



Service Level Agreement

Cloud Avenue

1. Definitions

Complementary to the definitions as per General Terms and Conditions and "Cloud Avenue" Service Description, the following specific definitions shall apply with respect to this Service Level Agreement.

Availability Rate, unless otherwise stated for a specific Feature, refers to the rate defined by the following formula:

$$\text{Availability Rate} = \frac{(t_{month} - t_{downtime} + t_{exclusion})}{t_{month}}$$

where:

- t_{month} is the time during which the concerned Feature is subscribed for the month involved
- $t_{downtime}$ is the Downtime of the concerned Feature for the month involved
- $t_{exclusion}$ is the Downtime of the concerned Feature for the month involved during which quality-of-service commitments of the Provider are not applicable as per contractual provisions.

Availability Rates are measured on a calendar month basis. The Availability Rate of a Managed Function is the average of the Availability Rates of the Service Units that constitute it.

Backup Service refers to VM backup and restore mechanisms. When a VM carries a backup policy, it is deemed "protected".

Compute Services refers to the Virtual Datacenter (vDC) and vCenter on Demand (vCoD) Features regardless of the class of service.

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. A downtime is logged by a ticket with Priority P1.

Fault Repair Time refers to the time elapsed between an Incident ticket's opening and its resolution, minus the periods during which the Provider's engagements do not apply.

General Terms and Conditions refers to the Provider' general terms and conditions for Cloud Services.

Guaranteed Availability Rate (or **GAR**) refers to the Availability Rate below which the Provider undertakes to pay penalties to the Customer, for a given Managed Function, in accordance with the level of support to which the Customer has subscribed.

Guaranteed Fault Repair Time (or GFRT) 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.

Guaranteed Response Time (or GRT) 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.

HA Dual Room (High Availability Dual Room) refers to an architecture in very high availability distributed over two rooms of the Datacenters of Val de Reuil.

HA Dual Site (High Availability Dual Site) refers to a very high availability architecture distributed over two Datacenters (Val de Reuil and Chartres).

Portal Services refers to all portals used by the Customer to manage the Cloud Avenue Service and its Features.

Priority refers to the following levels used by the Provider to classify Incident tickets:

- **Priority 1 (or P1)**: complete loss of Service for multiple 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.
- **Priority 3 (or P3)**: Services provided with delay or minor difficulties. 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 not applicable.

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's engagements do not apply.

2. Purpose of the document

The purpose of the present Service Level Agreement is to set forth the conditions on which the Provider commits on quality-of-service for the "Cloud Avenue" Service, in application of General Terms and Conditions, and the Cloud Specific Terms.

3. Conditions of application

3.1. The quality-of-service commitments described in article « Commitments and penalties » of the present document apply in accordance with General Terms and Conditions.

3.2. Limitations specific to each Feature, if any, are specified in article « Commitments and penalties ».

3.3. Under no circumstances may Service Credits granted to the Customer exceed 15% of the monthly recurring charge invoiced to the Customer for the Service in question and month in question.

3.4. The same incident shall not entitle to Service Credits under both the GAR and the GFRT. If both are applicable, the Customer will be granted the higher of the two Service Credits.

3.5. The service quality commitments do not apply in the cases excluded by the General Terms and Conditions, and in the following cases:

- Failure to comply with the resource limits allocated to a VM specified in the Service Description.
- The VMware tools versions installed on the Customer's VMs are not up to date.
- The Customer's application processes conflict with the backup processes at the time of their execution.

4. Commitments and penalties

4.1. Portal Services

4.1.1. Commitments

For Portal Services: the Provider is committed to ensuring an Availability Rate of 99.80% for each portal. The list of portals concerned is as follows:

- VCenter APIs and interface
- The vCloud Director Portal (VCD) and its APIs: technical portal to access all the settings of the organization and its vDCs; one portal per Datacenter.
- The Container Service Extender (CSE) component and its APIs, presented in the VCD portal
- The vRealize Operations (vROPS) component and its APIs, presented in the VCD portal
- vCloud Director Availability (vCDA) and its APIs, presented in the VCD portal
- vRealize Automation (vRA) and its APIs
- The NetBackup Self-Service portal: technical portal to manage backups and restores of VMs; one portal per Datacenter.
- Cloud Avenue APIs

4.1.2. Availability Rate Calculation

The availability of the portals is measured by the Provider via external probes that check availability every 5 minutes.

4.1.3. Penalties

If during a month, the Availability Rate of at least one portal is strictly less than 99.80% and greater than or equal to 98.0%, the Provider undertakes to issue a Service Credit of €50 excluding tax for the relevant month.

If during a month, the Availability Rate of at least one portal is strictly below 98.0%, the Provider undertakes to issue a service credit of €100 excluding tax for the relevant month.

4.2. Compute Services

4.2.1. Virtual Datacenter (vDC)

a) Commitments

For Compute Services: the Provider is committed to ensuring an Organization's vDC Availability Rate according to the availability class:

vDC Availability Class	Availability Rate Commitment
One Room	99,95%
Dual Room	99,99%

b) Availability Rate Calculation

The Interruption of Compute Services of a vDC is characterized by the combination of the following conditions:

- Loss of external connectivity and/or permanent inaccessibility.
- Inability to create new VMs in the vDC, via the portal or API, for any reason other than reaching the limit of the vDC's capacity in GHz, RAM or storage.
- vDC storage unreachable.
- Inability to start VMs created in the vDC.

c) Penalties

If during a month, the Compute Services Availability Rate of a Customer vDC is strictly lower than the contractual Availability Rate commitment, the Provider undertakes to issue a Service Credit equal to the percentage mentioned in the table below of the total amount, excluding taxes, of the consumption of Compute Services associated with the vDC concerned for the relevant month.

vDC Availability Class	Availability Rate	Percentage of Service Credit
One Room	99,95% or more	0%
	from 99% included to 99,95% excluded	10%
	Strictly under 99%	25%
High Availability Dual Room	99,99% or more	0%
	from 99,95% included to 99,99% excluded	10%
	strictly less than 99,95%	25%

4.2.2. vCenter On Demand

a) Commitments

For the **vCenter On Demand** Service, the Service Provider undertakes to ensure a Service* Availability Rate according to the availability class:

vCoD availability class	Availability Rate Commitment
Single Site	99,95%
HA Dual Room or Dual Site	99,99%

*vCoD service is defined as a set of compute, storage and network components.

b) Availability Rate Calculation

The vCoD service is considered available as long as there is no impact on the availability of VMs and their storage, even if one or more nodes are unavailable (failure or maintenance).

The Availability Rate calculation excludes scheduled maintenance periods.

The capacity management performed by the customer must respect a ceiling (vCPU and RAM) in order to allow the shutdown of one of the nodes of the cluster (maintenance and fault tolerance).

c) Penalties

If during a month, the Customer's vCoD Service Availability Rate is strictly lower than the contractual Availability Rate commitment, the Service Provider undertakes to issue a Service Credit equal to the percentage mentioned in the table below of the total amount, excluding taxes, of the consumption of the Compute Services associated with the vCoD concerned for the month in question.

Availability class of the vCoD service	Availability rate	Percentage of Service Credit
One Room and Single Site	99,95% or more	0%
	from 99% inclusive to 99.95% excluded	10%
	strictly less than 99%	25%
Dual Room or Dual Site , with a minimum of 6 servers equally distributed between the two locations	99,95% or more	0%
	from 99% inclusive to 99.95% excluded	10%
	strictly less than 99%	25%

4.3. QoS Appliance

4.3.1. Commitments

For the QoS Appliance service, the Provider is committed to ensuring an Availability Rate of 99.95%.

4.3.2. Availability Rate Calculation

The availability of QoS Appliances is measured by the Provider via internal monitoring probes.

4.3.3. Penalties

If during a month, the availability rate of a Customer's QoS Appliance Service is strictly less than 99.9%, the Provider undertakes to issue a Service Credit equal to the percentage mentioned in the table below of the total amount excluding taxes of the consumption of the QoS Appliance Service associated with the Organization concerned for the month in question.

Availability Rate	Percentage of Service Credit
99,9% or more	0%
from 98% included to 99,9% excluded	20%
strictly less than 98%	30%

4.4. Backup Service

4.4.1. Commitments

For the Backup Service: the Provider is committed to ensuring an Availability Rate of 99.9%.

4.4.2. Availability Rate Calculation

The Backup Service Availability Rate is represented by the success rate of backups, which is defined as follows:

$$\text{Backup Success Rate} = \frac{\text{number of successful backup}}{\text{total number of backups run}}$$

This calculation is valid for all the Organization's "protected" VMs and for a given month.

Any backup made in a case excluded in article 3, or concerning a VM exceeding the size limit authorized in the Service Description is deemed successful.

In any event, in the event of no backups within the month, the Availability Rate is deemed to be 100%.

4.4.3. Penalties

If during a month, the availability rate of a Customer's Backup Service is strictly less than 99.9%, the Provider undertakes to issue a Service Credit equal to the percentage mentioned in the table below of the total amount excluding taxes of the consumption of the Backup Service associated with the Organization concerned for the relevant month.

Availability Rate	Percentage of Service Credit
99,9% or more	0%
from 98% included to 99,9% excluded	10%
strictly less than 98%	25%

4.5. Object Storage Service

4.5.1. Commitments

For the Object Storage Service: The Provider is committed to ensuring an Availability Rate of 99.9%.

4.5.2. Availability Rate Calculation

The Object Storage Service Availability Rate is represented by the success rate of S3 API calls, which is defined as follows:

$$\text{S3 API Calls Success Rate} = \frac{\text{number of successful S3 API calls}}{\text{total number of S3 API calls}}$$

This calculation is valid for the entire storage account that is the object of the Organization and for a given month.

4.5.3. Penalties

If during a month, the availability rate of a Customer's Object Storage Service is strictly less than 99.9%, the Provider undertakes to issue a Service Credit equal to the percentage mentioned in the table below of the total amount excluding taxes of the consumption of the Object Storage Service associated with the Organization concerned for the relevant month.

Availability Rate	Percentage of Service Credit
99,9% or more	0%
from 98% included to 99,9% excluded	10%
strictly less than 98%	25%

4.5.4. Availability Service Level for OpenShift Hypershift Services

The Availability Service Level for the OpenShift Hypershift Service applies to the control plane components managed and operated by Orange as part of the Hosted Control Plane (Hypershift) architecture. It does not cover any Customer-managed components such as the Ingress Routers, or application workloads hosted within the Customer environment. AvScheduled maintenance windows and unavailability caused by Customer actions or infrastructure are excluded from the calculation of Downtime.

Availability Service Level for OpenShift Hypershift Control Plane

Component	Description	Availability Rate
API Server	The OpenShift API endpoint of the Hosted Control Plane is reachable and processes requests normally.	99.90 %
Control Plane Core (scheduler, controller, etcd)	Core components of the Hosted Control Plane are operational and synchronized.	99.90 %
Konnectivity Tunnel	The network tunnel between the control plane and hosted worker nodes is active and stable.	99.80 %

4.6. Guaranteed Fault Repair Time (GFRT)

4.6.1. Commitments

The Provider undertakes to provide the following Fault Repair Time for all services (Compute Services, QoS Appliance, Backup Service, Object Storage Service), excluding portals, depending on the level of support subscribed to, the time being counted only over the periods covered:

Guaranteed Fault Repair Time					
Support level		Standard	Bronze	Silver	Gold & managed tenants*
Severity of the incident	Severity 1	no commitment	8h	8h	4h
	Severity 2	no commitment	no commitment	no commitment	no commitment
	Severity 3	no commitment	no commitment	no commitment	no commitment

*Managed tenants
Customers subscribing to Managed Applications benefit from a dedicated support level, with commitments and obligations identical to those defined under the GOLD Support level.

4.6.2. Calculation of the gap

The "Fault Repair Gap" is calculated as follows for each ticket and each Service Unit concerned:

$$\text{Fault Repair Gap} = \frac{\text{Measured Fault Repair Time} - \text{Guaranteed Fault Repair Time}}{\text{Guaranteed Fault Repair Time}}$$

4.6.3. Penalties

If, for an incident ticket, the Guaranteed Recovery Time is exceeded, the Provider undertakes to issue a Service Credit equal to the percentage indicated in the table below of the monthly recurring amount, excluding tax, invoiced to the Customer for the affected Organization for the relevant month, according to the Recovery Gap observed:

Fault Repair Gap	Percentage of Service Credit
more than 0% to 50%	2%
more than 50% up to 75%	5%
more than 75% up to 100%	10%
More than 100%	25%

4.7. Guaranteed Response Time (GRT)

4.7.1. Commitments

The Provider undertakes to provide the following Response Time for all services (Compute Services, QoS Appliance, Backup Service, Object Storage Service), excluding portals, depending on the level of support subscribed to, the time being counted only over the periods covered:

Guaranteed Response Time					
Support Level		Standard	Bronze	Silver	Gold
Severity of the incident	Severity 1	1 Working Day	2h	2h	1h
	Severity 2	no commitment	no commitment	no commitment	no commitment
	Severity 3	no commitment	no commitment	no commitment	no commitment

4.7.2. Calculation of the gap

The "Response Gap" is calculated as follows:

$$\text{Response Gap} = \frac{\text{Measured Response Time} - \text{Guaranteed Response Time}}{\text{Guaranteed Response Time}}$$

4.7.3. Penalties

If, for an incident ticket, the Guaranteed Response Time is exceeded, the Provider undertakes to issue a Service Credit equal to the percentage indicated in the table below of the monthly recurring amount excluding tax invoiced to the Customer for the Organization affected for the relevant month, according to the Response Gap observed:

Response Gap	Percentage of Service Credit
more than 0% to 50%	2%
more than 50% up to 75%	5%
more than 75% up to 100%	10%
More than 100%	25%