this Contract.

The Customer undertakes to comply with the PGSSI-S and guarantees that, pursuant to Article L1470-5 of the Public Health Code, the information systems used comply with the interoperability and security reference frameworks drawn up by the public interest grouping mentioned in Article L. 1111-24 of the Public Health Code.

9.2.1 Responsibilities of the Provider

The Provider guarantees the Customer that it has the certification required by the provisions of the Public Health Code for the hosting of personal health data.

The Provider having received any information, documents or data whatsoever, either prior to the Purchase Order or during its execution, may not, without authorization, communicate them to persons other than those who are entitled to know them.

The Provider is obligated to respect and preserve the confidentiality and security of the personal data that it may process. Therefore, the Provider undertakes to take technical and organizational security measures, taking into account the nature of the data and the risks presented by the hosting of the data, in order to preserve the security and integrity of the data and, in particular, to prevent the destruction, loss, alteration, disclosure or unauthorized access.

The Provider hereby declares that it has made its personnel and any subcontractors aware of these stipulations. In accordance with the Personal Data Protection article of the General Conditions *** REF? *** , the Provider undertakes not to process the personal health data it hosts for purposes other than the performance of the FE HDS service. The Provider specifically prohibits any use of this data for marketing, advertising, commercial or statistical purposes.

Section 9.1 presents the responsibility matrix between the Proivider and the Customer.