Business

# Technical appendix to Managed Applications Service Description Managed Applications on AWS

## Table des matières

# 1 Overview of the Service

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

- MANAGED BUSINESS APPLICATION ON AMAZON WEB SERVICES
- MANAGED CLOUD NATIVE SERVICES ON AMAZON WEB SERVICES

# 2 Managed Cloud Native Services on AWS

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

### 2.1.1 The cloud native services

One can typically distinguish 3 categories of services:

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

| User plane services | Data services | Other services |
|---|---|---|
| **Compute**<br>□ AWS Elastic Beanstalk<br>□ AWS Lambda<br>□ Amazon Elastic Compute Cloud (EC2)<br><br>**Networking**<br>□ Application Load Balancer<br>□ Route 53<br>□ AWS Network Firewall<br>□ AWS Direct Connect<br>□ Amazon Elastic Load Balancer (ELB) | **Storage**<br>□ Simple Storage Services (S3)<br>□ Elastic Block Store (EBS)<br><br>**Databases**<br>□ DynamoDB<br>□ Amazon RDS<br>□ Amazon DocumentDB<br>□ ElastiCache for Redis<br>□ Amazon MemoryDB for Redis<br>□ Amazon Aurora | **Management & Governance**<br>□ Trusted Advisor<br>□ AWS Backup<br>□ AWS CloudWatch<br>□ AWS Organizations<br><br>**Security management**<br>□ AWS Security Hub<br>□ Amazon Inspector |

**AWS Cloud Native services by category**

### 2.1.2 Tasks involved Cloud Native service management

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

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

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

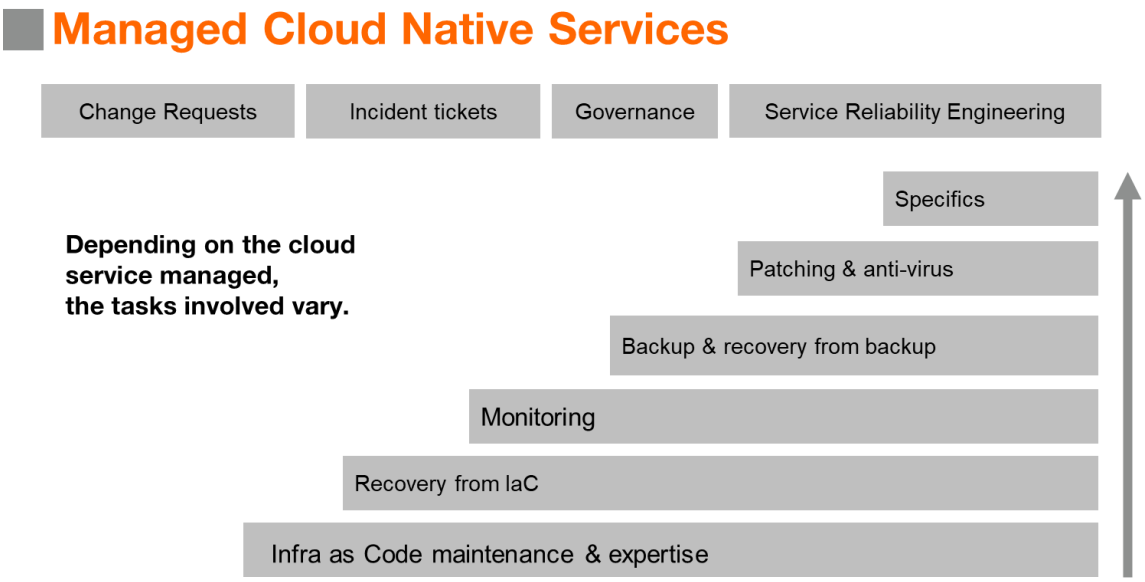


**Tasks involved in managed services for cloud native service**

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

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

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

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

### 2.1.1 Table of tasks involved in the management a Cloud Native service

| AWS service | Type | Configuration | Monitoring and alerts configured in Amazon CloudWatch | Backup configured in AWS Backup | Recovery procedure | Patch management | Antivirus management | Specificities |
|---|---|---|---|---|---|---|---|---|
| **Pre-requisite in case of** | | **Class 2, Class 4 when no AWS backup available for the service** | **Class 2** | **Class 2** | **Class 2 If different from a restore then Class 4, Class 5** | **Class 2** | **n/a** | |
| **Amazon Elastic Compute Cloud (EC2) - per instance** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | AWS Backup | From Backup | AWS Systems Manager Patch Manager | Orange Business Sophos (then TrendMicro) | Only supported OS versions |
| **Elastic Block Store (EBS) - included in Amazon EC2** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | AWS Backup | From Backup | n/a | n/a | Part of managed Amazon EC2 |
| **Auto Scaling - per Auto Scaling Group** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | | | n/a | n/a | Only supported OS versions |
| **Elastic Kubernetes** | Managed | Terraform or AWS | Amazon CloudWatch | n/a | From IaC | n/a | n/a | Patch management |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| **Service (EKS) - per cluster per vCPU** | | CloudFormation | | | | | | is included in the service |
| **Elastic Kubernetes Service (EKS) Fargate - per pod** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | n/a | From IaC | n/a | n/a | |
| **AWS Elastic Beanstalk - per Web Application** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | | From IaC | n/a | n/a | Python, Ruby, Java, .NET, PHP, Node.js, Go and Docker |
| **AWS Lambda - per 100 lines of code** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | n/a | From IaC | n/a | n/a | Customer to provide AWS Lambda code |
| **AWS Key Management Service (KMS) (mutualized HSM)** | Managed | Terraform Plan or AWS CloudFormation | Amazon CloudWatch | | From IaC | n/a | n/a | Natively Redundant |
| **Amazon Elastic Load Balancer** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | n/a | From IaC | n/a | n/a | |
| **Application Load Balancer - per Application Load Balancer** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | | From IaC | n/a | n/a | |
| **Route 53 - Per zone** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | n/a | From IaC | n/a | n/a | |
| **Cloud Front - per End Point** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | n/a | From IaC | n/a | n/a | |
| **AWS Direct Connect** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | n/a | Configured once, need to re-configure if an incident occurs | n/a | n/a | On-premises connection/routing is excluded from SoW |
| **Network ACLs** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | n/a | From IaC | n/a | n/a | Recovery from IaC is sow |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| **AWS Network Firewall** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | n/a | From IaC | n/a | n/a | Recovery from IaC is sow |
| **AWS WAF (web application firewall)** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | n/a | From IaC | n/a | n/a | Recovery is sow |
| **VPN Gateway - per connexion (cloud side MS only - link and e2e excluded)** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | n/a | From IaC | n/a | n/a | e2e link excluded, SIC required on top export is sow (e.g. Preshared keys) |
| **Network Security Groups - per 5 security groups with limitation of number of rules inside the security group** | Change mgt | Terraform or AWS CloudFormation | n/a | n/a | From IaC | n/a | n/a | |
| **VPC (up to 5) - included in managed tenant** | Change mgt | Terraform or AWS CloudFormation | n/a | n/a | From IaC | n/a | n/a | |
| **Simple Storage Services (S3)** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Optional: AWS Backup | From Data backup | n/a | n/a | SoW necessary for data backup |
| **DynamoDB - Dynamo DB table** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in continuous backup and on demand backup using AWS backup or Dynamo DB | From Backup | n/a | n/a | Execution of script provided by customer sow |
| **DynamoDB - per additional Dynamo DB table** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in continuous backup and on demand | From Backup | n/a | n/a | Execution of script provided by customer sow |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | | | | backup using AWS backup or Dynamo DB | | | | |
| **Amazon Aurora MySQL Compatible - per DB instance** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in continuous backup, snapshot using AWS Backup and backtrack | From Backup | IAC | n/a | Execution of script provided by customer sow |
| **Amazon Aurora MySQL Compatible - per additional DB instance** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in continuous backup, snapshot using AWS Backup and backtrack | From Backup | IAC | n/a | Execution of script provided by customer sow |
| **Amazon RDS - per DB instance** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in automated backup or manual backup | From Backup | IAC | n/a | Execution of script provided by customer sow |
| **Amazon RDS - per additional DB instance** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in automated backup or manual backup | From Backup | IAC | n/a | Execution of script provided by customer sow |
| **Amazon Aurora PostgreSQL Compatible - per DB instance** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in continuous backup and snapshot using AWS Backup | From Backup | IAC | n/a | Execution of script provided by customer sow |
| **Amazon Aurora PostgreSQL Compatible - per additional** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in continuous backup and snapshot using | From Backup | IAC | n/a | Execution of script provided by customer sow |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| **DB instance** | | | | AWS Backup | | | | |
| **Amazon Neptune - per DB instance** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in automated backup or manual backup | From Backup | IAC | n/a | Execution of script provided by customer sow |
| **Amazon Neptune - per additional DB instance** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in automated backup or manual backup | From Backup | IAC | n/a | Execution of script provided by customer sow |
| **ElastiCache for Redis- per node** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in automated backup or manual backup | From Backup | IAC | n/a | Execution of script provided by customer sow |
| **ElastiCache for Redis- per additional node** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in automated backup or manual backup | From Backup | IAC | n/a | Execution of script provided by customer sow |
| **Amazon MemoryDB for Redis - per node** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in automated backup or snapshot | From Backup | IAC | n/a | Execution of script provided by customer sow |
| **Amazon MemoryDB for Redis - per additional node** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Build-in automated backup or snapshot | From Backup | IAC | n/a | Execution of script provided by customer sow |
| **Amazon Keyspaces (for Apache Cassandra) - per Per Table within Keyspace** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Amazon Keyspaces PITR | From Backup | IAC | n/a | Execution of script provided by customer sow |
| **Amazon Keyspaces (for Apache Cassandra) - per additional** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | Amazon Keyspaces PITR | From Backup | IAC | n/a | Execution of script provided by customer sow |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| **Per Table within Keyspace** | | | | | | | | |
| **ElastiCache for Memcached - per node** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | not available | From IaC | IAC | n/a | Execution of script provided by customer sow |
| **ElastiCache for Memcached - per additional node** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | not available | From IaC | IAC | n/a | Execution of script provided by customer sow |
| **Amazon Quantum Ledger Database- per Table within ledger** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | not available / on-demand journal export feature | From IaC | IAC | n/a | Execution of script provided by customer sow |
| **Amazon Quantum Ledger Database- per additional Table within ledger** | Managed | Terraform or AWS CloudFormation | Amazon CloudWatch | not available/ on-demand journal export feature | From IaC | IAC | n/a | Execution of script provided by customer sow |
| **AWS MQ for Apache ActiveMQ** | Managed | Terraform or AWS CloudFormation | AWS CloudWatch | | From Terraform IaC or CloudFormation | n/a | n/a | |
| **Amazon Simple Queue Service** | Managed | Terraform or AWS CloudFormation | AWS CloudWatch | | From Terraform IaC or CloudFormation | n/a | n/a | |
| **AWS Amazon Simple Notification Service (SNS)** | Managed | Terraform or AWS CloudFormation | AWS CloudWatch | | From Terraform IaC or CloudFormation | n/a | n/a | |
| **AWS MQ for RabbitMQ** | Managed | Terraform or AWS CloudFormation | AWS CloudWatch | | From Terraform IaC or CloudFormation | n/a | n/a | |
| **Cognito** | | Terraform or AWS CloudFormation | | | From Infra as Code for user pools and Identity pools | n/a | n/a | |
| **AWS Directory Service** | Managed | Terraform or AWS | AWS CloudWatch | | From Terraform IaC or | n/a | n/a | |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| **(managed AD)** | | CloudFormation | | | CloudFormation | | | |
| **AWS API Gateway** | Managed | Terraform or AWS CloudFormation | AWS CloudWatch | | From Terraform IaC or CloudFormation | n/a | n/a | |
| **CloudWatch (option) - per managed ressource** | Change mgt | Terraform or AWS CloudFormation | n/a | n/a | From IaC | n/a | n/a | Optional based on change mgt |
| **SSM - per SSM Document** | Change mgt | | | | | n/a | n/a | custom SSM Document configured and maintained by customer |
| **Trusted Advisor** | Use & Change mgt | | | | | n/a | n/a | Used by the Service Reliability Engineer |
| **AWS Backup** | Use & Change mgt | Terraform or AWS CloudFormation | | | Backup policy from Terraform IaC. Data restored from backup. | n/a | n/a | Used by default for backup when available for the resource. Data restore is based on SoW |
| **AWS Code Star / Code Commit / Code Pipeline / Code Deploy / Code Artifact** | Use & Change mgt | | | | | n/a | n/a | SoW |
| **AWS Organizations** | Change mgt | Terraform or AWS CloudFormation | | | | n/a | n/a | |
| **AWS Management Console** | Use & Change mgt | | | | | n/a | n/a | |
| **AWS Security Hub** | SoW | | | | | n/a | n/a | Specific sow |
| **AWS Batch** | SoW | | | | | n/a | n/a | Specific sow |

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

### 2.1.2 Tooling used for cloud native managed services

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

| Process | Tool used by ORANGE BUSINESS MA delivery |
|---|---|
| | |
| Configuration of the infrastructure | Terraform script (Clouds Orange / Multicloud)<br>CloudFormation (optional)<br>GIT referential<br>CI / CD |
| Supervision solution | AWS CloudWatch with connector to Orange Business central Monitoring |
| Backup | AWS Backup (incl snapshots) |
| OS patching solution | AWS Systems Manager Patch Manager<br>Orange Business MA Patching Tool (Orange Business Cloud platforms and Multicloud)<br>Orange Business OS Factory |
| Antivirus solution | Orange Business MA Antivirus Tool Sophos (then TrendMicro) |
| Logging solution | Amazon CloudWatch Logs (on demand) |
| Recovery | From backup when applicable<br>From Terraform script in GitHub when applicable<br>Ideally from up to date Infra as code with CI/CD |
| Admin connectivity | VPN to Orange Business Administration Zone |
| Portal for access to MA contract, incident & change ITSM | Orange Business CloudStore |

### 2.1.3 General pre-requisites to the run of managed services

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

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

- Additional pre-requisites are required when transition is not the entire responsibility of ORANGE BUSINESS (not Full Build, ref to Build chapter of the document)

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

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

- Either based on AWS tools i.e. CloudFormation, Amazon CloudWatch
- Or based on generic multi-cloud tooling proposed by ORANGE BUSINESS e.g. CaasCad (Prometheus, Grafana, …)

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

### 2.1.4 Criteria for the run of a managed cloud native service component

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

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

# 3 The build of services & managed services on AWS

When the build effort is uncertain from pre-sales documentation, an assessment is proposed at the beginning of the build project by ORANGE BUSINESS Cloud Expert Services. During this assessment, the following tasks are performed:

- Collection of the architecture diagrams with dependencies, HLD, LLD of applications, and infrastructure to be managed and any other useful information.
- Check of the inventory of resources to be deployed and managed.
- Review for each of the dependence the remaining work requested to ORANGE BUSINESS for completing the build to reach readiness for the run. Review the criteria for a resource build to

qualify to a given model of build. Hence determining for each resource which build model applies: No build, Backend build, Operations Build or Full Build.

- Confirmation that the pre-required tools for operations are in place (or alternatively agreeing on a specific scope of work for different tooling if agreeable).
- Establishing requested responsibilities defined between the customer and ORANGE BUSINESS (RACI) for build and for the run.
- Identifying potential limitations on the managed application service if criteria are not met.

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

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

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

The build effort provided by ORANGE BUSINESS in the “backend build” includes integrating the alarms from AWS Monitoring to the ORANGE BUSINESS backend systems, capturing the procedural guides provided by the customer into the ORANGE BUSINESS knowledge repository

of operations, and operations readiness. It includes as well getting the administrative backend, the ORANGE BUSINESS ITSM, the portal and billing readiness for operations.

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

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

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

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

- Amazon CloudWatch for supervision with alerts
- AWS Backup configuration and deployment
- Update Manager configuration for VM patching
- Anti-virus configuration for VM
- Use of standard remedial and troubleshooting procedures on known incident for the cloud native service.
- Use of standard recovery procedures for the cloud native service.

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

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

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

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

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

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

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

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

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

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

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

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

## 3.5 Charging model for build

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

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

# 4 Detailed responsibilities and accountabilities

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

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

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

#### *4.1.1.1 RACI for Managed OS*

| Service Implementation | OBS | Customer | OBS | Customer | OBS | Customer |
|---|---|---|---|---|---|---|
| OS Server infrastructure implementation | Full Build | | Operations Build | | Backend Build | |
| Deployment of the infrastructure | R, A | I | I | R, A | I | R, A |
| Deployment of LAN components | R, A | I | I | R, A | I | R, A |

| | | | | | | |
|---|---|---|---|---|---|---|
| Deployment of DNS and NTP services | R, A | I | R, A | I | I | R, A |
| Backup tools for operations (AWS backup & AWS Snapshots) | R, A | I | R, A | I | I | R, A |
| Deployment of the OS patching solution (AWS Systems Manager Patch Manager) or OBS Patch management | R, A | I | R, A | I | I | R, A |
| Deployment of the Antivirus solution | R, A | I | R, A | I | SoW | SoW |
| Deployment of the supervision solution (Amazon CloudWatch) | R, A | I | R, A | I | I | R, A |
| Deployment of the supervision solution (Amazon CloudWatch) | R, A | I | R, A | I | I | R, A |
| Deployment of security groups and firewall rules | R, A | I | SoW | SoW | I | R, A |
| Recovery procedure (Infra as Code, restore, other…) | R, A | I | I | R, A | I | R, A |
| Testing and validation of infrastructure implementation | R | A | I | R, A | I | R, A |
| Testing and validation of AWS tooling implementation and lifecycle management | R | A | R | A | I | R, A |
| OS Server Implementation | | | | | | |
| Evaluation or deployment of the operating system | R, A | I | R, A | I | I | R, A |
| Deployment of new packages | R, A | I | R, A | I | R, A | I |
| Test and validation of operating system implementation for new packages | R, A | I | R, A | I | R, A | I |
| Service implementation documentation | | | | | | |
| Conception, architecture and low-level design for infrastructure | I | R, A | I | R, A | I | R, A |
| Implementation and operation documentation for infrastructure | R, A | I | I | R, A | I | R, A |
| Conception and low-level design for tooling (AWS) | R, A | I | R, A | I | I | R, A |
| Implementation & operation documentation for tooling (AWS) | R, A | I | R, A | I | I | R, A |

#### *4.1.1.2 RACI for Database as a Service*

| **Service Implementation** | **OBS** | **Customer** | **OBS** | **Customer** | **OBS** | **Customer** |
|---|---|---|---|---|---|---|
| Database aaS services conception and implementation | Full Build | | Operations Build | | Backend Build | |
| Maintenance of Infrastructure architecture referential | R, A | I | I | R, A | I | R, A |
| Maintenance of tooling configuration referential | R, A | I | R, A | I | I | R, A |
| Deployment of the infrastructure | R, A | I | I | R, A | I | R, A |
| Deployment of the supervision solution (Amazon CloudWatch) | R, A | I | R, A | I | I | R, A |
| Deployment of the logging solution (Amazon CloudWatch Logs) (optional) | R, A | I | R, A | I | I | R, A |
| Deployment of the backup solution (AWS Backup, Snapshot) | R, A | C, I | R, A | C, I | I | R, A |
| Recovery procedure for infrastructure from referential (Infra as code, restore from backup, other…) | R, A | C, I | I | R, A | I | R, A |
| Recovery procedure for tooling from referential (Infra as code, restore, other…) | R, A | C, I | R, A | C, I | I | R, A |

| | | | | | | |
|---|---|---|---|---|---|---|
| Testing and validation of infrastructure implementation | R, A | I | I | R, A | I | R, A |
| Testing and validation of tooling implementation and lifecycle management | R, A | I | R, A | C, I | I | R, A |
| Customer provided script execution on DB instance | R | A, I | R | A, I | R | A, I |
| OBS script execution on DB instance | R, A | C, I | R, A | C, I | R, A | C, I |
| Service implementation documentation | | | | | | |
| Conception, architecture, and low-level design for infrastructure | C, I | R, A | I | R, A | I | R, A |
| Implementation and operation documentation for infra | R, A | C, I | I | R, A | I | R, A |
| Conception and low-level design for tooling (AWS) | R, A | C, I | R, A | C, I | I | R, A |
| Implementation & operation documentation for tooling (AWS) | R, A | C, I | R, A | C, I | I | R, A |

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

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

#### *4.1.1.3 RACI for other Native Services managed*

| **Service Implementation** | **OBS** | **Customer** | **OBS** | **Customer** | **OBS** | **Customer** |
|---|---|---|---|---|---|---|
| Native service infrastructure implementation | Full Build | | Operations Build | | Backend Build | |
| Deployment of the infrastructure | R, A | I | I | R, A | I | R, A |
| Backup tools for operations (AWS backup) (1) | R, A | I | R, A | I | I | R, A |
| Deployment of the supervision solution (Amazon CloudWatch) (1) | R, A | I | R, A | I | I | R, A |
| Deployment of the logging solution (Amazon CloudWatch Logs) optional (1) | R, A | I | R, A | I | I | R, A |
| Deployment of security groups and firewall rules | R, A | I | SoW | SoW | I | R, A |
| Recovery procedure (Infra as Code, restore, other…) | R, A | I | I | R, A | I | R, A |

| | | | | | | |
|---|---|---|---|---|---|---|
| Testing and validation of infrastructure implementation | R | A | I | R, A | I | R, A |
| Testing and validation of AWS tooling implementation | R | A | R | A | I | R, A |
| Packages | | | | | | |
| Deployment of new packages (1) | I | R, A | I | R, A | I | R, A |
| Service implementation documentation | | | | | | |
| Conception, architecture, and low-level design for infrastructure | C, I | R, A | I | R, A | I | R, A |
| Implementation and operation documentation for infrastructure | R, A | I | I | R, A | I | R, A |
| Conception and low-level design for tooling (AWS) | R, A | I | R, A | I | I | R, A |
| Implementation & operation documentation for tooling (AWS) | R, A | I | R, A | I | I | R, A |

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

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

# 5 Detailed description per service (Extract)

## 5.1 Amazon API Gateway

### 5.1.1 Description

Amazon API Gateway is an AWS service for creating, publishing, maintaining, monitoring, and securing REST, HTTP, and WebSocket APIs at any scale.

### 5.1.2 Build to run service included in the OTC

#### *5.1.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.1.2.2 Build to run service*

- Refer to generic description.

### 5.1.3 RUN services included in the MRC

#### *5.1.3.1 Run service pre-requisite*

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

#### *5.1.3.2 KPI & alerts*

**Monitoring**

Yes

**KPI monitored**

Metrics supported for Amazon API Gateway:

- 4XXError
- 5XXError
- CacheHitCount
- CacheMissCount
- Count
- IntegrationLatency
- Latency

**Alerts observed**

- 4XXError
- 5XXError
- Latency

#### *5.1.3.3 Backup and restore*

Restore from Infra as Code. No native backup exists for this service.

#### *5.1.3.4 Disaster Recovery:*

No native Disaster Recovery is available for this service.

#### *5.1.3.5 AWS SLA High Availability*

The service is HA by design in a single AWS Region.

#### *5.1.3.6 Limitations & pre-requisite*

Whenever the API is customized, there should be procedures provided by the customer describing how to monitor and troubleshoot the API.

### 5.1.4 Charging model

| Work Unit |
| --- |
| Per API |

### 5.1.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| **Modify API behavior** | On quote |
| **Other changes** | Estimation in tokens based on time spent |

## 5.2 Elastic Load Balancing - Application Load Balancer (ALB)

### 5.2.1 Description

An Application Load Balancer functions at the application layer, the seventh layer of the Open Systems Interconnection (OSI) model. After the load balancer receives a request, it evaluates the listener rules in priority order to determine which rule to apply, and then selects a target from the target group for the rule action.

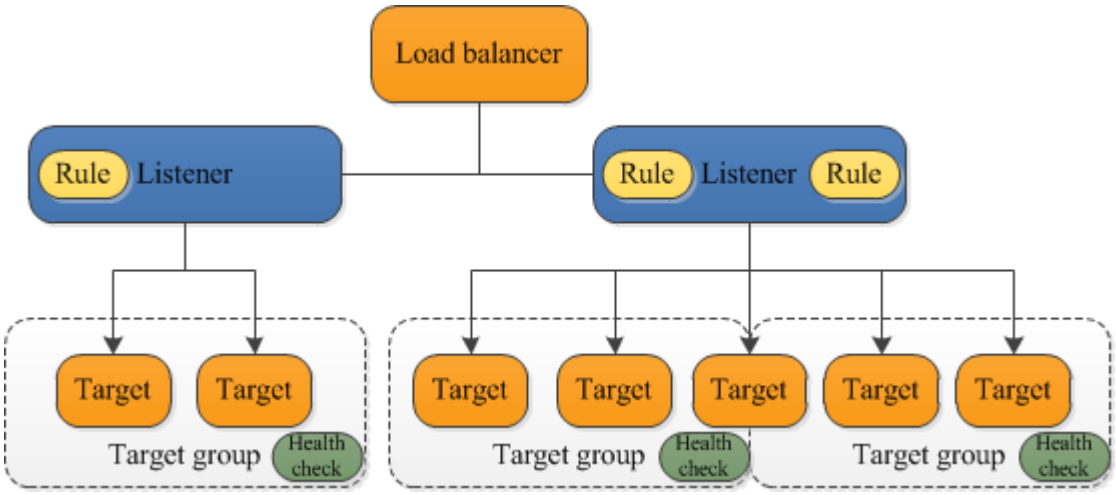


### 5.2.2 Build to run service included in the OTC

#### *5.2.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.2.2.2 Build to run service*

- Refer to generic description.

### 5.2.3 RUN services included in the MRC

#### *5.2.3.1 Run service pre-requisite*

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

#### *5.2.3.2 KPI & alerts*

**Monitoring**

Yes

**KPI monitored**

Metrics supported by Application Gateway:

- ActiveConnectionCount

- ClientTLSNegotiationErrorCount
- ConsumedLCUs
- DesyncMitigationMode_NonCompliant_Request_Count
- DroppedInvalidHeaderRequestCount
- ForwardedInvalidHeaderRequestCount
- GrpcRequestCount
- HTTP_Fixed_Response_Count
- HTTP_Redirect_Count
- HTTP_Redirect_Url_Limit_Exceeded_Count
- HTTPCode_ELB_3XX_Count
- HTTPCode_ELB_4XX_Count
- HTTPCode_ELB_5XX_Count
- HTTPCode_ELB_500_Count
- HTTPCode_ELB_502_Count
- HTTPCode_ELB_503_Count
- HTTPCode_ELB_504_Count
- IPv6ProcessedBytes
- IPv6RequestCount
- NewConnectionCount
- NonStickyRequestCount
- ProcessedBytes
- RejectedConnectionCount
- RequestCount
- RuleEvaluations

**Alerts observed:**

- HTTPCode_ELB_3XX_Count
- HTTPCode_ELB_4XX_Count
- HTTPCode_ELB_5XX_Count (Backend down)
- HTTPCode_ELB_500_Count
- HTTPCode_ELB_502_Count
- HTTPCode_ELB_503_Count
- HTTPCode_ELB_504_Count
- RejectedConnectionCount

#### *5.2.3.3 Backup and restore*

**Data backup and restore**

Can be exported from CI/CD Pipeline.

**Service restore**

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

#### *5.2.3.4 AWS SLA High Availability and Disaster Recovery inter-region*

The service is Highly Available by design by AWS.
There is no native Disaster Recovery for this service.

### 5.2.4 Charging model

| Work Unit |
|---|
| Per Application Load Balancer |

### 5.2.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| **Add/modify Backend** | 1 Token |
| **Certificate Installation** | 2 Tokens |
| **Add a pool member (if members are static)** | 2 Tokens |
| **Modify configuration (routing/Web Application Firewall)** | Estimation in tokens based on time spent |
| **Add a listener** | 2 Tokens |
| **Add a Target Group** | 2 Tokens |
| **Modify listener rule (Re-writing URL)** | Estimation in tokens based on time spent |
| **Other changes** | Estimation in tokens based on time spent |

## 5.3 Amazon Route 53

### 5.3.1 Description

Amazon Route 53 is a highly available and scalable Domain Name System (DNS) web service. You can use Route 53 to perform three main functions in any combination: domain registration, DNS routing, and health checking.

### 5.3.2 Build to run service included in the OTC

#### *5.3.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.3.2.2 Build to run service*

- Refer to generic description.

### 5.3.3 RUN services included in the MRC

#### *5.3.3.1 Run service pre-requisite*

- A referential file exists in the Git used by ORANGE BUSINESS which includes the reference configuration of the Route 53.
- This file can be executed with a CI/CD used by ORANGE BUSINESS and the execution has been tested successfully.

#### *5.3.3.2 Co-manage option*

For the Public Hosted Zones, ORANGE BUSINESS work with the customer for the public domain naming context.

For the Private Hosted Zones, a RACI must be done.

#### 5.3.3.3 Limitations & pre-requisite

For domain registration, we highly recommend that customer handles the intellectual property part and registers the domain wherever he wants. The domain hosting configuration on AWS will be handled by OBS.
ORANGE BUSINESS could be the registrar.
The Customer must be the registrant.

#### 5.3.3.4 KPI & alerts

**Monitoring**

Yes

**KPI monitored**

Global metrics supported by Route 53:

- ChildHealthCheckHealthyCount
- ConnectionTime
- HealthCheckPercentageHealthyHealthCheckStatus
- SSLHandshakeTime
- TimeToFirstByte

Hosted Zone metrics supported by Route 53:

- DNSQueries
- DNSSECInternalFailure
- DNSSECKeySigningKeysNeedingAction
- DNSSECKeySigningKeyMaxNeedingActionAge
- DNSSECKeySigningKeyAge

**Alerts observed**

**Default**

DNSQueries

**Optional**

DNSSECKeySigningKeyMaxNeedingActionAge to detect a security breach

#### 5.3.3.5 Backup and restore

**Data backup and restore**

No data to backup.

**Service restore**

Recovery will be from Infra as Code.

#### *5.3.3.6 AWS SLA High Availability and Disaster Recovery inter-region*

Route 53 is a native Global AWS service. The service is Highly Available by design by AWS. Disaster Recovery is native.

### 5.3.4 Charging model

| Work Unit |
|---|
| Per Hosted Zone |

### 5.3.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| **Create / update Zone** | 2 tokens |
| **Zone delegation*** | 4 tokens |
| **Configure inbound and outbound resolvers** | 8 tokens |
| **Other changes** | Estimation in tokens based on time spent |

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

## 5.4 AWS Lambda

### 5.4.1 Description

Lambda is a compute service that lets you run code without provisioning or managing servers. Lambda runs your code on a high-availability compute infrastructure and performs all of the administration of the compute resources, including server and operating system maintenance, capacity provisioning and automatic scaling, code monitoring and logging.

### 5.4.2 Build to run service included in the OTC

#### *5.4.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.4.2.2 Build to run service*

- Refer to generic description.

### 5.4.3 RUN services included in the MRC

#### *5.4.3.1 Run service pre-requisite*

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

#### *5.4.3.2 Co-manage option*

Yes, based on RACI determined during pre-sales or build jointly with the customer.

#### *5.4.3.3 KPI*

**KPIs monitored:**

**Invocation metrics**

- Invocations
- Errors
- DeadLetterErrors
- DestinationDeliveryFailures
- Throttles
- ProvisionedConcurrencyInvocations
- ProvisionedConcurrencySpilloverInvocations

**Performance metrics**

- Duration
- PostRuntimeExtensionsDuration
- IteratorAge
- OffsetLag

**Concurrency metrics**

- ConcurrentExecutions
- ProvisionedConcurrentExecutions
- ProvisionedConcurrencyUtilization
- UnreservedConcurrentExecutions

**Alerts Observed:**

Errors
Customized alerting can be added as an option based on customer needs.

#### *5.4.3.4 Backup and restore*

**Data backup and restore**

Backup is not used by default.

**Service restore**

By default, the Lambda Function code in the GIT is the referential and the Continuous Deployment chain workflow is used to deploy it. Shall a problem occur on a function, the Continuous Deployment chain is used to redeploy the Function from the version of reference in the GIT.

#### *5.4.3.5 AWS SLA High Availability*

The service is HA by design by AWS. The highly Availability of the business function depends on its design, its interfaces with other business functions and external services and its dependencies on operating systems, middleware, databases, micro-services, Kubernetes services, big data services, and cloud native services mainly event driven services (e.g. SNS, SQS, API Gateway, EventBridge, etc.).

#### *5.4.3.6 Disaster Recovery inter-region*

Redeploy in another region from Infra as Code.

### 5.4.4 Charging model

| Work Unit |
|---|
| Per package of 100 lines of Lambda function code |

### 5.4.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| **Activate / deactivate a function** | 2 tokens |
| **Deploy a new Lambda version** | 2 tokens |
| **Deploy a new Lambda** | Estimation in tokens based on time spent |
| **Develop and deploy a new Lambda** | Estimation in tokens based on time spent |

## 5.5 Simple Storage Service (S3)

### 5.5.1 Description

Amazon Simple Storage Service (Amazon S3) is an object storage service that offers industry-leading scalability, data availability, security, and performance.

### 5.5.2 Build to run service included in the OTC

#### *5.5.2.1 Build service pre-requisite*

- Please refer to generic description

#### *5.5.2.2 Build to run service*

- Build to run service for Simple Storage Service (S3) are necessary. They encompass the parameters setting e.g. Intelligent-Tiering, encryption, versioning, access policies, etc. Optionally, if an optional recurring managed service using S3 has been requested, build to run task will include the selection of KPIs to be observed and alerts to be set up based on KPI thresholds, or external calls to test the availability of the Simple Storage Service (S3). Please refer to generic build to run description.

### 5.5.3 RUN services included in the MRC

Run a managed Simple Storage Service (S3) service is optional. Depending on Customer's interest in monitoring the storage KPIs, in alerting based on KPIs, the Customer may request the service. By default, there is no recurring task proposed on Simple Storage Service (S3), but on demand changes and on demand investigations.

#### *5.5.3.1 Run service pre-requisite*

- A referential file exists in the Git including the reference configuration of the Simple Storage Service (S3).
- This file can be executed with a CI/CD and the execution has been tested successfully.

#### *5.5.3.2 Co-manage option*

Yes

#### *5.5.3.3 KPI & alerts*

**Monitoring**

S3 is monitored through AWS Health Dashboard.

This service is also monitored through the other services using it (CloudFront, Lambda Function, etc.).

#### *5.5.3.4 Backup and restore*

**Data backup**

Optional: Simple Storage Service is a highly available service. Backup is done only through replication in another region. Replication can be added as an option based on customer needs.
If the customer requests the replication in another, it will imply additional fees (network transfer and additional storage).

**Service restore**

Optional: subject to customer having ordered replication in another region.

#### *5.5.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

Yes, by default by AWS.

### 5.5.4 Charging model

| Work Unit |
| --- |
| Per S3 Bucket |

### 5.5.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |
| Change Access Tier | 2 Tokens |
| Other Changes | Estimation in tokens based on time spent |

## 5.6 CloudFront

### 5.6.1 Description

Amazon CloudFront is a web service that speeds up distribution of your static and dynamic web content, such as .html, .css, .js, and image files, to your users. CloudFront delivers your content through a worldwide network of data centers called edge locations. When a user requests content that you're serving with CloudFront, the request is routed to the edge location that provides the lowest latency (time delay), so that content is delivered with the best possible performance.

### 5.6.2 Build to run service included in the OTC

#### *5.6.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.6.2.2 Build to run service*

- Refer to generic description.

### 5.6.3 RUN services included in the MRC

#### *5.6.3.1 Run service pre-requisite*

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

#### 5.6.3.2 Co-manage option

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

#### 5.6.3.3 KPI & alerts

**Monitoring**

Yes
We can optionally configure CloudFront to create log files that contain detailed information about every user request that CloudFront receives.

**KPI monitored**

- Requests
- Bytes downloaded
- Bytes uploaded
- 4xx error rate
- 5xx error rate
- Total error rate

**Alerts observed**

- Requests
- 4xx error rate
- 5xx error rate
- Total error rate

#### 5.6.3.4 Backup and restore

**Data backup and restore**

Can be exported from Infra as Code

**Service restore**

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

#### 5.6.3.5 AWS SLA High Availability and Disaster Recovery

The service is Highly available by AWS (a global service).

### 5.6.4 Charging model

| Work Unit |
| --- |
| per CloudFront Distribution |

### 5.6.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |

| | |
|---|---|
| **Modify Origin** | 2 Tokens |
| **Customize HTTP headers** | 4 Tokens |
| **Modify cache Rules** | Estimation in tokens based on time spent |
| **Specify cache and compression settings** | Estimation in tokens based on time spent |
| **Specify the values to include in origin requests** | Estimation in tokens based on time spent |
| **Specify the HTTP headers to add to viewer responses** | Estimation in tokens based on time spent |
| **Other changes** | Estimation in tokens based on time spent |

## 5.7 AWS Key Management Service (AWS KMS)

### 5.7.1 Description

AWS Key Management Service (AWS KMS) is a managed service that makes it easy for you to create and control the cryptographic keys that are used to protect your data. AWS KMS uses hardware security modules (HSM) to protect and validate your AWS KMS keys under the FIPS 140-2 Cryptographic Module Validation Program, except in the China (Beijing) and China (Ningxia) Regions.

### 5.7.2 Build to run service included in the OTC

#### *5.7.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.7.2.2 Build to run service*

- Refer to generic description.

### 5.7.3 RUN services included in the MRC

#### *5.7.3.1 Run service pre-requisite*

- A referential file exists in the Git including the reference configuration of the AWS Key Management Service (AWS KMS).
- This file can be executed with a CI/CD and the execution has been tested successfully.

#### *5.7.3.2 KPI & alerts*

**KPIs monitored:**

- SecondsUntilKeyMaterialExpiration

**Alerts observed:**

- SecondsUntilKeyMaterialExpiration if the customer imports his keys

#### *5.7.3.3 Backup and restore*

**Data backup and restore**

There is no data to Backup. Keys durability is assured by AWS (Roll-back on deletion can be set up : deletion delay can be set up between 7 and 30 days).

#### *5.7.3.4 AWS SLA High Availability and Disaster Recovery inter-region*

High Availability is supported by AWS for this service.

#### 5.7.3.5 Security

Security recommendations can be part of an optional security scope of work based on customer request. By default, the MRC does not cover security recommendations.

### 5.7.4 Charging model

| Work Unit |
| --- |
| Per KMS key |

### 5.7.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |
| Add/remove key | 1 token |
| Configure access policy | 2 tokens |
| Configure AWS native services to use key KMS | Estimation in tokens based on time spent |
| Other changes | Estimation in tokens based on time spent |

## 5.8 Elastic Load Balancing - Network Load Balancer (NLB)

### 5.8.1 Description

A Network Load Balancer functions at the fourth layer of the Open Systems Interconnection (OSI) model. It can handle millions of requests per second. After the load balancer receives a connection request, it selects a target from the target group for the default rule. It attempts to open a TCP connection to the selected target on the port specified in the listener configuration.

### 5.8.2 Build to run service included in the OTC

#### 5.8.2.1 Build service pre-requisite

- Refer to generic description.

#### 5.8.2.2 Build to run service

- Refer to generic description.

### 5.8.3 RUN services included in the MRC

#### 5.8.3.1 Run service pre-requisite

- A referential file exists in the Git including the reference configuration of the Elastic Load Balancing – Network Load Balancer (NLB).
- This file can be executed with a CI/CD and the execution has been tested successfully.

#### 5.8.3.2 Co-manage option

No, ORANGE BUSINESS manages the Network Load Balancer

#### 5.8.3.3 KPI & alerts

**Monitoring**

Yes:

**KPI monitored**

- ActiveFlowCount
- ActiveFlowCount_TCP
- ActiveFlowCount_TLS
- ActiveFlowCount_UDP
- ClientTLSNegotiationErrorCount
- ConsumedLCUs
- ConsumedLCUs_TCP
- ConsumedLCUs_TLS
- ConsumedLCUs_UDP
- HealthyHostCount
- NewFlowCount
- NewFlowCount_TCP
- NewFlowCount_TLS
- NewFlowCount_UDP
- PeakBytesPerSecond
- PeakPacketsPerSecond
- ProcessedBytes
- ProcessedBytes_TCP
- ProcessedBytes_TLS
- ProcessedBytes_UDP
- ProcessedPackets
- TargetTLSNegotiationErrorCount
- TCP_Client_Reset_Count
- TCP_ELB_Reset_Count
- TCP_Target_Reset_Count
- UnHealthyHostCount

**Alerts observed**

UnHealthyHostCount (correlated with an EC2 instance Down)

#### *5.8.3.4 Backup and restore*

**Data backup and restore**

Not applicable. Load balancer does not store data.

**Service restore**

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

#### *5.8.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

The high availability is ensured by AWS.
There is no native DR for this service.
Maintaining a cross region Disaster Recovery requires specific design and subject to a specific additional charging.

### 5.8.4 Charging model

| Work Unit |
|---|
| **Per Network Load Balancer instance** |

### 5.8.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort | Impact on MRC |
|---|---|---|
| **Add a pool member (if members are static)** | 1 token | |
| **Add a listener** | 2 tokens | |
| **Add a Target Group** | 4 tokens | |
| **Other changes** | Estimation in tokens based on time spent | |

## 5.9 Virtual Private Cloud (VPC) - Security Group

### 5.9.1 Description

A security group acts as a virtual firewall for your instance to control inbound and outbound traffic. When you launch an instance in a VPC, you can assign up to five security groups to the instance. Security groups act at the instance level, not the subnet level. Therefore, each instance in a subnet in your VPC can be assigned to a different set of security groups.

At the basic level, managing Network Security group consists in building, deploying, and maintaining the Infra as Code for it and managing the changes.

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

### 5.9.2 Reporting

Optionally, AWS Config can be used on demand with extra fees.

### 5.9.3 Backup

Can be exported from Infra as Code.

### 5.9.4 Charging model

| Work Unit | OTC & MRC |
|---|---|
| **Network and security management services** | Custom, depending on agreed Scope of Work |

### 5.9.5 Changes catalogue – in Tokens, per act

OCB team has to validate rules before being applied for security compliance.
Example: for Frontend we can open port 80 or 443 and this can consume 1 token, for Backend we open port 3306 (1 token).

| Changes examples | Effort |
|---|---|
| **Add / modify / delete Security group (up to 5 rules) excluding dependencies*** | 2 tokens |
| **Add / modify / delete Security rules (up to 5 rules) excluding dependencies*** | 1 token |
| **Other changes** | Estimation in tokens based on time spent |

*Dependencies include all triggered applications like, EC2, AWS Firewall, AWS DB services and other native services.

## 5.10 Amazon Elastic Compute Cloud (EC2) and OS

### 5.10.1 Description

The Managed Service for EC2 is called Managed OS. ORANGE BUSINESS manages both the OS and the EC2.

Amazon Elastic Compute Cloud (Amazon EC2) provides scalable computing capacity in the Amazon Web Services (AWS) Cloud. Using Amazon EC2 eliminates your need to invest in hardware up front, so you can develop and deploy applications faster. You can use Amazon EC2 to launch as many or as few virtual servers as you need, configure security and networking, and manage storage. Amazon EC2 enables you to scale up or down to handle changes in requirements or spikes in popularity, reducing your need to forecast traffic.

### 5.10.2 Build to run service included in the OTC

#### *5.10.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.10.2.2 Build to run service*

- Refer to generic description.

### 5.10.3 RUN services included in the MRC

#### *5.10.3.1 Run service pre-requisite*

- A referential file exists in the Git including the reference configuration of the EC2.
- This file can be executed with a CI/CD and the execution has been tested successfully.

#### *5.10.3.2 KPI & alerts*

**KPI monitored for Instances:**

- CPUUtilization
- DiskReadOps
- DiskWriteOps
- DiskReadBytes
- DiskWriteBytes
- MetadataNoToken

- NetworkIn
- NetworkOut
- NetworkPacketsIn
- NetworkPacketsOut

Other metrics can be collected from CloudWatch agent (Deployed in each EC2 instance) like MemoryUsage and DiskUsage.

**Alerts observed:**

Alert on CPU, MemoryUsage and DiskUsage.

Alerts also on Status Checks: SystemStatusChecks and InstanceStatusChecks (aggregated in StatusCheckMetric)

Optional: Depending on the criticity of the application we might activate Detailed monitoring: With Basic monitoring. Data is available automatically in 5-minute periods. Using Detailed monitoring, data is available in 1-minute periods. To get this level of data, you must specifically enable it for the instance.

Activating Detailed Monitoring will be charged by AWS.

### *5.10.3.3 OS patching*

**AWS Systems Manager Patch Manager**

For managed OS, ORANGE BUSINESS leverages AWS Systems Manager Patch Manager for the patching of the Operating System (OS).

Behavior: With AWS Systems Manager Patch Manager, patches are decided by Amazon and all patches are to be applied if mandatory for the EC2 for Windows and Linux.

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

### *5.10.3.4 Antivirus*

For managed OS, ORANGE BUSINESS leverages its central anti-virus system based on Sophos. This requires the installation of the anti-virus agent on the OS for each EC2 as well as the VPN connectivity to ORANGE BUSINESS Centralized Administration Zone. ORANGE BUSINESS systems allows for central reporting on Malware from its backend console system.

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

### *5.10.3.5 Backup and restore*

**Data backup and restore**

By default, ORANGE BUSINESS leverages AWS Backup on the EC2 for Managed OS. The configuration of AWS Backup pattern as well as retention period shall be agreed with the Customer prior to the RUN. The first backup is full. The following backups are incremental. You can the frequency of the backup. As example: 1 x backup per week, 1x incremental backup per day per EC2. The retention period depends on customer request. AWS charges will be calculated based on change rate.

Restore of EC2 are performed from the backup.

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

#### 5.10.3.6 AWS SLA High Availability and Disaster Recovery inter-region

Service is Highly Available within a single Availability Zone.

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

This service is covered by AWS Backup which enables the creation of backup copies across AWS Regions.

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

#### 5.10.3.7 Administration tasks tracing

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

The VPN connectivity to the ORANGE BUSINESS Administration Zone necessary for the management.

#### 5.10.3.8 Login on to the Virtual Machine

For Windows OS based EC2, access shall be granted by the Customer to ORANGE BUSINESS managed application operations staff through a domain account configured with proper privilege groups.

For Linux OS based EC2, an encrypted key is created and provided to ORANGE BUSINESS managed application operations staff to log onto the VM.

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

#### 5.10.3.9 Logs

Log management is not included in the managed OS / managed EC2 service.
Optionally it can be activated through AWS CloudWatch Logs through Change Request process.

#### 5.10.3.10 Security

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

#### 5.10.3.11 Limitations

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

### 5.10.4 Charging model

| Work Unit |
|---|
| Per EC2 Instance |

### 5.10.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| **Create an EC2 Instance** | 4 Tokens |
| **Create an EC2 instance integrated to AD** | 2 Tokens |
| **Modify/delete Security Groups** | 2 Tokens |

| | |
|---|---|
| **Extend and existing volume** | 3 Tokens |
| **Attach a new volume** | 2 Tokens |
| **Modify EC2** | Estimation in tokens based on time spent |
| **Delete EC2** | 2 Tokens |
| **Start/Stop/Restart EC2** | 1 Token |
| **Other changes** | Estimation in tokens based on time spent |

## 5.11 Web Application Firewall (WAF)

### 5.11.1 Description

AWS Web Application Firewall (WAF) provides centralized protection of your web applications from common exploits and vulnerabilities. Web applications are increasingly targeted by malicious attacks that exploit commonly known vulnerabilities. SQL injection and cross-site scripting are among the most common attacks.

### 5.11.2 Build to run service included in the OTC

#### *5.11.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.11.2.2 Build to run service*

- Refer to generic description.

### 5.11.3 RUN services included in the MRC

#### *5.11.3.1 Run service pre-requisite*

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

#### *5.11.3.2 Co-manage option*

No, ORANGE BUSINESS manages the WAF.
The customer provides the WAF rules to ORANGE BUSINESS Team who will review, configure and apply them.

#### *5.11.3.3 KPI & alerts*

**Monitoring**

Yes

**KPI monitored**

- AllowedRequests
- BlockedRequests
- CountedRequests
- CaptchaRequests
- RequestsWithValidCaptchaToken
- PassedRequests

**Alerts observed**

BlockedRequests

#### *5.11.3.4 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.11.3.5 Backup and restore*

**Data backup and restore**

No persistent data to backup up

**Service restore**

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

#### *5.11.3.6 AWS SLA High Availability and Disaster Recovery inter-region*

The service is Highly available by design by AWS.
WAF is a global AWS service. Disaster Recovery is native.

#### *5.11.3.7 Network and security managed services*

Additional Network and Security Managed services might be added optionally depending on Scope of Work.

### 5.11.4 Charging model

| Work Unit |
|---|
| **Access Control list (ACL)** |

### 5.11.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| **Add already existing rule** | 2 tokens |
| **modify/delete rule/rules (up to 5)** | 2 tokens |
| **Create a simple rule** | Estimation in tokens based on time spent |
| **Create a complex rule** | Estimation in tokens based on time spent |

| | |
|---|---|
| **Other changes** | Estimation in tokens based on time spent |

# 5.12 Amazon Elastic File System

## 5.12.1 Description

Amazon EFS provides simple, scalable, elastic file storage for use with compute instances on the AWS Cloud and on-premises servers.

## 5.12.2 Build to run service included in the OTC

### *5.12.2.1 Build service pre-requisite*

- Please refer to generic description

### *5.12.2.2 Build to run service*

- Please refer to generic description.

## 5.12.3 RUN services included in the MRC

### *5.12.3.1 Run service pre-requisite*

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

### *5.12.3.2 Co-manage option*

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

### *5.12.3.3 KPI & alerts*

**Monitoring**

Yes

**KPI monitored**

- PermittedThroughput
- ClientConnections
- StorageBytes
- TotalIOBytes

**Alerts observed**

StorageBytes

### *5.12.3.4 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report on the total storage size.

### *5.12.3.5 Backup and restore*

**Data backup and restore**

AWS Backup native service is used for Backup. The Backup is at a block level: If the customer wants to restore a specific file, he must implement his own file-level backup solution.

**Service restore**

Restore from Backup

#### *5.12.3.6 AWS SLA High Availability and Disaster Recovery inter-region*

Yes, by default by AWS. There is no native disaster recovery.

### 5.12.4 Charging model

| Work Unit |
| --- |
| per EFS filesystem |

### 5.12.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |
| Change Access Tier | Estimation in tokens based on time spent |
| Change permission policy (for encrypted file systems) | Estimation in tokens based on time spent |
| Other Changes | Estimation in tokens based on time spent |

## 5.13 AWS Elastic Beanstalk

### 5.13.1 Description

AWS Elastic Beanstalk is an easy-to-use AWS service for deploying and managing applications developed with Python, Ruby, Java, .NET, PHP, Node.js and Go and Docker on familiar servers such as Apache, Nginx, Passenger, and IIS. Elastic Beanstalk reduces management complexity by automatically handling the details of capacity provisioning, load balancing, scaling, and application health monitoring.

### 5.13.2 Build to run service included in the OTC

#### *5.13.2.1 Build service pre-requisite*

- Please refer to generic description.

#### *5.13.2.2 Build to run service*

- Please refer to generic description.

### 5.13.3 RUN services included in the MRC

#### *5.13.3.1 Run service pre-requisite*

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

#### *5.13.3.2 Co-manage option*

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

#### *5.13.3.3 KPI & alerts*

**Monitoring**

Yes

**KPI monitored**

- EnvironmentHealth
- InstancesSevere
- InstancesDegraded
- InstancesWarning
- InstancesInfo
- InstancesOk
- InstancesPending
- InstancesUnknown
- InstancesNoData
- ApplicationRequestsTotal
- ApplicationRequests5xx
- ApplicationRequests4xx
- ApplicationRequests3xx
- ApplicationRequests2xx
- ApplicationLatencyP10
- ApplicationLatencyP50
- ApplicationLatencyP75
- ApplicationLatencyP85
- ApplicationLatencyP90
- ApplicationLatencyP95
- ApplicationLatencyP99
- ApplicationLatencyP99.9
- InstanceHealth

### Available metrics—Linux

- CPUIrq
- CPUIdle
- CPUUser
- CPUSystem
- CPUSoftirq
- CPUIowait
- CPUNice
- LoadAverage1min
- RootFilesystemUtil

### Available metrics—Windows

- CPUIdle
- CPUUser
- CPUPriveleged

**Alerts observed**

Alert on InstanceHealth

Alert on CPUIdle and CPUUser for both Linux and Windows

#### *5.13.3.4 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.13.3.5 Backup and restore*

**Data backup and restore**

There is no data to backup.

**Service restore**

Recovery will be from Infra as Code.

#### *5.13.3.6 AWS SLA High Availability and Disaster Recovery inter-region*

Disaster Recovery is optional for this service. The service will be rebuilt in another region based on Infra as Code.

### 5.13.4 Charging model

| Work Unit |
|---|
| per Web Application |

### 5.13.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| **Add a custom domain on an AWS Elastic Beanstalk** | 3 Tokens |
| **Configure a connection string to access another resource** | 2 Tokens |
| **Deploy a new version of an existing webapp** | 2 Tokens |
| **Migrate an on-Premises webapp on AWS Elastic Beanstalk** | Estimation in tokens based on time spent |
| **Move Elastic Beanstalk in another region** | Estimation in tokens based on time spent |
| **Create and deploy a new webapp** | Estimation in tokens based on time spent |
| **Other Changes** | Estimation in tokens based on time spent |

## 5.14 Amazon GuardDuty

### 5.14.1 Description

Amazon GuardDuty is a threat detection service that continuously monitors your AWS accounts and workloads for malicious activity and delivers detailed security findings for visibility and remediation.

It is highly recommended to activate Amazon GuardDuty regardless of the AWS services used.

### 5.14.2 Build to run service included in the OTC

#### *5.14.2.1 Build service pre-requisite*

- Please refer to generic description

#### *5.14.2.2 Build to run service*

- Please refer to generic description.

### 5.14.3 RUN services included in the MRC

#### *5.14.3.1 Run service pre-requisite*

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

#### *5.14.3.2 Co-manage option*

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

#### *5.14.3.3 KPI & alerts*

Amazon GuardDuty analyzes and processes the following Data sources: VPC Flow Logs, AWS CloudTrail management event logs, CloudTrail S3 data event logs, and DNS logs.

An action is taken if a threat is detected (trigger a lambda function, event management, workflow logs).

By default, we will set CloudWatch events.

Optionally, a Lambda Function for automatic remediation can be requested by customer. The estimation will be based on time spent.

#### *5.14.3.4 Backup and restore*

**Data backup and restore**

There is no native backup for this service.

**Service restore**

Recovery will be from Infra as Code.

#### *5.14.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

The service is highly available by design by AWS.

There is no native Disaster Recovery for this service.

### 5.14.4 Charging model

| Work Unit |
|---|
| per security threat |

### 5.14.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| Add Trusted IP List or threat list | 2 Tokens |

| | |
|---|---|
| **Modify, Delete an existing rule** | Estimation in tokens based on time spent |
| **Other Changes** | Estimation in tokens based on time spent |

# 5.15 Amazon MQ

## 5.15.1 Description

Amazon MQ is a managed message broker service. A message broker allows software applications and components to communicate using various programming languages, operating systems, and formal messaging protocols. Currently, Amazon MQ supports Apache ActiveMQ and RabbitMQ engine types.

## 5.15.2 Build to run service included in the OTC

### *5.15.2.1 Build service pre-requisite*

- Refer to generic description.

### *5.15.2.2 Build to run service*

- Refer to generic description.

## 5.15.3 RUN services included in the MRC

### *5.15.3.1 Run service pre-requisite*

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

### *5.15.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

### *5.15.3.3 KPI & alerts*

### *ActiveMQ*

**Monitoring**

Yes

**KPI monitored**

**Amazon MQ for ActiveMQ metrics:**

- BurstBalance
- CpuCreditBalance
- CpuUtilization
- CurrentConnectionsCount
- EstablishedConnectionsCount
- HeapUsage
- InactiveDurableTopicSubscribersCount
- JobSchedulerStorePercentUsage
- JournalFilesForFastRecovery
- JournalFilesForFullRecovery
- NetworkIn
- NetworkOut
- OpenTransactionCount
- StorePercentUsage
- TempPercentUsage
- TotalConsumerCount
- TotalMessageCount
- TotalProducerCount
- VolumeReadOps
- VolumeWriteOps

**ActiveMQ destination (queue and topic) metrics:**

- ConsumerCount
- EnqueueCount
- EnqueueTime
- ExpiredCount

- DispatchCount
- DequeueCount
- InFlightCount
- ReceiveCount
- MemoryUsage
- ProducerCount
- QueueSize
- TotalEnqueueCount
- TotalDequeueCount

**Alerts observed**

Specific ActiveMQ alerts: CpuUtilization and HeapUsage

For Message Broker alerts: EnqueueCount and EnqueueTime

### *RabbitMQ*

**Monitoring**

Yes

**KPI monitored**

**RabbitMQ broker metrics:**

- ExchangeCount
- QueueCount
- ConnectionCount
- ChannelCount
- ConsumerCount
- MessageCount
- MessageReadyCount
- MessageUnacknowledgedCount
- PublishRate
- ConfirmRate
- AckRate

**RabbitMQ node metrics:**

- SystemCpuUtilization
- RabbitMQMemLimit
- RabbitMQMemUsed
- RabbitMQDiskFreeLimit
- RabbitMQDiskFree
- RabbitMQFdUsed

**RabbitMQ queue metrics:**

- ConsumerCount
- MessageReadyCount
- MessageUnacknowledgedCount
- MessageCount

**Alerts observed**

Message Broker alerts: AckRate

Node alerts: SystemCpuUtilization, RabbitMQMemUsed and RabbitMQDiskFree

#### *5.15.3.4 Backup and restore*

**Data backup and restore**

There is no data to backup.

**Service restore**

Recovery will be from Infra as Code. The messages in queue when the incident occurs won't be recovered.

#### *5.15.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

There is no native Disaster Recovery for this service.
Optionally, the DR can be customized by design. We recommend activating this option for Production Workloads that require High Availability and message durability.

### 5.15.4 Charging model

| Work Unit |
| --- |
| Per broker |

### 5.15.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |
| Upgrade broker | 6 Tokens |
| Reboot broker | 2 Tokens |
| Change maintenance window | 1 Token |
| Change Broker configuration* | Estimation in tokens based on time spent |
| Other changes | Estimation in tokens based on time spent |

*Each broker has his own specificities that can be configured.

## 5.16 Amazon Simple Notification Service (SNS)

### 5.16.1 Description

Amazon Simple Notification Service (Amazon SNS) is a managed service that provides message delivery from publishers to subscribers (also known as producers and consumers). Publishers communicate asynchronously with subscribers by sending messages to a topic, which is a logical access point and communication channel. Clients can subscribe to the SNS topic and receive published messages using a supported endpoint type, such as Amazon Kinesis Data Firehose, Amazon SQS, AWS Lambda, HTTP, email, mobile push notifications, and mobile text messages (SMS).

### 5.16.2 Build to run service included in the OTC

#### *5.16.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.16.2.2 Build to run service*

- Refer to generic description.

### 5.16.3 RUN services included in the MRC

#### *5.16.3.1 Run service pre-requisite*

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

#### *5.16.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.16.3.3 KPI & alerts*

**Monitoring**

optional

**KPI monitored**

- NumberOfMessagesPublished
- NumberOfNotificationsDelivered
- NumberOfNotificationsFailed
- NumberOfNotificationsFilteredOut
- NumberOfNotificationsFilteredOut-InvalidAttributes
- NumberOfNotificationsFilteredOut-NoMessageAttributes
- NumberOfNotificationsRedrivenToDlq
- NumberOfNotificationsFailedToRedriveToDlq
- PublishSize
- SMSMonthToDateSpentUSD
- SMSSuccessRate

**Alerts observed**

Alert on NumberOfNotificationsFailed.

Optionally, other alerts could be observed. The selection of these additional alerts depends on the Application's requirements.

#### *5.16.3.4 Backup and restore*

**Data backup and restore**

There is no data to backup.

**Service restore**

Recovery will be from Infra as Code.

#### *5.16.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

The service is Highly Available by default by AWS.

### 5.16.4 Charging model

| Work Unit |
| --- |
| Per SNS topic |

### 5.16.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |
| Adding a subscription | 1 token |
| Other changes | Estimation in tokens based on time spent |

## 5.17 Amazon Simple Queue Service (SQS)

### 5.17.1 Description

Amazon Simple Queue Service (Amazon SQS) offers a secure, durable, and available hosted queue that lets you integrate and decouple distributed software systems and components. Amazon SQS offers common constructs such as dead-letter queues and cost allocation tags.

### 5.17.2 Build to run service included in the OTC

#### *5.17.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.17.2.2 Build to run service*

- Refer to generic description.

### 5.17.3 RUN services included in the MRC

#### *5.17.3.1 Run service pre-requisite*

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

#### *5.17.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.17.3.3 KPI & alerts*

**Monitoring**

optional

**KPI monitored**

- ApproximateAgeOfOldestMessage
- ApproximateNumberOfMessagesDelayed
- ApproximateNumberOfMessagesNotVisible
- ApproximateNumberOfMessagesVisible
- NumberOfEmptyReceives
- NumberOfMessagesDeleted
- NumberOfMessagesReceived
- NumberOfMessagesSent
- SentMessageSize

**Alerts observed**

Alert on ApproximateAgeOfOldestMessage.

Optionally, other alerts could be observed. The selection of these additional alerts depends on the Application's requirements.

#### *5.17.3.4 Backup and restore*

**Data backup and restore**

There is no data to backup.

**Service restore**

Recovery will be from Infra as Code.

#### *5.17.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

The service is Highly Available by default by AWS. There is no native Disaster Recovery.

### 5.17.4 Charging model

| Work Unit |
| --- |
| Per SQS Queue |

### 5.17.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |
| Adding a subscription | 1 Token |

| | |
|---|---|
| Other changes | Estimation in tokens based on time spent |

# 5.18 Amazon CloudWatch – basic monitoring with class 2 transition

## 5.18.1 Description

Amazon CloudWatch is a monitoring and observability service. CloudWatch provides you with data and actionable insights to monitor your applications, respond to system-wide performance changes, and optimize resource utilization. CloudWatch collects monitoring and operational data in the form of logs, metrics, and events. You get a unified view of operational health and gain complete visibility of your AWS resources, applications, and services running on AWS and on-premises. You can use CloudWatch to detect anomalous behavior in your environments, set alarms, visualize logs and metrics side by side, take automated actions, troubleshoot issues, and discover insights to keep your applications running smoothly.

## 5.18.2 Build to run service included in the OTC

### *5.18.2.1 Build to run service pre-requisite*

The pre-requisite to Amazon CloudWatch basic monitoring with class 2 transition is that it has been configured by the Customer including

- Resources monitored
- CloudWatch Agent deployed on the resources when applicable
- Metrics and alerts forwarded to Amazon CloudWatch
- Performance dashboards using CloudWatch Dashboards

### *5.18.2.2 Build to run service*

For Amazon CloudWatch basic monitoring with class 2 transition, the build to run service included in the OTC consists in integrating the alerts into ORANGE BUSINESS supervision backend.

## 5.18.3 RUN services included in the MRC

### *5.18.3.1 Run service pre-requisite*

- The resource monitored is in the inventory Scope of Work of managed service: infrastructure resource, middleware resource, application resource, database resource, Kubernetes cluster resource, microservice resource, etc.…
- A referential file exists in the Git including the reference configuration of Amazon CloudWatch.
- This file can be executed with a CI/CD and the execution has been tested successfully.

### *5.18.3.2 KPI & alerts*

**Monitoring**

Yes

**Alerts observed**

- Alerts defined in Amazon CloudWatch for resources in the Scope of Work of managed services.

### *5.18.3.3 Monitoring service*

As part of the Application basic monitoring service, ORANGE BUSINESS operations will monitor the alerts, raise tickets and inform the Customer on incident. The basic service excludes remedial of incident.

### *5.18.3.4 Backup and restore*

**Backup and restore of Amazon CloudWatch:** N/A

**Service restore of Amazon CloudWatch:** The configuration of Amazon CloudWatch can be recovered from Infrastructure-as-code if its configuration has been done through infrastructure as code.

**Backup and restore of resources monitored by Amazon CloudWatch:** N/A

**Restore from IaC for resources monitored by Amazon CloudWatch:** N/A

#### *5.18.3.5 Limitations & pre-requisite*

The Amazon CloudWatch basic monitoring service is monitoring only.

### 5.18.4 Charging model

| Work Unit |
| --- |
| Per managed resource |

### 5.18.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |
| Other changes | Estimation in tokens based on time spent |

## 5.19 AWS Backup – basic backup with class 2 transition

### 5.19.1 Description

AWS Backup enables you to centralize and automate data protection across AWS services and hybrid workloads. AWS Backup offers a cost-effective, fully managed, policy-based service that further simplifies data protection at scale. AWS Backup also helps you support your regulatory compliance or business policies for data protection. Together with AWS Organizations, AWS Backup enables you to centrally deploy data protection policies to configure, manage, and govern your backup activity across your organization's AWS accounts and resources, including Amazon Elastic Compute Cloud (Amazon EC2) instances, Amazon Elastic Block Store (Amazon EBS) volumes, Amazon Simple Storage Service (Amazon S3) buckets, Amazon Relational Database Service (Amazon RDS) databases (including Amazon Aurora clusters), Amazon DynamoDB tables, Amazon Neptune databases, Amazon DocumentDB (with MongoDB compatibility) databases, Amazon Elastic File System (Amazon EFS) file systems, Amazon FSx for Lustre file systems, Amazon FSx for Windows File Server file systems, and AWS Storage Gateway volumes, and VMware workloads on premises and in VMware Cloud™ on AWS.

### 5.19.2 Build to run service included in the OTC

#### *5.19.2.1 Build to run service pre-requisite*

The pre-requisite to AWS backup – backup with class 2 transition is that it has been configured by the Customer including

- Resources backed up when applicable
- Backup configured and VSS deployed when applicable
- Metrics and alerts on backup status forwarded to Amazon CloudWatch

#### *5.19.2.2 Build to run service*

For AWS Backup with class 2 transition, the build to run service included in the OTC consists in integrating the alerts on backup status into ORANGE BUSINESS supervision backend.

### 5.19.3 RUN services included in the MRC

#### *5.19.3.1 Run service pre-requisite*

- The resource backed-up is in the inventory Scope of Work of managed services supported by AWS Backup
- A referential file exists in the Git including the reference configuration of AWS Backup.
- This file can be executed with a CI/CD and the execution has been tested successfully.

#### *5.19.3.2 KPI & alerts*

**Monitoring**

Yes

**Job status monitored KPIs**

- CREATED
- PENDING
- RUNNING
- ABORTED
- COMPLETED
- FAILED
- EXPIRED

**Alerts observed**

- Alerts on ABORTED and FAILED for each resource supported by AWS Backup

#### *5.19.3.3 Backup service*

As part of the AWS Backup service, ORANGE BUSINESS operations will monitor the alerts related to backup status, raise tickets and inform the Customer on incident. The basic service excludes data recovery. Data Recovery is requested through a change request.

#### *5.19.3.4 Backup and restore*

**Backup and restore of AWS Backup:** N/A

**Service restore of AWS Backup:** The configuration of AWS Backup can be recovered from Infrastructure-as-code if its configuration has been done through infrastructure as code.

#### *5.19.3.5 Limitations & pre-requisite*

AWS Backup native service is used for Backup. The Backup provided is a block-level one: If the customer wants to restore a specific file, he needs to implement his own file-level backup solution.

### 5.19.4 Charging model

| Work Unit |
| --- |
| Per managed resource |

### 5.19.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |

| | |
|---|---|
| **Change the configuration of backup plan** | 3 Tokens |
| Other changes | Estimation in tokens based on time spent |

# 5.20 Amazon Elastic Container Service (ECS)

## 5.20.1 Description

Amazon Elastic Container Service (ECS) is a highly scalable, high performance container management service that supports Docker containers and allows you to easily run applications on a managed cluster of Amazon Elastic Compute Cloud (Amazon EC2) instances. Amazon ECS eliminates the need for you to install, operate, and scale your own cluster management infrastructure. With simple API calls, you can launch and stop container-enabled applications, query the complete state of your cluster, and access many familiar features like security groups, Elastic Load Balancing, Amazon Elastic Block Store (EBS) volumes, and Identity Access Management (IAM). roles. You can use Amazon ECS to schedule container placement across your cluster based on your resource needs and availability requirements. You can also integrate your own scheduler or third-party schedulers to meet business or application specific requirements.

## 5.20.2 Build to run service included in the OTC

### *5.20.2.1 Build service pre-requisite*

- Refer to generic description.

### *5.20.2.2 Build to run service*

- Refer to generic description.

## 5.20.3 RUN services included in the MRC

### *5.20.3.1 Run service pre-requisite*

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

### *5.20.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

### *5.20.3.3 KPI & alerts*

**Monitoring**

optional

**KPI monitored**

- CPUReservation
- CPUUtilization
- MemoryReservation
- MemoryUtilization
- GPUReservation

When Amazon ECS runs containers on top of EC2 instances, EC2 metrics will be collected as well (Please refer to EC2 section).

When “EC2 launch type” is used with Linux container instances, the Amazon ECS container agent relies on Docker stats metrics to gather CPU and memory data for each container running on the instance. For burstable performance instances (T3, T3a, and T2 instances), the CPU utilization metric may reflect different data compared to instance-level CPU metrics.

**Alerts observed**

- CPUReservation
- CPUUtilization
- MemoryReservation
- MemoryUtilization
- GPUReservation (optional, only if the application requires GPU)

Alerts for EC2 when Amazon ECS runs containers on top of EC2 instances (please refer to EC2 section)

#### *5.20.3.4 Backup and restore*

**Data backup and restore**

There is no data to backup.

**Service restore**

Recovery will be from Infra as Code.

#### *5.20.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

When Amazon ECS runs containers on top of EC2, High Availability depends on the service configuration and is optional.

When Amazon ECS is used with Amazon Fargate, the service is natively Highly Available.
Disaster Recovery requires specific configuration and is optional.

### 5.20.4 Charging model

| Work Unit |
| --- |
| Per Docker Image |

### 5.20.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| Modify a container Version | 4 Tokens |
| Adjust the CPU and Memory | 2 Tokens |
| Create ECS cluster | 6 Tokens |
| Task deployment | 4 Tokens |
| Service deployment | 4 Tokens |
| Other changes | Estimation in tokens based on time spent |

## 5.21 Elastic Container Registry (ECR)

### 5.21.1 Description

Amazon Elastic Container Registry (Amazon ECR) is an AWS managed container image registry service that is secure, scalable, and reliable.

### 5.21.2 Build to run service included in the OTC

#### *5.21.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.21.2.2 Build to run service*

- Refer to generic description.

### 5.21.3 RUN services included in the MRC

#### *5.21.3.1 Run service pre-requisite*

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

#### *5.21.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.21.3.3 KPI & alerts*

**Monitoring**

optional

**KPI monitored**

- CallCount

**Alerts observed**

No alerts observed

#### *5.21.3.4 Backup and restore*

**Data backup and restore**

There is no native backup for this service. In case of an issue with ECR, all Docker images will be lost.

**Service restore**

Recovery will be from Infra as Code.

#### *5.21.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

The service is highly available by design by AWS.
Disaster Recovery is optional. Set up multi-region backup (synchronize with another region) can be requested by the customer. This will have an impact on storage cost.

### 5.21.4 Charging model

| Work Unit |
| --- |
| Per Docker Image |

### 5.21.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |
| Set up access to image | 2 Tokens |
| Other changes | Estimation in tokens based on time spent |

## 5.22 AWS Directory Service

### 5.22.1 Description

AWS Directory Service provides multiple ways to use Microsoft Active Directory (AD) with other AWS services. Directories store information about users, groups, and devices, and administrators use them to manage access to information and resources. AWS Directory Service provides multiple directory choices for customers who want to use existing Microsoft AD or Lightweight Directory Access Protocol (LDAP)–aware applications in the cloud. It also offers those same choices to developers who need a directory to manage users, groups, devices, and access.

### 5.22.2 Build to run service included in the OTC

#### *5.22.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.22.2.2 Build to run service*

- Refer to generic description.

### 5.22.3 RUN services included in the MRC

#### *5.22.3.1 Run service pre-requisite*

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

#### *5.22.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.22.3.3 KPI & alerts*

**Monitoring**

Yes

**KPI monitored**

- Active
- Creating
- Deleted
- Deleting
- Failed
- Impaired
- Inoperable
- Requested
- RestoreFailed
- Restoring
- Processor
- Memory
- Logical Disk
- Network Interface

- LDAP searches
- Binds
- DNS queries
- Directory reads
- Directory writes

**Alerts**

Alert on Failed, Impaired and Inoperable.

We will also trigger alerts on Processor and Memory.

Optionally, other alerts could be requested by the customer based on quote.

#### *5.22.3.4 Backup and restore*

**Data backup and restore**

Optionally, Plugging AD Backup tools can be requested by the customer.

**Service restore**

Recovery will be from Infra as Code.

#### *5.22.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

The service is highly available by design by AWS.
There is no native Disaster Recovery for this service.
Multi-Region replication is only supported for the **Enterprise Edition** of AWS Managed Microsoft AD. You can use automated multi-Region replication in all Regions where AWS Managed Microsoft AD is available.

### 5.22.4 Charging model

| Work Unit |
| --- |
| Per AD |

### 5.22.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |
| Modify the maintenance window | 1 Token |
| Set trusted relationship | Estimation in tokens based on time spent |
| Share Directory Service with another AWS account | 6 Tokens |
| Other changes | Estimation in tokens based on time spent |

## 5.23 Amazon Cognito

### 5.23.1 Description

Amazon Cognito handles user authentication and authorization for your web and mobile apps. With user pools, you can easily and securely add sign-up and sign-in functionality to your apps. With identity pools (federated identities), your apps can get temporary credentials that grant users access to specific AWS resources, whether the users are anonymous or are signed in.

### 5.23.2 Build to run service included in the OTC

#### *5.23.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.23.2.2 Build to run service*

- Refer to generic description.

### 5.23.3 RUN services included in the MRC

#### *5.23.3.1 Run service pre-requisite*

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

#### *5.23.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.23.3.3 KPI & alerts*

**Monitoring**

yes

**KPI monitored**

- SignUpSuccesses
- SignUpThrottles
- SignInSuccesses
- SignInThrottles
- TokenRefreshSuccesses
- TokenRefreshThrottles
- FederationSuccesses
- FederationThrottles
- CallCount
- ThrottleCount

**Alerts observed**

Alert on SignInThrottles

Optionally, other alerts can be observed.

#### *5.23.3.4 Backup and restore*

**Data backup and restore**

No native Backup option is provided. Optionally, the customer can request automatic accounts backup.

**Service restore**

Recovery will be from Infra as Code for user pools and identity pools.
Optionally, a customized backup can be provided on quote.
If the customer chooses to activate this option, the infrastructure cost will be impacted.

#### *5.23.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

The service is Highly available by design by AWS. There is no native Disaster Recovery.
Optionally, implement the customized backup solution cross-region.

### 5.23.4 Charging model

| Work Unit |
| --- |
| Per Pool |

### 5.23.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |

| | |
|---|---|
| Add IAM role for Cognito | 2 Tokens |
| Create a User pool | 6 Tokens |
| Create Identity pool | 6 Tokens |
| Other changes | Estimation in tokens based on time spent |

## 5.24 Amazon DynamoDB

### 5.24.1 Description

Amazon DynamoDB is a fully managed, serverless, key-value NoSQL database designed to run high-performance applications at any scale. DynamoDB offers built-in security, continuous backups, automated multi-Region replication, in-memory caching, and data export tools.

### 5.24.2 Build to run service included in the OTC

#### *5.24.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.24.2.2 Build to run service*

- Refer to generic description.

### 5.24.3 RUN services included in the MRC

#### *5.24.3.1 Run service pre-requisite*

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

#### *5.24.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.24.3.3 KPI & alerts*

**Monitoring**

yes

**Available metrics**

- AccountMaxReads
- AccountMaxTableLevelReads
- AccountMaxTableLevelWrites
- AccountMaxWrites
- AccountProvisionedReadCapacityUtilization
- AccountProvisionedWriteCapacityUtilization
- AgeOfOldestUnreplicatedRecord
- ConditionalCheckFailedRequests
- ConsumedChangeDataCaptureUnits
- ConsumedReadCapacityUnits
- ConsumedWriteCapacityUnits
- FailedToReplicateRecordCount
- MaxProvisionedTableReadCapacityUtilization
- MaxProvisionedTableWriteCapacityUtilization
- OnlineIndexConsumedWriteCapacity
- OnlineIndexPercentageProgress
- OnlineIndexThrottleEvents
- PendingReplicationCount
- ProvisionedReadCapacityUnits
- ProvisionedWriteCapacityUnits
- ReadThrottleEvents
- ReplicationLatency
- ReturnedBytes
- ReturnedItemCount
- ReturnedRecordsCount
- SuccessfulRequestLatency

- SystemErrors
- TimeToLiveDeletedItemCount
- ThrottledPutRecordCount
- ThrottledRequests
- TransactionConflict
- UserErrors
- WriteThrottleEvents

**Alerts observed**

Alert on SystemErrors reveals an issue with a Dynamo DB table.

Alert on UserErrors reveals an issue on the application side.

In case of multi-region replication, alert on ReplicationLatency.

Alert on AccountProvisionedReadCapacityUtilization and AccountProvisionedWriteCapacityUtilization explains an application latency.

Optionally, other alerts can be observed.

### *5.24.3.4 Backup and restore*

**Data backup and restore**

DynamoDB offers two methods to back up your table data. Continuous backups with point-in-time recovery (PITR) provide an ongoing backup of your table for the preceding 35 days. You can restore your table to the state of any specified second in the preceding five weeks. On-demand backups create snapshots of your table to archive for extended periods to help you meet corporate and governmental regulatory requirements.

- **On demand backup:** On-demand backup allows you to create full backups of your Amazon DynamoDB table at specified points in time. Recovery Point Objective will depend on the backup frequency chosen by the customer. This option is suitable for long-term retention and archival. It can help you to comply with regulatory requirements.
  There are two options available for creating and managing DynamoDB on-demand backups:
  - AWS Backup service
  - DynamoDB

  With AWS Backup, you can configure backup policies and monitor activity for your AWS resources and on-premises workloads in one place. Using DynamoDB with AWS Backup, you can copy your on-demand backups across AWS accounts and Regions, add cost allocation tags to on-demand backups, and transition on-demand backups to cold storage for lower costs. DynamoDB charges for on-demand backups based on the storage size of the table (table data and local secondary indexes). The size of each backup is determined at the time of each backup

request. The total backup storage size billed each month is the sum of all backups of DynamoDB tables. DynamoDB monitors the size of on-demand backups continuously throughout the month to determine your backup charges.

- **Continuous Backup:** With continuous backups, you can restore your AWS Backup-supported resource by rewinding it back to a specific time that you choose, within 1 second of precision (going back a maximum of 35 days). This built-in feature protects against accidental writes or deletes. Continuous backup works by first creating a full backup of your resource, and then constantly backing up your resource's transaction logs. **PITR** restore works by accessing your full backup and replaying the transaction log to the time that you tell AWS Backup to recover.
  The **Recovery Point Objective** (RPO) is close to zero.
  DynamoDB charges for PITR based on the size of each DynamoDB table (table data and local secondary indexes) on which it is enabled. DynamoDB monitors the size of your PITR-enabled tables continuously throughout the month to determine your backup charges and continues to bill you until you disable PITR on each table.

We could set both options to have a longer retention period and to minimize the RPO.
By default, we will set "on demand backup".
Restore will be done from backup depending on the option chosen by the customer: full backup or point in time recovery.

**Service restore**

Recovery will be from Infra as Code.

#### *5.24.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

DynamoDB automatically spreads the data and traffic for your tables over a sufficient number of servers to handle your throughput and storage requirements, while maintaining consistent and fast performance.

All your data is stored on solid-state disks (SSDs) and is automatically replicated across multiple Availability Zones in an AWS Region, providing built-in high availability and data durability.

You can use global tables to keep DynamoDB tables in sync across AWS Regions.

In same region, High Availability is a built-in feature.
in other regions, we will optionally use global tables to replicate tables across regions.

### 5.24.4 Charging model

| Work Unit |
| --- |
| per Dynamo DB table |

### 5.24.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |
| Edit table capacity | 2 tokens |
| Update table class | 2 tokens |
| Create snapshot | 2 tokens |

| | |
|---|---|
| Delete table | Estimation in tokens based on table size |
| Create item | Estimation in tokens based on number of rows |
| Create index | Estimation in tokens based on table size |
| Create replica | Estimation in tokens based on table size |
| Other changes | Estimation in tokens based on time spent |

# 5.25 ElastiCache for Redis

## 5.25.1 Description

Amazon ElastiCache is a fully managed, in-memory caching service supporting flexible, real-time use cases. You can use ElastiCache for caching, which accelerates application and database performance, or as a primary data store for use cases that don't require durability like session stores, gaming leaderboards, streaming, and analytics.
Built on open-source Redis and compatible with the Redis APIs, ElastiCache for Redis works with your Redis clients and uses the open Redis data format to store your data.

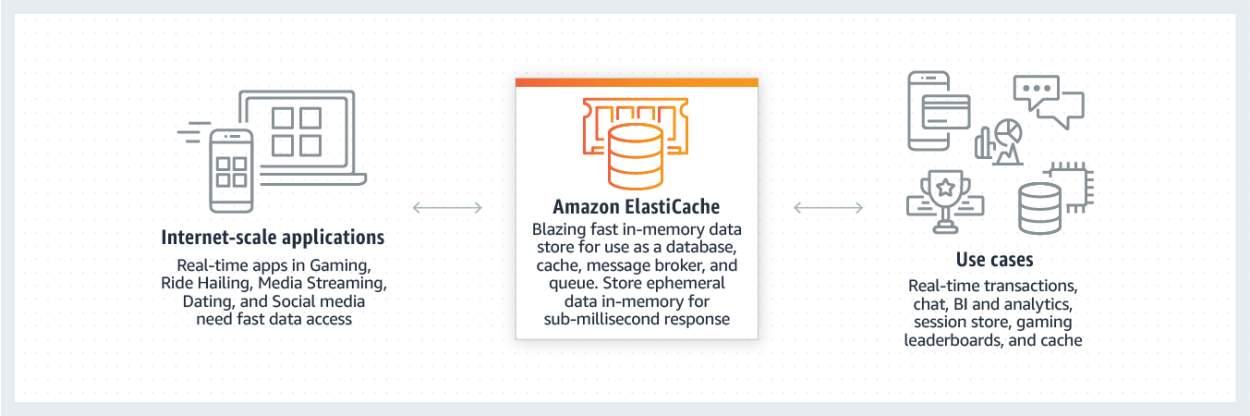


## 5.25.2 Build to run service included in the OTC

### *5.25.2.1 Build service pre-requisite*

- Refer to generic description.

### *5.25.2.2 Build to run service*

- Refer to generic description.

## 5.25.3 RUN services included in the MRC

### *5.25.3.1 Run service pre-requisite*

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

### *5.25.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.25.3.3 KPI & alerts*

**Monitoring**

yes

**Available metrics**

- CPUUtilization
- CPUCreditBalance
- CPUCreditUsage
- FreeableMemory
- NetworkBytesIn
- NetworkBytesOut
- NetworkPacketsIn
- NetworkPacketsOut
- NetworkBandwidthInAllowanceExceeded
- NetworkConntrackAllowanceExceeded
- NetworkLinkLocalAllowanceExceeded
- NetworkBandwidthOutAllowanceExceeded
- Network Packets Per Second Allowance Exceeded
- SwapUsage
- ActiveDefragHits
- AuthenticationFailures
- BytesUsedForCache
- BytesReadFromDisk
- BytesWrittenToDisk
- CacheHits
- CacheMisses
- CommandAuthorizationFailures
- CacheHitRate
- CurrConnections
- CurrItems
- CurrVolatileItems
- DatabaseMemoryUsagePercentage
- DatabaseMemoryUsageCountedForEvictPercentage
- DB0AverageTTL
- EngineCPUUtilization
- Evictions
- GlobalDatastoreReplicationLag
- IsPrimary
- KeyAuthorizationFailures
- KeysTracked
- MemoryFragmentationRatio
- NewConnections
- NumItemsReadFromDisk
- NumItemsWrittenToDisk
- PrimaryLinkHealthStatus
- Reclaimed
- ReplicationBytes
- ReplicationLag
- SaveInProgress

**Alerts observed**

- CPUUtilization (Host Level)

- EngineCPUUtilization (Node Level) {analyze the load of the Redis process}
- Evictions (Node Level) {if the number is very high we need to increase maxmemory limit}
- CurrConnections (Node Level)
- ReplicationLag (Node Level)
- DatabaseMemoryUsagePercentage (Node Level)

#### *5.25.3.4 Backup and restore*

**Data backup and restore**

ElasticCache for Redis offers two methods to back up your table data:

- **Automatic Backup:** For any Redis cluster, you can enable automatic backups. When automatic backups are enabled, ElastiCache creates a backup of the cluster on a daily basis. Automatic backups can help guard against data loss. In the event of a failure, you can create a new cluster, restoring your data from the most recent backup. The result is a warm-started cluster, preloaded with your data and ready for use. The minimum length for the backup window is 60 minutes. The maximum backup retention limit is 35 days.

- **Manual backups:** In addition to automatic backups, you can create a manual backup at any time. Unlike automatic backups, which are automatically deleted after a specified retention period, manual backups do not have a retention period after which they are automatically deleted. You must manually delete any manual backup. Even if you delete a cluster or node, any manual backups from that cluster or node are retained. Manual backups are useful for testing and archiving.

**Service restore**

Recovery will be from Infra as Code.

#### *5.25.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

Beginning with Redis version 3.2, you have the ability to create one of two distinct types of Redis clusters (API/CLI: replication groups). A Redis (cluster mode disabled) cluster always has a single shard (API/CLI: node group) with up to 5 read replica nodes. A Redis (cluster mode enabled) cluster has up to 500 shards with 1 to 5 read replica nodes in each.
With cluster mode enabled, your Redis Cluster gains enhanced scalability and high availability.
In addition, Amazon ElastiCache offers multiple Availability Zone (Multi-AZ) support with auto failover that enables you to set up a cluster with one or more replicas across zones. In the event of a failure on the primary node, Amazon ElastiCache for Redis automatically fails over to a replica to ensure high availability.
You can enable multi-AZ only on Redis (cluster mode disabled) clusters that have at least one available read replica. Clusters without read replicas do not provide high availability or fault tolerance.
Creating a Replication Group Using an Available Redis (Cluster Mode Disabled) is optional.

### 5.25.4 Charging model

| Work Unit |
|---|
| per node |

### 5.25.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| setup specific alert | 1 token |
| purge cache | 1 token |
| configuration modification | 1 token |
| Other changes | Estimation in tokens based on time spent |

# 5.26 Amazon MemoryDB for Redis

## 5.26.1 Description

MemoryDB for Redis is a durable, in-memory database service that delivers ultra-fast performance. It is purpose-built for modern applications with microservices architectures. MemoryDB is compatible with Redis, a popular open-source data store, enabling you to quickly build applications using the same flexible and friendly Redis data structures, APIs, and commands that they already use today.

## 5.26.2 Build to run service included in the OTC

### *5.26.2.1 Build service pre-requisite*

- Refer to generic description.

### *5.26.2.2 Build to run service*

- Refer to generic description.

## 5.26.3 RUN services included in the MRC

### *5.26.3.1 Run service pre-requisite*

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

### *5.26.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

### *5.26.3.3 KPI & alerts*

**Monitoring**

yes

**Available metrics**

- CPUUtilization
- FreeableMemory
- NetworkBytesIn
- NetworkBytesOut
- NetworkPacketsIn
- NetworkPacketsOut
- SwapUsage
- ActiveDefragHits
- AuthenticationFailures
- BytesUsedForMemoryDB
- CommandAuthorizationFailures
- CurrConnections
- CurrItems
- DatabaseMemoryUsagePercentage
- DB0AverageTTL
- EngineCPUUtilization
- Evictions
- IsPrimary
- KeyAuthorizationFailures
- KeyspaceHits

- KeyspaceMisses
- KeysTracked
- MaxReplicationThroughput
- MemoryFragmentationRatio
- NewConnections
- PrimaryLinkHealthStatus
- Reclaimed
- ReplicationBytes
- ReplicationDelayedWriteCommands
- ReplicationLag

**Alerts observed**

- CPUUtilization (Host Level)
- EngineCPUUtilization (Node Level) {analyze the load of the Redis process}
- Evictions (Node Level) {if the number is very high, we need to increase maxmemory limit}
- CurrConnections (Node Level)
- ReplicationLag (Node Level)
- DatabaseMemoryUsagePercentage (Node Level)

#### *5.26.3.4 Backup and restore*

**Data backup and restore**

MemoryDB for Redis clusters automatically back up data to a multi-AZ transactional log, but you can choose to create point-in-time snapshots of a cluster either periodically or on-demand. These snapshots can be used to recreate a cluster at a previous point or to seed a brand-new cluster. The snapshot consists of the cluster's metadata, along with all of the data in the cluster. All snapshots are written to Amazon Simple Storage Service (Amazon S3), which provides durable storage. At any time, you can restore your data by creating a new MemoryDB cluster and populating it with data from a snapshot.

**Service restore**

Recovery will be from Infra as Code.

### 5.26.4 Charging model

| Work Unit |
|---|
| per node |

### 5.26.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| Modify Associated subnets | 1 token |
| Modify parameter groups | 1 token |
| Modify Node type | 1 token |
| Modify Security Groups | 1 token |
| Modify ACL (Access Control List) | 1 token |
| Modify Snapshot | 1 token |
| Modify Maintenance window | 1 token |
| Take snapshot | Estimation in tokens based on table size |

| | |
|---|---|
| Other changes | Estimation in tokens based on time spent |

# 5.27 Amazon Neptune

## 5.27.1 Description

Amazon Neptune is a fast, reliable, fully managed graph database service that makes it easy to build and run applications that work with highly connected datasets. Neptune supports the popular graph query languages Apache TinkerPop Gremlin, the W3C's SPARQL, and Neo4j's openCypher, enabling you to build queries that efficiently navigate highly connected datasets. Neptune powers graph use cases such as recommendation engines, fraud detection, knowledge graphs, drug discovery, and network security.

## 5.27.2 Build to run service included in the OTC

### *5.27.2.1 Build service pre-requisite*

- Refer to generic description.

### *5.27.2.2 Build to run service*

- Refer to generic description.

## 5.27.3 RUN services included in the MRC

### *5.27.3.1 Run service pre-requisite*

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

### *5.27.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

### *5.27.3.3 KPI & alerts*

**Monitoring**

yes

**Available metrics**

- BufferCacheHitRatio: The percentage of requests that are served by the buffer cache. Cache misses add significant latency to query execution. If the cache hit ratio is below 99.9% and latency is an issue for your application, consider upgrading the instance type to cache more data in memory.
- CPU utilization: Percentage of computer processing capacity used. High values for CPU consumption might be appropriate, depending on your query-performance goals.
- Freeable memory: How much RAM is available on the DB instance, in megabytes. Neptune has its own memory manager, so this metric may be lower than you expect. A good sign that you should consider upgrading your instance class to one with more RAM is if queries often throw out-of-memory exceptions.
- BackupRetentionPeriodStorageUsed
- ClusterReplicaLag
- ClusterReplicaLagMaximum
- ClusterReplicaLagMinimum
- EngineUptime
- GremlinRequestsPerSec

- GremlinWebSocketOpenConnections
- LoaderRequestsPerSec
- MainRequestQueuePendingRequests
- NetworkReceiveThroughput
- NetworkThroughput
- NetworkTransmitThroughput
- NumTxCommitted
- NumTxOpened
- NumTxRolledBack
- SnapshotStorageUsed
- SparqlRequestsPerSec
- StatsNumStatementsScanned
- TotalBackupStorageBilled
- TotalRequestsPerSec
- TotalClientErrorsPerSec
- TotalServerErrorsPerSec
- VolumeBytesUsed
- VolumeReadIOPs
- VolumeWriteIOPs

**Alerts observed**

- ClusterReplicaLag
- CPUUtilization
- FreeableMemory
- TotalClientErrorsPerSec
- TotalServerErrorsPerSec
- BufferCacheHitRatio

#### *5.27.3.4 Backup and restore*

**Data backup and restore**

Neptune backs up your cluster volume automatically and retains restore data for the length of the backup retention period. Neptune backups are continuous and incremental so you can quickly restore to any point within the backup retention period. No performance impact or interruption of database service occurs as backup data is being written. You can specify a backup retention period, from 1 to 35 days, when you create or modify a DB cluster.

If you want to retain a backup beyond the backup retention period, you can also take a snapshot of the data in your cluster volume. Storing snapshots incurs the standard storage charges for Neptune.

**Service restore**

Recovery will be from Infra as Code.

#### *5.27.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

A Neptune DB cluster is fault tolerant by design. The cluster volume spans multiple Availability Zones in a single AWS Region, and each Availability Zone contains a copy of the cluster volume data. This functionality means that your DB cluster can tolerate a failure of an Availability Zone without any loss of data and only a brief interruption of service.

### 5.27.4 Charging model

| Work Unit |
| --- |
| per Database instance |

### 5.27.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| **Create clone** | 1 token |
| **Restore to point in time, depends on DB size** | Estimation in tokens based on Database size |
| **Upgrade** | Estimation in tokens based on time spent |
| **Operational (start, stop, reboot, failover)** | 1 token |
| **Modify and change configuration** | 1 token |
| **Create reader replica instance** | 1 token |
| Other changes | Estimation in tokens based on time spent |

## 5.28 Amazon Keyspaces (for Apache Cassandra)

### 5.28.1 Description

Amazon Keyspaces (for Apache Cassandra) is a scalable, highly available, and managed Apache Cassandra–compatible database service. With Amazon Keyspaces, you don't have to provision, patch, or manage servers, and you don't have to install, maintain, or operate software.

Amazon Keyspaces is serverless, so you pay for only the resources that you use, and the service automatically scales tables up and down in response to application traffic.

### 5.28.2 Build to run service included in the OTC

#### *5.28.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.28.2.2 Build to run service*

- Refer to generic description.

### 5.28.3 RUN services included in the MRC

#### *5.28.3.1 Run service pre-requisite*

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

#### *5.28.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.28.3.3 KPI & alerts*

**Monitoring**

yes

**Available metrics**

- AccountMaxTableLevelReads (The maximum number of read capacity units that can be used by a table of the account.)

- AccountMaxTableLevelWrites (The maximum number of write capacity units that can be used by a table of the account.)
- AccountProvisionedReadCapacityUtilization (The percentage of provisioned read capacity units utilized by an account.)
- AccountProvisionedWriteCapacityUtilization (The percentage of provisioned write capacity units utilized by an account.)
- ConditionalCheckFailedRequests
- ConsumedReadCapacityUnits (The number of read capacity units consumed over the specified time period.)
- ConsumedWriteCapacityUnits (The number of write capacity units consumed over the specified time period.)
- MaxProvisionedTableReadCapacityUtilization (The maximum percentage of provisioned read capacity units utilized by the highest provisioned read table of the account.)
- MaxProvisionedTableWriteCapacityUtilization
- PerConnectionRequestRateExceeded (Requests to Amazon Keyspaces that exceed the per-connection request rate quota)
- ProvisionedReadCapacityUnits
- ProvisionedWriteCapacityUnits
- ReadThrottleEvents (Requests to Amazon Keyspaces that exceed the provisioned read capacity for a table.)
- ReturnedItemCount (The number of rows returned by multi-row SELECT queries during the specified time period.)
- StoragePartitionThroughputCapacityExceeded (Requests to an Amazon Keyspaces storage partition that exceed the throughput capacity of the partition.)
- SuccessfulRequestLatency (The successful requests to Amazon Keyspaces during the specified time period.)
- SystemErrors (The requests to Amazon Keyspaces that generate a ServerError during the specified time period.)
- TTLDeletes (The units consumed to delete or update data in a row by using Time to Live (TTL).)
- UserErrors (Requests to Amazon Keyspaces that generate an InvalidRequest error during the specified time period.)
- WriteThrottleEvents (Requests to Amazon Keyspaces that exceed the provisioned write capacity for a table.)

**Alerts observed**

- SystemErrors (usually indicates an internal service error.)
- UserErrors (usually indicates a client-side error, such as an attempt to update a nonexistent table)
- AccountProvisionedReadCapacityUtilization
- AccountProvisionedWriteCapacityUtilization

### *5.28.3.4 Backup and restore*

**Data backup and restore**

Point-in-time recovery (PITR) helps protect your Amazon Keyspaces tables from accidental write or delete operations by providing you continuous backups of your table data.

For example, suppose that a test script writes accidentally to a production Amazon Keyspaces table. With point-in-time recovery, you can restore that table's data to any second in time since PITR was enabled within the last 35 days. If you delete a table with point-in-time recovery enabled, you can query for the deleted table's data for 35 days (at no additional cost), and restore it to the state it was in just before the point of deletion.

Point-in-time operations have no performance or availability impact on the base table and restoring a table doesn't consume additional throughput.

Amazon Keyspaces PITR uses two timestamps to maintain the time frame for which restorable backups are available for a table.

- Earliest restorable time – Marks the time of the earliest restorable backup. The earliest restorable backup goes back up to 35 days or when PITR was enabled, whichever is more recent. The maximum backup window of 35 days can't be modified.

- Current time – The timestamp for the latest restorable backup is the current time. If no timestamp is provided during a restore, current time is used.

When PITR is enabled, you can restore to any point in time between EarliestRestorableDateTime and CurrentTime. You can only restore table data to a time when PITR was enabled.

**Service restore**

Recovery will be from Infra as Code.

#### *5.28.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

Amazon Keyspaces replicates data automatically three times in multiple AWS Availability Zones within the same AWS Region for durability and high availability.

### 5.28.4 Charging model

| Work Unit |
|---|
| Per Table within Keyspace |

### 5.28.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| Create a Keyspace | 1 token |
| Create table | 1 token |
| Restore PITR | Estimation in tokens based on table size |
| Other changes | Estimation in tokens based on time spent |

## 5.29 ElastiCache for Memcached

### 5.29.1 Description

Amazon ElastiCache for Memcached is a Memcached-compatible in-memory key-value store service that can be used as a cache or a data store. It delivers the performance, ease-of-use, and simplicity of Memcached. ElastiCache for Memcached is fully managed, scalable, and secure - making it an ideal candidate for use cases where frequently accessed data must be in-memory. It is a popular choice for use cases such as Web, Mobile Apps, Gaming, Ad-Tech, and E-Commerce.

### 5.29.2 Build to run service included in the OTC

#### *5.29.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.29.2.2 Build to run service*

- Refer to generic description.

### 5.29.3 RUN services included in the MRC

#### *5.29.3.1 Run service pre-requisite*

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

#### *5.29.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.29.3.3 KPI & alerts*

**Monitoring**

yes

**KPI monitored**

**Host-Level Metrics:**

- CPUUtilization
- CPUCreditBalance
- CPUCreditUsage
- FreeableMemory
- NetworkBytesIn
- NetworkBytesOut
- NetworkPacketsIn
- NetworkPacketsOut
- NetworkBandwidthInAllowanceExceeded
- NetworkConntrackAllowanceExceeded
- NetworkLinkLocalAllowanceExceeded
- NetworkBandwidthOutAllowanceExceeded
- Network Packets Per Second Allowance Exceeded
- SwapUsage

**Cache Node Level:**

- BytesReadIntoMemcached
- BytesUsedForCacheItems
- BytesWrittenOutFromMemcached
- CasBadval
- CasHits
- CasMisses
- CmdFlush
- CmdGets
- CmdSet
- CurrConnections
- CurrItems
- DecrHits
- DecrMisses
- DeleteHits
- DeleteMisses
- Evictions
- GetHits
- GetMisses

- IncrHits
- IncrMisses
- Reclaimed

**Alerts observed**

- CPUUtilization (Host Level): Memcached is multi-threaded, this metric can be as high as 90%. If you exceed this threshold, scale your cache cluster up by using a larger cache node type, or scale out by adding more cache nodes.

- Evictions (Node Level): This is a cache engine metric. We recommend that you determine your own alarm threshold for this metric based on your application needs. If you exceed your chosen threshold, scale your cluster up by using a larger node type, or scale out by adding more nodes.

- CurrConnections (Node Level): This is a cache engine metric. We recommend that you determine your own alarm threshold for this metric based on your application needs. An increasing number of CurrConnections might indicate a problem with your application; you will need to investigate the application behavior to address this issue.

#### *5.29.3.4 Backup and restore*

**Data backup and restore**

The backup feature is not available for Memcached Clusters.

**Service restore**

Recovery will be from Infra as Code.

#### *5.29.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

High Availability is not supported because since the service does not support replication.
When running the Memcached engine, you have the following options for minimizing the impact of a failure.

There are two types of failures to address in your failure mitigation plans:

- Node failure
- Availability Zone failure.

1. Mitigating Node Failures: spread your cached data over more nodes. Because Memcached does not support replication, a node failure will always result in some data loss from your cluster.

2. Mitigating Availability Zone Failures: locate your nodes in as many Availability Zones as possible.

### 5.29.4 Charging model

| Work Unit |
|---|
| Per node |

### 5.29.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| Modify Engine version compatibility | 2 tokens |
| Modify VPC Security Group | 2 tokens |
| Modify Parameter group | 2 tokens |

| | |
|---|---|
| Modify Maintenance Window | 2 tokens |
| Modify Topic for SNS Notification | 2 tokens |
| Reboot | 1 token |
| Other changes | Estimation in tokens based on time spent |

# 5.30 Amazon Aurora PostgreSQL Compatible

## 5.30.1 Description

Amazon Aurora PostgreSQL is a fully managed, PostgreSQL-compatible, and ACID-compliant relational database engine that combines the speed and reliability of high-end commercial databases with the simplicity and cost-effectiveness of open-source databases. Aurora PostgreSQL is a drop-in replacement for PostgreSQL and makes it simple and cost-effective to set up, operate, and scale your new and existing PostgreSQL deployments, thus freeing you to focus on your business and applications.

## 5.30.2 Build to run service included in the OTC

### *5.30.2.1 Build service pre-requisite*

- Refer to generic description.

### *5.30.2.2 Build to run service*

- Refer to generic description.

## 5.30.3 RUN services included in the MRC

### *5.30.3.1 Run service pre-requisite*

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

### *5.30.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

### *5.30.3.3 KPI & alerts*

**Monitoring**

yes

**Available metrics**

**Cluster level Metrics:**

- AuroraGlobalDBDataTransferBytes
- AuroraGlobalDBProgressLag
- AuroraGlobalDBReplicatedWriteIO
- AuroraGlobalDBReplicationLag
- AuroraGlobalDBRPOLag
- AuroraVolumeBytesLeftTotal
- BacktrackChangeRecordsCreationRate
- BacktrackChangeRecordsStored
- BackupRetentionPeriodStorageUsed
- ServerlessDatabaseCapacity

- SnapshotStorageUsed
- TotalBackupStorageBilled
- VolumeBytesUsed
- VolumeReadIOPs
- VolumeWriteIOPs

**Instance Level Metrics:**

- AbortedClients
- ActiveTransactions
- AuroraBinlogReplicaLag
- AuroraReplicaLag
- AuroraReplicaLagMaximum
- AuroraReplicaLagMinimum
- BacktrackWindowActual
- BacktrackWindowAlert
- BlockedTransactions
- BufferCacheHitRatio
- CommitLatency
- CommitThroughput
- CPUCreditBalance
- CPUCreditUsage
- CPUUtilization
- DatabaseConnections
- DDLLatency
- DDLThroughput
- Deadlocks
- DeleteLatency
- DeleteThroughput
- DiskQueueDepth
- DMLLatency
- DMLThroughput
- EBSByteBalance%
- EBSIOBalance%
- EngineUptime
- FreeableMemory
- FreeLocalStorage
- InsertLatency
- InsertThroughput
- LoginFailures
- MaximumUsedTransactionIDs
- NetworkReceiveThroughput
- NetworkThroughput
- NetworkTransmitThroughput
- NumBinaryLogFiles
- Queries
- RDSToAuroraPostgreSQLReplicaLag
- ReadIOPS
- ReadLatency
- ReadThroughput
- ReplicationSlotDiskUsage
- ResultSetCacheHitRatio
- RollbackSegmentHistoryListLength
- RowLockTime
- SelectLatency
- SelectThroughput
- StorageNetworkReceiveThroughput

- StorageNetworkThroughput
- StorageNetworkTransmitThroughput
- SumBinaryLogSize
- SwapUsage
- TransactionLogsDiskUsage
- UpdateLatency
- UpdateThroughput
- WriteIOPS
- WriteLatency
- WriteThroughput

**Alerts observed**

- WriteLatency
- ReadLatency
- FreeableMemory
- Deadlocks
- CPUUtilization
- DatabaseConnections
- BlockedTransactions
- BufferCacheHitRatio
- CommitLatency
- AbortedClients
- AuroraGlobalDBReplicationLag

#### *5.30.3.4 Backup and restore*

**Data backup and restore**

Aurora backs up your cluster volume automatically and retains restore data for the length of the backup retention period.

Aurora backups are continuous and incremental so you can quickly restore to any point within the backup retention period.

You can specify a backup retention period, from 1 to 35 days, when you create or modify a DB cluster.

Aurora backups are stored in Amazon S3.

AWS backup can be used as snapshot backup.

If you want to retain a backup beyond the backup retention period, you can also take a snapshot of the data in your cluster volume.

Restore will be done from backup to a new instance if restoring from backup or snapshot.

**Service restore**

Recovery will be from Infra as Code.

#### *5.30.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

An Aurora DB cluster is fault tolerant by design.
The cluster volume spans multiple Availability Zones in a single AWS Region, and each Availability Zone contains a copy of the cluster volume data.
This functionality means that your DB cluster can tolerate a failure of an Availability Zone without any loss of data and only a brief interruption of service.

### 5.30.4 Charging model

| Work Unit |
|---|
| per Database Instance |

### 5.30.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| Provision database | 2 tokens |
| Reboot an instance | 2 tokens |
| Delete an instance | 2 tokens |
| Instance failover | 2 tokens |
| Take snapshot of an instance | 2 tokens |
| Stop & start a cluster | 2 tokens |
| Delete a cluster | 2 tokens |
| Add reader Instance | 2 tokens |
| Add AWS region at cluster level | 2 tokens |
| Create clone at cluster level | 2 tokens |
| Restore a cluster to point in Time | Estimation in tokens based on the database size |
| Modify cluster configuration | 1 token |
| Upgrade a database | Estimation in tokens based on time spent |
| Minor Version patching | Estimation in tokens based on time spent |
| Export database snapshot to S3 | Estimation in tokens based on the size of extracted data |
| Other changes | Estimation in tokens based on time spent |

## 5.31 Amazon Aurora MySQL Compatible

### 5.31.1 Description

Amazon Aurora is a relational database management system (RDBMS) built for the cloud with full MySQL and PostgreSQL compatibility.

### 5.31.2 Build to run service included in the OTC

#### *5.31.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.31.2.2 Build to run service*

- Refer to generic description.

### 5.31.3 RUN services included in the MRC

#### *5.31.3.1 Run service pre-requisite*

- A referential file exists in the Git used by ORANGE BUSINESS which includes the reference configuration of the service.

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

#### *5.31.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.31.3.3 KPI & alerts*

**Monitoring**

yes

**Available metrics**

**Cluster level Metrics:**

- AuroraGlobalDBDataTransferBytes
- AuroraGlobalDBProgressLag
- AuroraGlobalDBReplicatedWriteIO
- AuroraGlobalDBReplicationLag
- AuroraGlobalDBRPOLag
- AuroraVolumeBytesLeftTotal
- BacktrackChangeRecordsCreationRate
- BacktrackChangeRecordsStored
- BackupRetentionPeriodStorageUsed
- ServerlessDatabaseCapacity
- SnapshotStorageUsed
- TotalBackupStorageBilled
- VolumeBytesUsed
- VolumeReadIOPs
- VolumeWriteIOPs

**Instance Level Metrics:**

- AbortedClients
- ActiveTransactions
- AuroraBinlogReplicaLag
- AuroraReplicaLag
- AuroraReplicaLagMaximum
- AuroraReplicaLagMinimum
- BacktrackWindowActual
- BacktrackWindowAlert
- BlockedTransactions
- BufferCacheHitRatio
- CommitLatency
- CommitThroughput
- CPUCreditBalance
- CPUCreditUsage
- CPUUtilization
- DatabaseConnections
- DDLLatency
- DDLThroughput
- Deadlocks
- DeleteLatency
- DeleteThroughput
- DiskQueueDepth

- DMLLatency
- DMLThroughput
- EBSByteBalance%
- EBSIOBalance%
- EngineUptime
- FreeableMemory
- FreeLocalStorage
- InsertLatency
- InsertThroughput
- LoginFailures
- MaximumUsedTransactionIDs
- NetworkReceiveThroughput
- NetworkThroughput
- NetworkTransmitThroughput
- NumBinaryLogFiles
- Queries
- RDSToAuroraPostgreSQLReplicaLag
- ReadIOPS
- ReadLatency
- ReadThroughput
- ReplicationSlotDiskUsage
- ResultSetCacheHitRatio
- RollbackSegmentHistoryListLength
- RowLockTime
- SelectLatency
- SelectThroughput
- StorageNetworkReceiveThroughput
- StorageNetworkThroughput
- StorageNetworkTransmitThroughput
- SumBinaryLogSize
- SwapUsage
- TransactionLogsDiskUsage
- UpdateLatency
- UpdateThroughput
- WriteIOPS
- WriteLatency
- WriteThroughput

**Alerts observed**

- WriteLatency
- ReadLatency
- FreeableMemory
- Deadlocks
- CPUUtilization
- DatabaseConnections
- BlockedTransactions
- BufferCacheHitRatio
- CommitLatency
- AbortedClients
- AuroraGlobalDBReplicationLag

#### *5.31.3.4 Backup and restore*

**Data backup and restore**

Aurora backs up your cluster volume automatically and retains restore data for the length of the backup retention period.

Aurora backups are continuous and incremental so you can quickly restore to any point within the backup retention period.

You can specify a backup retention period, from 1 to 35 days, when you create or modify a DB cluster.

Aurora backups are stored in Amazon S3.

AWS backup can be used as snapshot backup.

If you want to retain a backup beyond the backup retention period, you can also take a snapshot of the data in your cluster volume.

Restore will be done from backup to a new instance if restoring from backup or snapshot.

With Amazon Aurora MySQL-Compatible Edition, you can backtrack a DB cluster to a specific time, without restoring data from a backup.
Backtracking it lets you quickly move an Aurora database to a prior point in time without needing to restore data from a backup.
This lets you quickly recover from user errors, such as dropping the wrong table or deleting the wrong row.

**Service restore**

Recovery will be from Infra as Code.

#### *5.31.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

An Aurora DB cluster is fault tolerant by design.
The cluster volume spans multiple Availability Zones in a single AWS Region, and each Availability Zone contains a copy of the cluster volume data.
This functionality means that your DB cluster can tolerate a failure of an Availability Zone without any loss of data and only a brief interruption of service.

### 5.31.4 Charging model

| Work Unit |
|---|
| per Database Instance |

### 5.31.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| Provision database | 2 tokens |
| Reboot an instance | 2 tokens |
| Delete an instance | 2 tokens |
| Instance failover | 2 tokens |
| Take snapshot of an instance | 2 tokens |
| Stop & start a cluster | 2 tokens |
| Delete a cluster | 2 tokens |
| Add reader Instance | 2 tokens |
| Add AWS region at cluster level | 2 tokens |
| Create clone at cluster level | 2 tokens |

| | |
|---|---|
| Restore a cluster to point in Time | Estimation in tokens based on the database size |
| Modify cluster configuration | 1 token |
| Export database snapshot to S3 | Estimation in tokens based on the size of extracted data |
| Upgrade a database | Estimation in tokens based on time spent |
| Minor Version patching | Estimation in tokens based on time spent |
| Other changes | Estimation in tokens based on time spent |

# 5.32 Amazon Quantum Ledger Database

## 5.32.1 Description

Amazon QLDB is a new class of database that helps eliminate the need to engage in the complex development effort of building your own ledger-like applications. With QLDB, the history of changes to your data is immutable—it can't be altered, updated, or deleted. And using cryptography, you can easily verify that there have been no unintended changes to your application's data. QLDB uses an immutable transactional log, known as a journal. The journal is append-only and is composed of a sequenced and hash-chained set of blocks that contain your committed data.

## 5.32.2 Build to run service included in the OTC

### *5.32.2.1 Build service pre-requisite*

- Refer to generic description.

### *5.32.2.2 Build to run service*

- Refer to generic description.

## 5.32.3 RUN services included in the MRC

### *5.32.3.1 Run service pre-requisite*

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

### *5.32.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

### *5.32.3.3 KPI & alerts*

**Monitoring**

yes

**Available metrics**

- JournalStorage
- IndexedStorage
- ReadIOs
- WriteIOs

- CommandLatency
- IsImpaired
- OccConflictExceptions
- Session4xxExceptions
- Session5xxExceptions
- SessionRateExceededExceptions

**Alerts observed**

- JournalStorage
- IndexedStorage
- ReadIOs
- WriteIOs
- CommandLatency

#### *5.32.3.4 Backup and restore*

**Data backup and restore**

QLDB doesn't provide a dedicated backup and related restore feature at this time.

QLDB provides an on-demand journal export feature. You can access the contents of your journal by exporting journal blocks from your ledger into an Amazon Simple Storage Service (Amazon S3) bucket. You can use this data for various purposes such as data retention, analytics, and auditing.

**Service restore**

Recovery will be from Infra as Code.

#### *5.32.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

The service is Highly Available by design by AWS.

QLDB journal storage features synchronous replication to multiple Availability Zones on transaction commits. This ensures that even a full Availability Zone failure of journal storage would not compromise data integrity or the ability to maintain an active service. Additionally, the QLDB journal features asynchronous archives to fault-tolerant storage. This feature supports disaster recovery in the highly unlikely event of simultaneous storage failure for multiple Availability Zones.
QLDB doesn't provide an automated recovery feature for logical corruption scenarios at this time.

### 5.32.4 Charging model

| Work Unit |
| --- |
| per Table within ledger |

### 5.32.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |
| Create Ledger | 1 token |
| Other changes | Estimation in tokens based on time spent |

# 5.33 Microsoft SQL Server on Amazon RDS

## 5.33.1 Description

SQL Server is a relational database management system developed by Microsoft. Amazon RDS for SQL Server makes it easy to set up, operate, and scale SQL Server deployments in the cloud. Amazon RDS frees you up to focus on application development by managing time-consuming database administration tasks including provisioning, backups, software patching, monitoring, and hardware scaling.
Amazon RDS supports DB instances running several versions and editions of Microsoft SQL Server. For the full list of supported versions, editions, and RDS engine versions, see Microsoft SQL Server versions on Amazon RDS.
For information about licensing for SQL Server, see Licensing Microsoft SQL Server on Amazon RDS.
For information about SQL Server builds, see this Microsoft support article about the latest SQL Server builds.

## 5.33.2 Build to run service included in the OTC

### *5.33.2.1 Build service pre-requisite*

- Refer to generic description.

### *5.33.2.2 Build to run service*

- Refer to generic description.

## 5.33.3 RUN services included in the MRC

### *5.33.3.1 Run service pre-requisite*

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

### *5.33.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

### *5.33.3.3 KPI & alerts*

**Monitoring**

yes

**Available metrics**

- BinLogDiskUsage
- BurstBalance
- CPUUtilization
- CPUCreditUsage
- CPUCreditBalance
- DatabaseConnections
- DiskQueueDepth
- EBSByteBalance%
- EBSIOBalance%
- FailedSQLServerAgentJobsCount
- FreeableMemory
- FreeLocalStorage
- FreeStorageSpace
- MaximumUsedTransactionIDs

- NetworkReceiveThroughput
- NetworkTransmitThroughput
- OldestReplicationSlotLag
- ReadIOPS
- ReadIOPSLocalStorage
- ReadLatency
- ReadLatencyLocalStorage
- ReadThroughput
- ReadThroughputLocalStorage
- ReplicaLag
- ReplicationSlotDiskUsage
- SwapUsage
- TransactionLogsDiskUsage
- TransactionLogsGeneration
- WriteIOPS
- WriteIOPSLocalStorage
- WriteLatency
- WriteLatencyLocalStorage
- WriteThroughput
- WriteThroughputLocalStorage

**Alerts observed**

- DatabaseConnections
- FreeStorageSpace
- FreeableMemory
- ReadLatency
- ReadThroughput
- WriteLatency
- WriteThroughput
- ReadIOPS
- DiskQueueDepth
- WriteIOPS
- NetworkTransmitThroughput
- NetworkReceiveThroughput
- SwapUsage
- EBSByteBalance%
- EBSIOBalance%
- CPUSurplusCreditBalance
- CPUCreditUsage
- CPUCreditBalance
- CPUSurplusCreditsCharged
- CPUUtilization
- BurstBalance

### *5.33.3.4 Backup and restore*

**Data backup and restore**

Amazon RDS creates and saves automated backups of your DB instance during the backup window of your DB instance. RDS creates a storage volume snapshot of your DB instance, backing up the entire DB instance and not just individual databases. RDS saves the automated backups of your DB instance

according to the backup retention period that you specify. If necessary, you can recover your database to any point in time during the backup retention period.

You can also back up your DB instance manually, by manually creating a DB snapshot.

The first snapshot of a DB instance contains the data for the full DB instance. Subsequent snapshots of the same DB instance are incremental, which means that only the data that has changed after your most recent snapshot is saved.

You can set the backup retention period to between 0 and 35 days.

You can have up to 100 manual snapshots per region. Manual snapshot limits (100 per Region) do not apply to automated backups.

Restore will be done from backup depending on the option chosen by the customer: full backup or point in time recovery (automated backup).

**Service restore**

Recovery will be from Infra as Code.

#### *5.33.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

Multi-AZ deployments provide increased availability, data durability, and fault tolerance for DB instances. Multi-AZ deployments for SQL Server are implemented using SQL Server's native DBM or AGs (Availability Groups) technology.

### 5.33.4 Charging model

| Work Unit |
|---|
| per DB Instance |

### 5.33.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| Creating a database user | 1 Token |
| Dropping a Microsoft SQL Server database | 1 Token |
| Resetting the db_owner role password | 1 Token |
| Restoring license-terminated DB instances | 1 Token |
| Transitioning a Microsoft SQL Server database from OFFLINE to ONLINE | 1 Token |
| Enable Change data capture (CDC) | 1 Token |
| Enable and modify Database Mail | 1 Token |
| Native backup and restore | size * tokens |
| Amazon S3 file transfer | Estimation in tokens based on file size |
| Enable Microsoft Distributed Transaction Coordinator (MSDTC) | 1 Token |
| Enable Microsoft Business Intelligence (MSBI) | 1 Token |
| Enable Microsoft SQL Server Integration Services (SSIS) | 1 Token |
| Enable Microsoft SQL Server Reporting Services (SSRS) | 1 Token |

| | |
|---|---|
| SQL Server Audit (track the changes not related to the data (table creation, user creation, etc.) | 1 Token |
| Other changes | Estimation in tokens based on time spent |

## 5.34 Amazon RDS for MariaDB

### 5.34.1 Description

Amazon RDS supports an array of database engines to store and organize data among which MariaDB. It also helps with relational database management tasks, such as data migration, backup, recovery, and patching.

### 5.34.2 Build to run service included in the OTC

#### *5.34.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.34.2.2 Build to run service*

- Refer to generic description.

### 5.34.3 RUN services included in the MRC

#### *5.34.3.1 Run service pre-requisite*

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

#### *5.34.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.34.3.3 KPI & alerts*

**Monitoring**

Yes

**Available metrics**

- BinLogDiskUsage
- BurstBalance
- CPUUtilization
- CPUCreditUsage
- CPUCreditBalance
- DatabaseConnections
- DiskQueueDepth
- EBSByteBalance%
- EBSIOBalance%
- FailedSQLServerAgentJobsCount
- FreeableMemory
- FreeLocalStorage
- FreeStorageSpace
- MaximumUsedTransactionIDs
- NetworkReceiveThroughput

- NetworkTransmitThroughput
- OldestReplicationSlotLag
- ReadIOPS
- ReadIOPSLocalStorage
- ReadLatency
- ReadLatencyLocalStorage
- ReadThroughput
- ReadThroughputLocalStorage
- ReplicaLag
- ReplicationSlotDiskUsage
- SwapUsage
- TransactionLogsDiskUsage
- TransactionLogsGeneration
- WriteIOPS
- WriteIOPSLocalStorage
- WriteLatency
- WriteLatencyLocalStorage
- WriteThroughput
- WriteThroughputLocalStorage

**Alerts observed**

- DatabaseConnections
- FreeStorageSpace
- FreeableMemory
- ReadLatency
- ReadThroughput
- WriteLatency
- WriteThroughput
- ReadIOPS
- DiskQueueDepth
- WriteIOPS
- NetworkTransmitThroughput
- NetworkReceiveThroughput
- SwapUsage
- EBSByteBalance%
- EBSIOBalance%
- CPUSurplusCreditBalance
- CPUCreditUsage
- CPUCreditBalance
- CPUSurplusCreditsCharged
- CPUUtilization
- BurstBalance

### *5.34.3.4 Backup and restore*

**Data backup and restore**

Amazon RDS creates and saves automated backups of your DB instance during the backup window of your DB instance. RDS creates a storage volume snapshot of your DB instance, backing up the entire DB instance and not just individual databases. RDS saves the automated backups of your DB instance

according to the backup retention period that you specify. If necessary, you can recover your database to any point in time during the backup retention period.

You can also back up your DB instance manually, by manually creating a DB snapshot.

The first snapshot of a DB instance contains the data for the full DB instance. Subsequent snapshots of the same DB instance are incremental, which means that only the data that has changed after your most recent snapshot is saved.

You can set the backup retention period to between 0 and 35 days.

You can have up to 100 manual snapshots per region. Manual snapshot limits (100 per Region) do not apply to automated backups.

Restore will be done from backup depending on the option chosen by the customer: full backup or point in time recovery (automated backup).

**Service restore**

Recovery will be from Infra as Code.

#### *5.34.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

For your MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server database (DB) instances, you can use Amazon RDS Multi-AZ deployments. When you provision a Multi-AZ DB instance, Amazon RDS automatically creates a primary DB instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). In case of an infrastructure failure, Amazon RDS performs an automatic failover to the standby DB instance. Since the endpoint for your DB instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.

### 5.34.4 Charging model

| Work Unit |
|---|
| per DB Instance |

### 5.34.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| Provision a database | 2 Tokens |
| Reboot a DB instance | 2 Tokens |
| Delete a DB instance | 2 Tokens |
| Stop a DB instance | 2 Tokens |
| Create replica Database | 2 Tokens |
| Take a snapshot | 2 Tokens |
| Restore to point in time (Instance level) | 2 Tokens |
| Start database activity stream (Instance level) | 2 Tokens |
| Restore from S3 (Instance level) | Estimation in tokens based on database size |
| Other changes | Estimation in tokens based on time spent |

# 5.35 Amazon RDS for Oracle

## 5.35.1 Description

Amazon RDS is a managed service for relational databases, including Oracle. RDS is offered with AWS licensing or in a bring your own license (BYOL) model. Once you set up your Oracle database on RDS, you can use the AWS platform to monitor, configure, backup, secure, and scale your workloads.

## 5.35.2 Build to run service included in the OTC

### *5.35.2.1 Build service pre-requisite*

- Refer to generic description.

### *5.35.2.2 Build to run service*

- Refer to generic description.

## 5.35.3 RUN services included in the MRC

### *5.35.3.1 Run service pre-requisite*

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

### *5.35.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

### *5.35.3.3 KPI & alerts*

**Monitoring**

yes

**Available metrics**

- BinLogDiskUsage
- BurstBalance
- CPUUtilization
- CPUCreditUsage
- CPUCreditBalance
- DatabaseConnections
- DiskQueueDepth
- EBSByteBalance%
- EBSIOBalance%
- FailedSQLServerAgentJobsCount
- FreeableMemory
- FreeLocalStorage
- FreeStorageSpace
- MaximumUsedTransactionIDs
- NetworkReceiveThroughput
- NetworkTransmitThroughput
- OldestReplicationSlotLag
- ReadIOPS
- ReadIOPSLocalStorage
- ReadLatency

- ReadLatencyLocalStorage
- ReadThroughput
- ReadThroughputLocalStorage
- ReplicaLag
- ReplicationSlotDiskUsage
- SwapUsage
- TransactionLogsDiskUsage
- TransactionLogsGeneration
- WriteIOPS
- WriteIOPSLocalStorage
- WriteLatency
- WriteLatencyLocalStorage
- WriteThroughput
- WriteThroughputLocalStorage

**Alerts observed**

- DatabaseConnections
- FreeStorageSpace
- FreeableMemory
- ReadLatency
- ReadThroughput
- WriteLatency
- WriteThroughput
- ReadIOPS
- DiskQueueDepth
- WriteIOPS
- NetworkTransmitThroughput
- NetworkReceiveThroughput
- SwapUsage
- EBSByteBalance%
- EBSIOBalance%
- CPUSurplusCreditBalance
- CPUCreditUsage
- CPUCreditBalance
- CPUSurplusCreditsCharged
- CPUUtilization
- BurstBalance

### *5.35.3.4 Backup and restore*

**Data backup and restore**

Amazon RDS creates and saves automated backups of your DB instance during the backup window of your DB instance. RDS creates a storage volume snapshot of your DB instance, backing up the entire DB instance and not just individual databases. RDS saves the automated backups of your DB instance

according to the backup retention period that you specify. If necessary, you can recover your database to any point in time during the backup retention period.

You can also back up your DB instance manually, by manually creating a DB snapshot.

The first snapshot of a DB instance contains the data for the full DB instance. Subsequent snapshots of the same DB instance are incremental, which means that only the data that has changed after your most recent snapshot is saved.

You can set the backup retention period to between 0 and 35 days.

You can have up to 100 manual snapshots per region. Manual snapshot limits (100 per Region) do not apply to automated backups.

Restore will be done from backup depending on the option chosen by the customer: full backup or point in time recovery (automated backup).

**Service restore**

Recovery will be from Infra as Code.

#### *5.35.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

- **Automatic Host Replacement** – Amazon RDS will automatically replace the compute instance powering your deployment in the event of a hardware failure.

- **Multi-AZ Deployments** – A deployment option for your production DB Instances that enhances database availability while protecting your latest database updates against unplanned outages. When you create or modify your DB Instance to run as a Multi-AZ deployment, Amazon RDS will automatically provision and manage a “standby” replica in a different Availability Zone (independent infrastructure in a physically separate location). Database updates are made concurrently on the primary and standby resources to prevent replication lag. In the event of planned database maintenance, DB Instance failure, or an Availability Zone failure, Amazon RDS will automatically failover to the up-to-date standby so that database operations can resume quickly without administrative intervention. Prior to failover you cannot directly access the standby, and it cannot be used to serve read traffic.

### 5.35.4 Charging model

| Work Unit |
|---|
| per DB Instance |

### 5.35.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| Provision a database | 2 Tokens |
| Reboot a DB instance | 2 Tokens |
| Delete a DB instance | 2 Tokens |
| Stop a DB instance | 2 Tokens |
| Create replica Database | 2 Tokens |

| | |
|---|---|
| Take a snapshot | 2 Tokens |
| Restore to point in time (Instance level) | 2 Tokens |
| Start database activity stream (Instance level) | 2 Tokens |
| Restore from S3 (Instance level) | Estimation in tokens based on database size |
| Other changes | Estimation in tokens based on time spent |

# 5.36 Amazon RDS for PostgreSQL

## 5.36.1 Description

Amazon RDS for Oracle is a fully managed commercial database that makes it easy to set up, operate, and scale Oracle deployments in the cloud.
Amazon RDS currently supports the following versions of MariaDB:

## 5.36.2 Build to run service included in the OTC

### *5.36.2.1 Build service pre-requisite*

- Refer to generic description.

### *5.36.2.2 Build to run service*

- Refer to generic description.

## 5.36.3 RUN services included in the MRC

### *5.36.3.1 Run service pre-requisite*

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

### *5.36.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

### *5.36.3.3 KPI & alerts*

**Monitoring**

yes

**Available metrics**

- BinLogDiskUsage
- BurstBalance
- CPUUtilization
- CPUCreditUsage
- CPUCreditBalance
- DatabaseConnections
- DiskQueueDepth
- EBSByteBalance%
- EBSIOBalance%
- FailedSQLServerAgentJobsCount

- FreeableMemory
- FreeLocalStorage
- FreeStorageSpace
- MaximumUsedTransactionIDs
- NetworkReceiveThroughput
- NetworkTransmitThroughput
- OldestReplicationSlotLag
- ReadIOPS
- ReadIOPSLocalStorage
- ReadLatency
- ReadLatencyLocalStorage
- ReadThroughput
- ReadThroughputLocalStorage
- ReplicaLag
- ReplicationSlotDiskUsage
- SwapUsage
- TransactionLogsDiskUsage
- TransactionLogsGeneration
- WriteIOPS
- WriteIOPSLocalStorage
- WriteLatency
- WriteLatencyLocalStorage
- WriteThroughput
- WriteThroughputLocalStorage

**Alerts observed**

- CPU Utilization
- DB Connections
- Write IOPS
- Read IOPS
- Queue Depth
- Freeable Memory
- Swap Usage
- Write Latency
- Read Latency
- Write Throughput
- Read Throughput

### *5.36.3.4 Backup and restore*

**Data backup and restore**

Amazon RDS creates and saves automated backups of your DB instance during the backup window of your DB instance. RDS creates a storage volume snapshot of your DB instance, backing up the entire DB instance and not just individual databases. RDS saves the automated backups of your DB instance

according to the backup retention period that you specify. If necessary, you can recover your database to any point in time during the backup retention period.

You can also back up your DB instance manually, by manually creating a DB snapshot.

The first snapshot of a DB instance contains the data for the full DB instance. Subsequent snapshots of the same DB instance are incremental, which means that only the data that has changed after your most recent snapshot is saved.

You can set the backup retention period to between 0 and 35 days.

You can have up to 100 manual snapshots per region. Manual snapshot limits (100 per Region) do not apply to automated backups.

Restore will be done from backup depending on the option chosen by the customer: full backup or point in time recovery (automated backup).

**Service restore**

Recovery will be from Infra as Code.

#### *5.36.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

For your MySQL, MariaDB, PostgreSQL, Oracle, and SQL Server database (DB) instances, you can use Amazon RDS Multi-AZ deployments. When you provision a Multi-AZ DB instance, Amazon RDS automatically creates a primary DB instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). In case of an infrastructure failure, Amazon RDS performs an automatic failover to the standby DB instance. Since the endpoint for your DB instance remains the same after a failover, your application can resume database operation without the need for manual administrative intervention.

### 5.36.4 Charging model

| Work Unit |
|---|
| per DB Instance |

### 5.36.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
|---|---|
| Provision database | 2 tokens |
| Reboot an instance | 2 tokens |
| Delete an instance | 2 tokens |
| Instance failover | 2 tokens |
| Take snapshot of an instance | 2 tokens |
| Stop & start a cluster | 2 tokens |
| Delete a cluster | 2 tokens |
| Add reader Instance | 2 tokens |
| Add AWS region at cluster level | 2 tokens |
| Create clone at cluster level | 2 tokens |
| Restore a cluster to point in Time | Estimation in tokens based on the database size |
| Modify cluster configuration | 1 token |
| Upgrade a database | Estimation in tokens based on time spent |

| | |
|---|---|
| Minor Version patching | Estimation in tokens based on time spent |
| Export database snapshot to S3 | Estimation in tokens based on the size of extracted data |
| Other changes | Estimation in tokens based on time spent |

## 5.37 Amazon DocumentDB

### 5.37.1 Description

Amazon DocumentDB is an AWS database service that is fully managed and compatible with MongoDB. You can use this service to migrate and host MongoDB workloads and application data while working with native Mongo code, tools, and drivers.
Through DocumentDB you gain access to the following features:

- Automatic scaling to match the size of your storage needs. Scaling occurs in increments of 10GB up to 64TB.
- Ability to create up to 15 read replicas for higher throughput. Storage is shared so writes only need to be performed to centralized volumes, not duplicated across replicas.
- Enables you to scale memory and compute resources independently for greater flexibility and cost optimization.
- Operates in a Virtual Private Cloud (VPC) with firewalls for greater isolation and security.
- Supports encryption with keys managed in AWS Key Management Service (AWS KMS). Active data, backups, replicas, and snapshots are all encrypted.

### 5.37.2 Build to run service included in the OTC

#### *5.37.2.1 Build service pre-requisite*

- Refer to generic description.

#### *5.37.2.2 Build to run service*

- Refer to generic description.

### 5.37.3 RUN services included in the MRC

#### *5.37.3.1 Run service pre-requisite*

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

#### *5.37.3.2 Reporting*

By default, no. Reporting can be requested by customer through change request to have point in time report.

#### *5.37.3.3 KPI & alerts*

**Monitoring**

yes

**Available metrics**

Amazon CloudWatch metrics for Amazon DocumentDB are available at: **Monitoring Amazon DocumentDB with CloudWatch - Amazon DocumentDB**

**Alerts observed**

- DatabaseConnections
- DatabaseConnectionsMax
- NetworkThroughput
- CPUUtilization
- CPUCreditUsage
- CPUCreditBalance
- FreeLocalStorage
- FreeableMemory
- BufferCacheHitRatio
- VolumeBytesUsed
- VolumeReadIOPs
- VolumeWriteIOPs

### *5.37.3.4 Backup and restore*

**Data backup and restore**

Amazon DocumentDB (with MongoDB compatibility) continuously backs up your data to Amazon Simple Storage Service (Amazon S3) for 1–35 days so that you can quickly restore to any point within the backup retention period. Amazon DocumentDB also takes automatic snapshots of your data as part of this continuous backup process.

You can also retain backup data beyond the backup retention period by creating a manual snapshot of your cluster's data. The backup process does not impact your cluster's performance.

**Service restore**

Recovery will be from Infra as Code.

### *5.37.3.5 AWS SLA High Availability and Disaster Recovery inter-region*

You can achieve high availability and read scaling in Amazon DocumentDB (with MongoDB compatibility) by using replica instances. A single Amazon DocumentDB cluster supports a single primary instance and up to 15 replica instances. These instances can be distributed across Availability Zones within the cluster's Region. The primary instance accepts read and write traffic, and replica instances accept only read requests.

**Sharding Option:**

No. Amazon DocumentDB's distributed storage architecture is a different approach to scaling than MongoDB sharding.

### 5.37.4 Charging model

| Work Unit |
| --- |
| per DB Instance |

### 5.37.5 Changes catalogue – in Tokens, per act

| Changes examples | Effort |
| --- | --- |
| Adding a Replica to an Amazon DocumentDB Cluster | 1 token |
| Creating a Cluster Snapshot | 1 token |
| Restoring from a Snapshot | 1 token |
| Removing an Instance from a Cluster | 1 token |
| Deleting a Cluster | 1 token |
| Other changes | Estimation in tokens based on time spent |

# 6 End of the document