



Business

Flexible Engine

Security white paper

Version 3.2

Orange Business

11/08/2023

Document description

Properties				
Document title	Security White paper – Flexible Engine			
Version	3.2			
Author	Orange Cloud for Business Security			
Status	<input type="checkbox"/> In progress	<input type="checkbox"/> Reviewed	<input type="checkbox"/> Validated	<input checked="" type="checkbox"/> Approved
Date	August 11th 2023			

Document Classification

Classification	
Confidentiality	Orange Unrestricted

Version history

Version	Operation	Name	Date
1.0	Initial Release	Orange Cloud for Business Security	April 2018
1.1	First Revision	Orange Cloud for Business Security	May 2018
1.2	Second Revision	Orange Cloud for Business Security	August 2018
1.3	Third Revision	Orange Cloud for Business Security	September 2018
2.0	Second Release	Orange Cloud for Business Security	June 2019
2.1	Fourth Revision	Orange Cloud for Business Security	October 2019
3.1	Fifth Revision	Orange Cloud for Business Security	September 2021
3.2	Revision	Orange Cloud for Business Security	August 2023

Notice

All the information in this document is provided to the customer for informational purpose only. It represents Orange Business current and future product offers as of the date of issue of this document, which are subject to change.

© Copyright 2023 – Orange Business All rights reserved

1. Introduction

1.1. Purpose of the document

This document is the Security white paper of the Cloud Computing services of “Flexible Engine”, by Orange Business. It describes the main technical and organizational security measures applied by Orange Business to guarantee the security of Flexible Engine services and the protection of customers’ data.

1.2. Document organization

- This document is organized into the following chapters:
- Chapter 1 is the document introduction;
- Chapter 2 is a summary of the benefits of Cloud Orange in matters of security;
- Chapter 3 explains the Shared Responsibility Model for Flexible Engine (as Cloud Service Provider) and the customers (as Tenants owners);
- Chapter 4 lists the security measures used, organized according to the sections of ISO 27002;
- Chapter 5 details the extra security measures specific to Flexible Engine;
- Appendix A describes the tier classification of a datacenter.

2. The security strengths of Orange Business’ cloud services

This chapter lists the main security benefits of Orange Business’ Cloud services:

- **A Trusted Partner:** Orange Business is certified ISO 9001, ISO27001¹ and ISAE 3402 type II² (SOC1 report, former SAS70) and MTCS level 3³. These certifications provide cloud customers with assurance from 3rd party auditors that relevant cloud computing security practices and controls are in place and demonstrated.
- **Data location:** Orange Business guarantees the location of customers’ data and backups in one or more given countries.
- **Physical security of datacenters:** The security of all datacenters is periodically assessed to comply with Orange security and availability requirements. All datacenters used for Flexible Engine are certified ISO27001, ISAE 3402 Type II (ex-SAS 70) and comply with TIER III or TIER IV requirements.
- **Operator experience:** The Cloud services have been designed on the basis of Orange’s experience as a critical infrastructure operator and as one of the French leaders in security integration and consulting with Orange Cyberdefense.

¹ Associated certificate: <https://certificats-attestations.afnor.org/certification=335181233155>

² ISAE 3402: standard ensuring customers of externalized services of the reliability of the internal control mechanism of their service provision. In particular, the standard covers the physical security of datacenters.

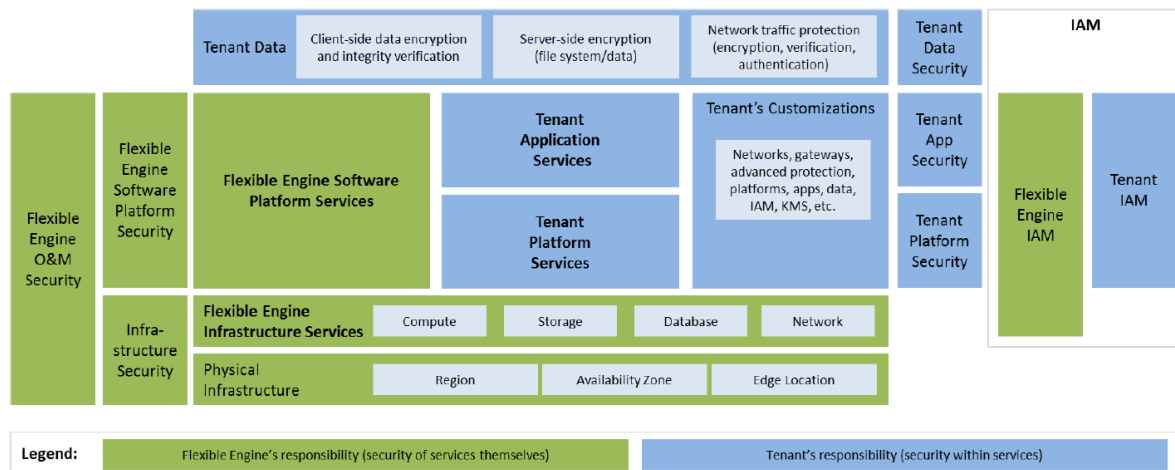
³ Associated certificate: <https://www2.imda.gov.sg/regulations-and-licensing-listing/ict-standards-and-quality-of-service/IT-Standards-and-Frameworks/ComplianceAndCertification>

- **Support and maintenance operation security:** the operating centers are located in France and abroad and have ISO 27001 security certification, ensuring the confidentiality of Orange customers' information. Administration hosts allow to strictly secure and control the administration rights, privileges and actions.
- **Administration portal security:** customers are provided with portals allowing them to manage their Cloud services. Access to these portals is realized through personal accounts, actions are logged and data flows are protected and encrypted. Portal security is regularly checked with intrusion tests.
- **Connectivity with company VPNs:** Cloud customers' virtual environments can be securely interconnected with a Business VPN MPLS, providing greater end-to-end network isolation, together with a bandwidth and network service quality that the Internet cannot steadily provide.
- **Protection against intrusion and denial of service:** The infrastructure used to operate Flexible Engine cloud services is protected against intrusions and DoS attacks. Customers can also be provided additional security protection on their virtual infrastructure or application (WAF, IDS, DDoS).
- **Service continuity of Orange Cloud services:** all Orange Cloud services come with high service availability rates (described in the service descriptions of each offer). Furthermore, all our Cloud services benefit from entirely redundant infrastructure, with high-availability mechanisms. Any local interruption is consequently transparent: network (Internet access, VPN access, firewall), administration portal, virtualization, storage, etc. These measures are detailed in the rest of this document.

3. Shared responsibility model

In three primary cloud service delivery mechanisms defined by NIST⁴, as infrastructure as a service (IaaS), platform as a service (PaaS) and software as service (SaaS), the tenants maintain complete control over their content, services, data location, encryption as well as access control. As a CSP, Orange Business is responsible for maintaining the security of applications, platform and infrastructure to ensure confidentiality, integrity and availability of Flexible Engine. This shared responsibility model is fundamental to understand the respective roles of the tenants and Flexible Engine in the context of cloud security principles.

⁴ 4 NIST: National Institute of Standards and Technology



As shown in the above figure, the primary responsibilities of Flexible Engine are developing and operating the physical infrastructure of Flexible Engine datacenters; the IaaS, PaaS, and SaaS services provided by Flexible Engine; and the built-in security functions of a variety of services.

The primary responsibilities of the tenant are customizing the configuration and operating the virtual network, platform, application, data, management, security, and other cloud services to which a tenant subscribes on Flexible Engine, including its customization of Flexible Engine services according to its needs as well as the O&M of any platform, application, and IAM services that the tenant deploys on Flexible Engine.

3.1. Flexible Engine security responsibilities

Flexible Engine is responsible for protecting the security of our IaaS, PaaS and SaaS services, as well as the physical environments of the Flexible Engine global datacenters on which our cloud services operate. Flexible Engine is not only responsible for the security, performance of the infrastructure, but also the overall security compliance of our infrastructure and services to respective standards and regulations.

In addition to protecting the global infrastructure, Flexible Engine is also maintaining the products that are considered managed services, such as RDS, MRS, Workspace, and several

others. Following functions, but not limited to, are implemented fulfil the O&M security requirements of Flexible Engine:

- Rapid security incident detection, isolation, and response to ensure fast recovery of cloud services.
- Vulnerability management mechanism to track, test, validate and implement across the platform.
- Secure configuration and version upkeep of cloud services.
- Basic security tasks like operating system update and middleware patching, firewall configuration and disaster recovery for managed services.

3.2. Tenant security responsibilities

Within Flexible Engine, tenants are responsible for security inside the IaaS, PaaS, and SaaS cloud services which they subscribe and provision. This includes the security configurations

and operations of Elastic Cloud Servers (ECS), Virtual Private Cloud (VPC), security groups, advanced security services, key and identity management. Tenants should pay attention to but not limited to following requirements when they are using Flexible Engine services.

- Security configurations of tenant-managed services, such as virtual firewalls, gateways, advanced security services, and security management tasks (patch management, hardening).
- Protection of Flexible Engine account credentials and provision of individual user accounts with IAM and setting up multi-factor authentication, secure data transfer protocols in accordance with industry best practices.
- Adequate testing of cloud services that are deployed on Flexible Engine to prevent adverse effects on applications and to minimize business impact.
- Creating user accounts based on least privilege principle and enforces segregation of duties, utilizing all available security features provided by Flexible Engine such as logging, access control lists, permissions on applicable services, such as RDS, OBS, IMS.
- While Flexible Engine strives to meet regulatory and industry security compliance as a CSP, tenants are responsible to comply any application and service that they deploy and operates on Flexible Engine that is not part of our services.

3.3. Cloud Alliance member responsibilities

In order to serve multinational corporations with presence across different regions other than Flexible Engine, Orange has become a member of Cloud Alliance which is the combination of majors IT and Telco actors with Huawei (Equipment and Cloud Provider) for creating a world class leader in Cloud Service. This unique alliance allows its customers to benefit of a worldwide footprint for their cloud ambitions with no less than 19 Availability Zones. A Single Sign On combination allows an alliance customer to operate through all partners solutions on a shared and secure network across the world.

All members of the Cloud Alliance have based their security approach on the “Shared Responsibility Model” and are using Huawei OpenStack public cloud platform. Huawei Cloud in particular has started global certification programs⁵ including ISO 27001, ISO 27017, ISO 27018, PCI-DSS, CSA-STAR, SOC 1/2/3 and regional programs such as Singapore Multi-Tier Cloud Security (MTCS), China Trusted Cloud Service (TRUCS) by the Data Center Alliance (DCA) and the China Academy of Information and Communications Technology (CAICT).

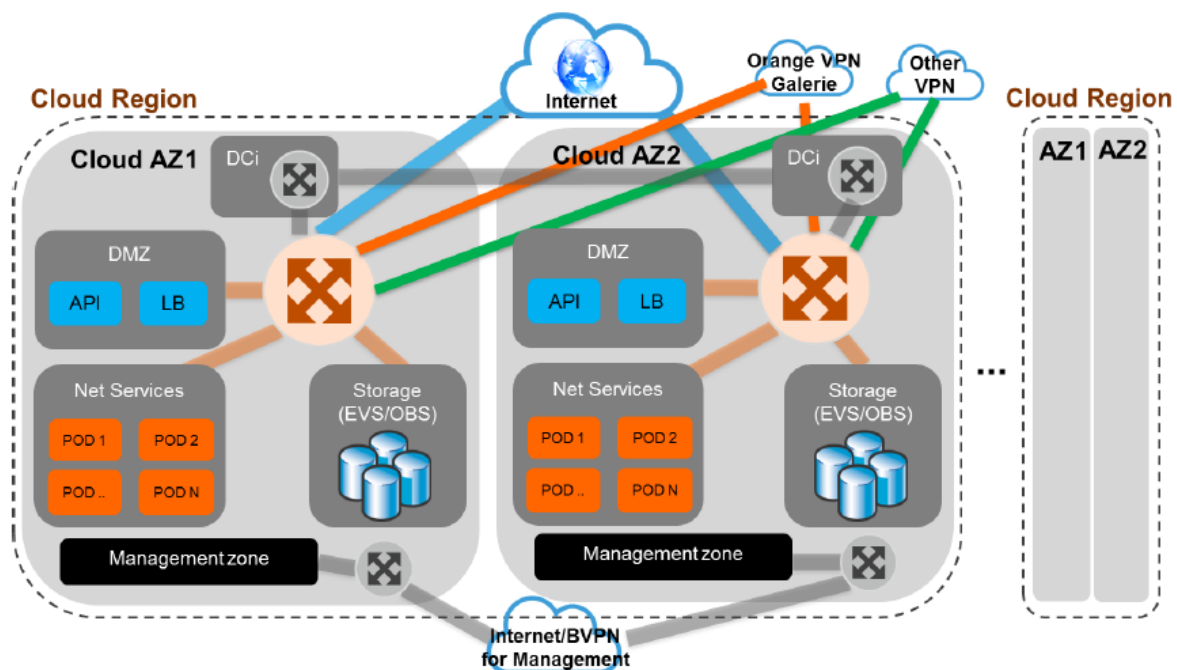
Cloud Alliance member providers will ensure the data transfer among the resources owned by the customers is secure and access controls are applied consistently regardless of availability zones and regions.

⁵ <https://www.huaweicloud.com/en-us/securecenter/compliance.html>

4. Common Security measures

Security measures are organized into 14 chapters, corresponding to the 14 sections in the ISO 27002 standard version 2013.

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



System virtualization

The Flexible Engine virtualization platform is based on the OpenSource XEN virtualization engine and KVM⁶ on CentOS. In addition to the standard security of this solution, the hypervisor has been hardened to strengthen its partitioning:

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

Segregated access to the stored data (block and object) is ensured by an application layer (OpenStack Cinder for storage block, compatible with AWS S3 for object storage) which allows data access only to data owners or the corresponding storage. In addition, data

⁶ Kernel-based Virtual Machine (KVM) is an open source virtualization technology built into Linux®. Specifically, KVM turns Linux into a hypervisor that allows a host machine to run multiple, isolated virtual environments called guests or virtual machines (VMs).

written to the infrastructure is not recoverable once it is deleted by the client or when the corresponding virtual infrastructure is terminated by the customer. These mechanisms are regularly tested thanks to intrusion tests carried out by Orange Cyber Defense and its partners; whose skills are recognized. Physical disks needing to be replaced are destroyed by grinding through a traceable and certified process.

Network virtualization

Virtual Private Cloud (VPC) features, carried by the Neutron OpenStack component, provide a logical partitioning of communications on the user network. Any form of network traffic that is not naturally authorized on the customers tenant is not processed by the devices supporting the client's virtual network, preventing any use of spoofing technologies. All tenants traffic are routed northwise/upwards where routing is done under layers of security provided by industry level firewalling with state of the art firewalling device.

4.1. Security policy

This chapter corresponds to section 5 in ISO 27002 version 2013.

Orange Cloud services are subject to security policies from Orange group or more dedicated to the Cloud organization. The latter are aligned information system security standards: ISO 27001 and ISO 27002, MTCS, SOC2 and SecNumCloud. The policies are at least yearly reviewed and regularly updated.

4.2. Organization of security

This chapter corresponds to section 6 in ISO 27002 version 2013.

This part describes the security organization of Orange's Cloud offers.



Security organization of Orange's Cloud offers

Business Security Officer






For customers with strong security requirements, a security expert (Business Security Officer) can be appointed to be the central point of contact of customer's CISO (Chief Information Security Officer) providing advice and detailed security reporting.

Orange Cloud for Business Security Department

The Orange Cloud for Business Security department ensures that security is observed in all the Orange services. Its tasks are, for example:

- To define security policy and set the guidelines to be applied to Cloud services;
- To carry out and update security risk analyses (ISO 27005 type) for Cloud services;
- To manage the security action plans resulting from risk analysis;
- To guarantee the homogeneity of security solutions deployed in order to ensure their effectiveness and durability;
- To capitalize on Orange's knowledge of Cloud security;
- To offer support and advice, to provide expertise in Cloud security issues;
- To train personnel and have them made aware of security procedures;
- To provide a security watch service and work with international authorities on Cloud security. Thus, Orange Business is a participant in the Cloud Security Alliance⁷ (CSA), ITU-T and ISO; it acts at a European level as a member of the Cloud Security Expert Group of ENISA⁸.
- To monitor the security level of services by conducting regular security audits (internal and/or external).

The Orange Cloud for Business Security department is run by the Chief Security Officer. It is made up of security experts involved, and is organized as described below:

	Guarantee security of OCB organization	Reinforce Security Value of our services	Billable security services
	CSO/CISO role	Business Security Officer (BSO) for major services	Support to sales/presales Dedicated BSO
	Organisational certifications (ISO, ISAE, SOC)	Service certifications (ISO, SecNumCloud, MTCS)	Staff security certifications (CISSP, ISO27005...)
	Security policies Compliance (GDPR, NIS, SOX...)	Security White paper	Security Assurance Plan Security questionnaires
	Audit and Control Security Incident Mgt Vulnerability Mgt Access Control Mgt Security Risk Mgt Transversal Security projects	Security incident follow-up Vulnerability follow-up SecurityInTTM process (risks) Penetration tests	Security incident follow-up Vulnerability follow-up On-demand audits and advice

⁷ Cloud Security Alliance (CSA): <https://cloudsecurityalliance.org/>

⁸ ENISA: European Union Agency for Network and Information Security (ENISA) - www.enisa.europa.eu

Orange Global Security

Orange Global Security (DSEC), attached to the General Secretary of the Group, initiates and coordinates security actions over the whole scope of the Orange Group. The department has the responsibility among other tasks:

- To define Security Policies (general, sectorial);
- To define the policies on personnel awareness and training in security and to promote security in the company strategy and culture;
- To interface with the main security and data protection entities:
 - French National Defense and Government authorities (SGDN, ANSSI, etc.);
 - French justice, police and gendarmerie authorities (in liaison with Orange Legal Management);
 - the CNIL (French National Commission for Data Protection and Freedom, data protection regulator) (in liaison with Orange Legal Management and General Secretary);
 - Orange Risk Audit and Monitoring Management (DACR), the National Center for the Security of Information Systems (Orange entity);
 - the CERTs (Computer Emergency Response Teams), in particular the CERT-IST.
- To centralize security incident information and consolidate the dashboard;
- To coordinate security certification programs such as ISMS ISO27001 certification;
- To conduct audits and assessments and to follow up on audits and assessments conducted on the initiative of the entities.

4.3. Human resource-related security

This chapter corresponds to section 7 in ISO 27002 version 2013.

4.3.1. Personnel listing

A list of the staff contributing to the Cloud services build and operations is kept updated as part of the input/output management process. This list is used as a reference base when reviewing access authorizations. Personnel's access rights are modified/removed whenever they change position or leave the company. Personnel with material access to operations are subjected to regular background checks.

4.3.2. Awareness

Information System security awareness training is given to each new Orange staff member and all subcontractors joining the Orange teams. This training is regularly updated.

Each entity manager keeps their staff informed of security-related risks and of the legal and regulatory obligations. They are assisted by experts (inside or outside their entity).

Each entity manager has their staff made aware of, and ensures compliance with, ethical and responsible behavior in various internal or external communication situations (emails, web, telephone, removable storage, seminars, travel, etc.).

Each entity manager makes sure that their internal staff is competent in matters of security, and if necessary arranges extra training.

4.3.3. Written undertaking of Staff

Staff confidentiality undertaking

Employees have a general obligation of discretion written into their contract of employment and/or into their collective agreement. The following is an extract from the Orange employment contract:

"Article 10: Professional secrecy and the duty of discretion
Like all Orange staff, Mr./Ms. _____ is bound by absolute professional secrecy. This stipulation applies during the performance and suspension of the employment contract and also after its termination and concerns in particular the techniques deployed by Orange, together with all design studies and work carried out in the company.
It also applies to all confidential information, in particular that which is not usually communicated to the public, which Mr./Ms. _____ may encounter in the performance of their functions or merely by their presence in Orange.
Failure to observe this secrecy requirement constitutes misconduct and may incur civil and/or criminal prosecution. "

For subcontractors, this undertaking is signed by the employee via their employer.

Acceptable use charter for computer resources

Similarly, each employee must, according to internal regulations, comply with the acceptable use charter for computer resources.

4.4. Asset management

This chapter corresponds to section 8 in ISO 27002 version 2013.

4.4.1. Asset inventory

The list of support assets (routers, switches, firewalls, storage bays, servers, virtual machines, etc.) in Cloud offers is kept in the Orange internal mapping tools. These tools trace both the components of the Orange Cloud infrastructure and the components of the customers' environments (e.g.: Virtual Machine, virtual firewalls, etc.)

These inventories are updated as part of the change management process.

4.4.2. Protective measures for support assets

Classification of infrastructure resources

Infrastructure resources are classified on a 4-level scale, defined in the Orange Global Security Policy, and this qualifies the harm to Orange and its customers in the event of an incident involving the resource:

- Level 4: vital impact, very high (even vital) stakes corresponding to inadmissible risks;
- Level 3: critical impact, very high stakes corresponding to risks whose effects must be limited;
- Level 2: substantial impact, moderately high stakes and risks controlled with limited harm to the company;
- Level 1: no or practically no impact.

All shared Cloud infrastructure resources hosted in our datacenters (routers, switches, storage and backup bays, servers, etc.) are considered to be at least level 3 (i.e. level 3 or

level 4). Measures for the protection of these critical resources are described throughout this document.

Classification, marking and protection of documentation

Documentation is classified and then marked in accordance with the Orange internal rules defined in the document "Document marking". Four levels are defined here: free circulation (unrestricted), Orange internal, Orange confidential and Orange secret.

Operational documents of Cloud projects / services are stored on file servers. Access is protected by personal access code and authorization management on a need-to-know basis.

Classification of customer information

The classification of customer information hosted on Orange's Cloud services is the customer's responsibility.

4.4.3. Data deletion

Protection of de-allocated customer data

Customer data located on de-allocated resources cannot be accessed by another customer. This protection is provided by the internal mechanisms of the products used by Orange. For example, OpenStack has an internal process to erase the resource before allocating it to another customer.

Deletion of customer data at termination of contract

When the contract terminates, all the resources allocated to customers are de-allocated, and the customer data become unusable, as explained in the previous paragraph.

Secure disposal or re-use of equipment

All items of equipment containing storage media will be checked to ensure that any sensitive data and licensed software has been removed or securely overwritten prior to disposal. Orange has established security policies detailing the sanitization and disposal procedures for handling storage devices.

4.5. Access control

This chapter corresponds to section 9 in ISO 27002 version 2013.

Orange provides access control for the environments that are contractually its responsibility. Access control for systems and applications managed by the customer is the customer's responsibility.

4.5.1. Access control for Orange operators

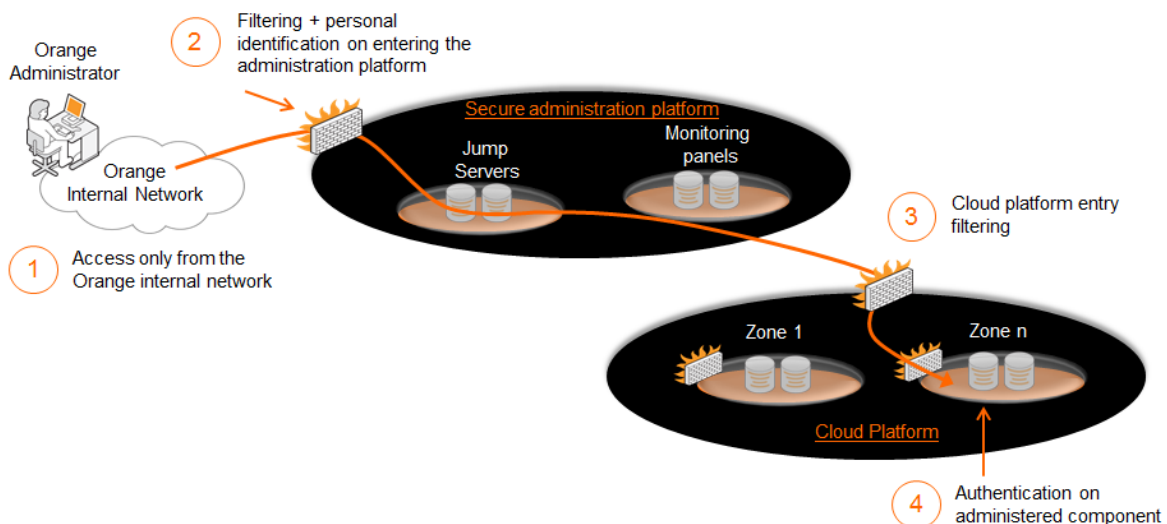
Means of access

Orange operators manage the Cloud infrastructure (servers, storage bays, network equipment, security equipment, etc.) as follows:

1. Access from the Orange internal network. Cloud infrastructure administration is only possible from the Orange internal network.

2. Filtering and personal authentication on entering the administration platform. The Cloud platform is protected from the Orange internal network by an intermediate security zone called a bastion host. This step ensures:
 - Network filtering;
 - Personal strong authentication;
 - Logging of sessions on a centralized log server;
 - Authorization: the administrator can only access resources authorized for their profile.
3. Authentication and authorization on managed cloud environments. The he authorized administrator are given exactly the privileges necessary to their duties. All their actions are logged on a system ensuring non-repudiation.

Administration access is always done securely via the SSL/TLS protocol (workflow examples: HTTPS, SSH).



Password authentication and management

Connections with operating accounts comply with the following password management policy:

- Complexity
 - The password must have at least 12 characters.
 - The password must contain at least one upper case letter, one lower case letter, one numerical digit and one special character;
 - A dictionary attack is done to ensure that the password could not be bruteforced in case of compromise
- Renewal
 - The password must be renewed at least every 3 months;
 - The password must differ from the 5 previous passwords.
- Blocking
 - The account must be automatically blocked after 5 failed authentication attempts;
 - The account must be automatically blocked if the password expiry date has passed.

- Password protection
 - The login screens must be accessible only through encrypted and secure communications;
 - The password entered must not be displayed in clear text;
 - The password must not be stored in clear text (e.g. only its hash must be stored).
- Session duration
 - Inactive sessions must be ended automatically after a period of time determined for each service.
- Logging
 - Authentication-related events must be noted in the log.

Strong authentication is ensured with an additional token either soft or hard.

Procedure for assigning and revoking authorizations for Orange operators

The procedure for assigning/revoking an account and its associated rights, called CARM (Controlling Access Rights Management), is applied and monitored. This procedure is an ISMS security measure (Information Security Management Systems), certified ISO 27001 and deployed by Orange.

Shared accounts are prohibited.

Review of rights for Orange operators

A review of authorizations and rights is carried out every semester. This activity consists in comparing the Human Resource file, describing the roles of each person, with the actual rights assigned in the systems. If non-legitimate rights are discovered, suitable measures are applied:

- Suspension of the account;
- Examination of traces of use to reveal any incidents;
- Corrective measures (updating procedures, having management made aware of the issue, etc.).

These reviews are the subject of a report and are audited yearly by external auditors.

4.5.2. Access control for Orange customers

Customers log in to cloud environments for the purpose of administration or access to application services. This access is via the Internet and/or private network (customer VPN) according to the services and options selected by the customer.

Customer access is always done securely via the SSL/TLS protocol

- Server authentication (administration portals, application servers) is systematically done via an X.509 "server" certificate issued by a recognized certification authority.
- Customer authentication is based on a login/password previously sent securely to each customer. For certain services, strong authentication is possible (use of a software token generating a one-time password, use of the X.509 "customer" certificate).
- Network workflows are systematically encrypted (AES algorithm).

The customer bears responsibility for the security of the authentication details sent to them by Orange. Within the possibilities offered by the tools, Orange sets a minimum complexity for all passwords handled by the customer.

The customer can then manage access accounts himself (creation/deletion, user profiles, administrator profiles). The associated logins and rights are the customer's responsibility.

4.6. Cryptography

This chapter corresponds to section 10 in ISO 27002 version 2013.

Cryptographic sequences - All cryptographic sequences used in Cloud services (e.g.: AES 256, SHA-256, TLS 1.2) are based on market standards with a proven security level. The ENISA10 technical study, "Algorithms, Key Size and Parameters Report, 2013 Recommendations" has been taken into account in choosing cryptographic sequences. Administration workflows are systematically encrypted using the SSL/TLS protocol.

Certificate management – Server authentication certificates are based on certification authorities (e.g.: VeriSign, Thawte, etc.) recognized by the main Web browsers. Authentication certificates for our internal administration services (accessible only by Orange operators) can be based on X.509 certificates generated by an Orange internal certification authority.

Token – Software tokens can be used by Orange Business operators. Certain Cloud services also allow the use of tokens to strengthen customer authentication.

Encryption of hard drives of Orange operators – To ensure the security of data and operations, all workstations (PC, laptops) hard disks are encrypted using McAfee Endpoint Encryption for PC (eePC) solution. Sensitive data can also be encrypted using ZoneCentral encryption tool. ZoneCentral is a security product for encrypting the data on hard drives, emails, and containers with access reserved for authorized and identified users only. The ZoneCentral tool has been certified EAL3+ by ANSSI (ANSSI-CC-2012/07) issued on February 13th, 2012.

4.7. Physical and environmental security

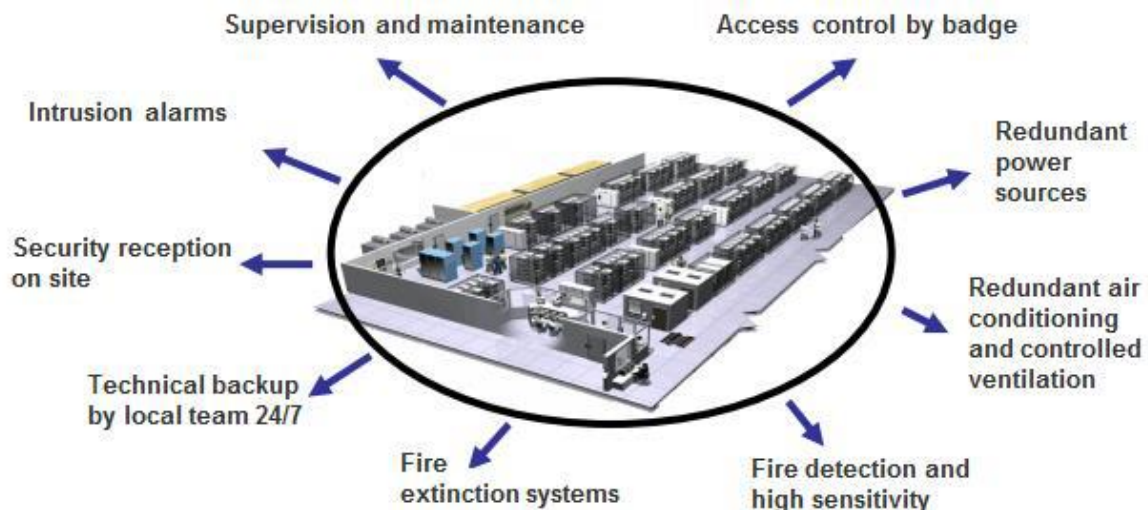
This chapter corresponds to section 11 in ISO 27002 version 2013.

The purpose of physical and environmental security is to protect the information system of Orange and its customers from:

- Environmental threats and harm: explosions, earthquakes, fire, water damage, electrical breakdown, air conditioning breakdown, telecommunication breakdown, etc.
- The threat of physical intrusions: access of non-authorized persons into the room, theft of hard drives containing sensitive information, etc.

Physical and environmental security of our datacenters

Our Cloud services (infrastructures, backups, etc.) are hosted in several datacenters with following security measures:



ISAE 3402 - All datacenters are certified ISAE 3402 Type II (ex-SAS 70). ISAE 3402 is a standard giving customers of externalized services an assurance of the reliability of the internal control mechanism of their service provision. This independent audit report gives Orange Business customers the assurance that the services they benefit from meet the requirements of the Sarbanes-Oxley act in the USA and meet the key objectives of security, change management and service continuity.

ISO 27001 - all datacenters used for Flexible Engine are certified ISO27001. ISO27001 is the most recognized security management standard. It provides customers with the assurance that the physical security of their data in datacenter is well managed.

TIER classification 11 - The datacenters for Flexible Engine are built and operated to be compatible with TIER III requirements (cf. definition in annex A).

Tolerance to breakdowns and maintenance - The technical architectures used enable the following operating constraints to be met:

- No impact on load when the first fault occurs.
- High breakdown tolerance, up to 2 major faults without impact.
- No interruption to service during maintenance operations.

The management modes and the processes used ensure the security of personnel, buildings and the equipment in them.

Physical and environmental security of our operating platforms

Operating platforms are places where Orange operates its Cloud services. The various platforms are:

- MSC (Major Service Centre) in Egypt: level 1 and level 2 support for infrastructures driving Flexible Engine (firewall, servers, networks, etc.);
- MSC in France : Preproduction, level 3 support for Cloud services.

The following operating platforms have ISO 27001 security certification from AFNOR: the Cesson-Sévigné site, and the MSCs in Mauritius, India and Egypt. The functional scope of

the certification is the deployment, supply and support of managed services and communication solutions.

In supplying cloud services, Orange Business may have to transfer customer data outside the European Union to the MSCs mentioned above. Within Orange Business, an “Agreement for the transfer of customers’ personal data” has been in force since January 2nd, 2013. Under the terms of this Agreement, the personal data that customers (as processing managers) transfer to Orange Business (as subcontractor or subsequent subcontractor) benefit from an adequate level of protection as required by the European Commission. Each legal entity of Orange Business outside the European Union in a country where an adequate level of protection is not provided as required by the European Commission, undertakes, in signing this Agreement, to comply with European Directive 95/46/EC, thereby agreeing to be bound by the standard clauses of the European Commission.

4.8. Operational security

This chapter corresponds to section 12 in ISO 27002 version 2013.

4.8.1. Management of operational procedures

ISO 20000 and ITIL certification

The management of services by Orange Business is ISO 20000₁₂ certified. The requirements of ISO 20000 certification are aligned with the best practices in the latest version of ITIL (v3 2011) for processes such as service provision, relations management, problem-solving, testing and production launch, and stringent security requirements.

Change management

Change management is carried out according to the ITILv3 model; this concerns security-related changes in particular. For each service, change management is performed under the responsibility of a change manager.

- “Standard” changes, i.e. pre-identified in a catalog, follow a simplified process and do not require a prior study by a security expert.
- Changes identified as “non-standard” (developments) are formalized in RFCs (Request For Change) and are discussed in a CAB meeting (Change Advisory Board). RFCs are subject to a prior security study by a security expert from the Cloud Security Competence Center, who gives their opinion. Security is thus represented on the CAB; the CAB has the authority to prohibit a change judged to be dangerous or not conforming to security policy, and it can do this quite independently.
- CAB meetings are held regularly (at least once a month). Urgent requests for change (e.g.: critical security update) can be processed immediately at an ECAB meeting (Emergency Change Advisory Board).

The change management implemented by Orange thus ensures:

- Minimal impact of change on operational services and users;
- Standardized methods, procedures and control mechanisms;

- Identification of personnel authorized to request a change;
- The control of change for compliance with security policy;
- Formalization of the correct assessment of the impact, priority, advantages and risks of a change;
- Definition of the change categories and associated implementation times;
- Management of the priorities of changes according to risks and impacts;
- Improvement in the quality of information and communication;
 - Ensuring that all the parties concerned have been involved to limit incidents related to the changes;
 - The configuration database is supplied with correct information;
 - Documentation of all changes;
 - Proactive communication of scheduled outages to users/customers.

Capacity management

Capacity management is carried out in accordance with the ITILv3 model under the responsibility of a “capacity manager”. The aim is to ensure a constant level of service for customers. Capacity management enables customers’ future requirements to be anticipated by analyzing the measurements and trends in the consumption of Cloud infrastructure resources.

Capacity management oversees and acts mainly on:

- CPU/RAM resources;
- Network resources (bandwidth, address space);
- Storage and backup resources;
- Software licenses.

For each service, the “capacity manager” regularly compiles statistics with various indicators specific to each service (technical measurements, business forecasts).

4.8.2. Protection against harmful codes

Customer environment

Certain Cloud services natively include a dedicated antivirus for customer environments. This service is, save for some specific customers, confined to Microsoft Windows systems since the risk / benefit ratio has not been deemed favorable for Linux systems.

The antivirus agent installed on Windows systems regularly updates its virus signature database (at least once a day) from a specific updating server located in a security zone dedicated to customer machine updates. This virus signature database server is itself updated in real time on public servers available on the Internet from the antivirus publisher.

Orange infrastructure

All Microsoft Windows systems have an antivirus; this concerns both the servers deployed within the infrastructures and administration machines. Except for some exposed servers, Linux machines do not have an antivirus.

The antivirus agent installed on Windows systems regularly updates its virus signature database (at least once a day) from a specific internal Orange updating server located in a

security zone totally isolated from customer environments. This virus signature database server is updated in real time on public servers available on the Internet from the antivirus publisher.

Orange internal security zones are also protected by the use of Intrusion Detection System (IDS) probes operated by a Security Operating Center (SOC).

4.8.3. Backup management

Customer environment

The data in customer environments are backed up in accordance with a backup policy specific to each service and according to the customer's options. In general, customers' data are backed up automatically and periodically with a retention period depending on the services

The backups can be of different kinds according to the services and options. Example:

- File level backup, particularly for application data;
- Image level backup for virtual machines;
- Raw data backup (file server).

Recovery methods vary according to offerings:

- Recovery independently by the customer;
- Recovery by Orange in response to a request for change via the administration portal.

Orange infrastructure

All the configurations of cloud infrastructure equipment are regularly backed up. This concerns in particular:

- Servers: Windows, Linux, application data (databases in particular);
- Network equipment (routers, switches);
- Security equipment (UTM firewall, VPN appliance);
- Storage equipment.

Backups for static configurations are carried out on each configuration change. The backup operation is controlled by the change management process.

All configuration backups are stored on third party sites (distinct datacenters) to guarantee the availability of configuration data in case of a major incident on the production site. Configuration backups for new datacenters are carried out on the same production site but in different computer rooms. Backups are stored securely (virtual partitioning) on a need-to-know basis.

4.8.4. Security logging and supervision

Orange systematically oversees its Cloud infrastructures and services, particularly as regards security. Events are logged for all the components, including security components. The purposes of logging are:

- To detect and analyze security incidents;

- To meet legal obligations;
- Troubleshooting.

All components are supervised 24/7 and standby duties are arranged in case of critical incidents, especially when security is affected.

Nature of security logs

Events that are logged are defined component by component according to two criteria:

- legal requirements: In France, the LCEN13 and CPCE14.
- security requirements:
 - events related to administration for all the components,
 - events affecting security components (e.g. Firewall logs).

In Europe, logging and log supervision are carried out in compliance with the General Data Protection Regulation (GDPR).

Centralization of security logs

All the cloud infrastructure components feed back their logs periodically or in real time to collection servers located in specific security zones. The integrity of security logs is protected and the logs are only accessible by administrators in charge of security.

Pro-active monitoring of security logs – The security log collection servers enable human supervision of security events in accordance with processes defined during the design phase. For example, failure to feed back logs, a sudden increase in the number of events or a particular event occurring may be subject to analysis and processing. This supervision can be offered to customers as an option.

Single time reference

Timestamping of Logs, including security logs, is done with reference to a single time for all components in the cloud infrastructures. The time service is provided by various NTP servers in specific security zones. The time reference is supplied by a dedicated physical server (root server with GPS antenna) installed in Orange Business' internal service zones.

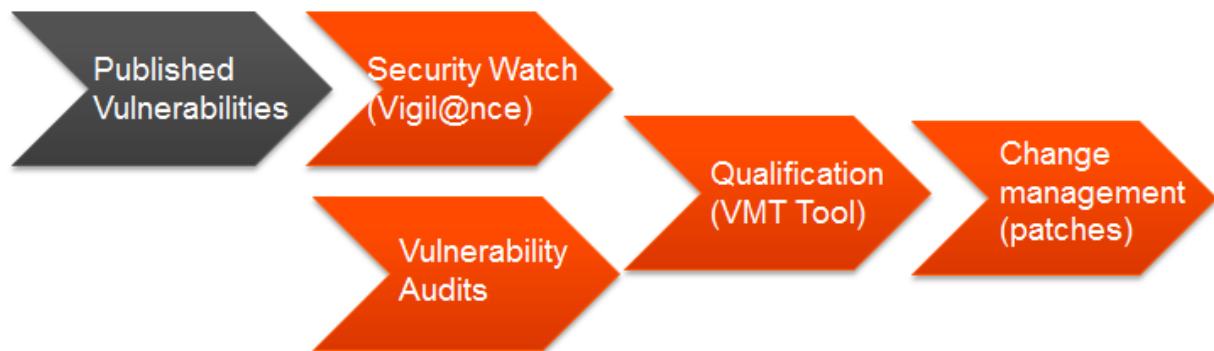
Customer consultation of logs

With certain services, the customer can access the Logs for their environment, including application and firewall logs. Access to logs is secure and restricted to each customer's environment

4.8.5. Management of technical vulnerabilities

The main purpose of vulnerability management is to keep the Cloud infrastructures and services secure. Vulnerabilities are identified during a security watch and vulnerability scans. They are then qualified with a patch to be deployed in production.

The security watch, monitoring of patch deployment and reporting are the responsibility of "Security Operation Managers".



Security watch

All components are subject to a permanent security watch aimed at identifying technical vulnerabilities that may affect the security of infrastructures and Cloud services.

The security watch is supplied with information from several channels:

- Security bulletins from product publishers (virtualization hypervisor, network and security equipment, OS, etc.);
- Watch reports: Vigil@nce, CVE, NIST NVD, CERT-FR etc.;
- Customer feedback;
- Feedback from our vulnerability audits and scans.

Vigil@nce is the security watch tool of the Orange group. It also provides an inventory of the various products (hardware and software versions) used within each service, based on the CMDB. Thus, the “Operations Security Managers” for the various services are alerted as soon as a vulnerability is likely to affect the services for which they are responsible.

Vulnerability audits

Regular audits (or permanent ones on the most critical infrastructures) are launched on our production platforms both to measure the state of vulnerability of our platforms and to potentially identify vulnerabilities that have not been analyzed by the watch. To do this, Orange uses solutions from third party leaders in the vulnerability scan market (Qualys, Nessus).

Vulnerability qualification

Security vulnerabilities and patches are qualified with the Orange internal “Vulnerability Management Tool (VMT)” under the responsibility of the “Operations Security Manager”. The Operations Security Manager decides the level of criticality of each vulnerability and the action plan to be followed (no action, change of configuration, installation of a security patch). They can optionally co-opt engineering teams or other security managers (e.g.: Engineering Security Manager, Chief security officer) to decide on the processing of a vulnerability.

Application of security patches

All qualified security patches (change in configuration, security patch) are deployed in the production launch process (change management):

- Vulnerabilities with low/medium criticality are managed as minor or major production launches (CAB). The scheduling is adapted to Orange Business' operational constraints.
- Vulnerabilities with high/critical criticality are managed as urgent production launches (ECAB).

Equipment inventory files are updated after applying patches.

Reporting

The Operations Security Manager prepares an internal monthly report on the security watch and vulnerabilities processed over the period. This report is produced using the Orange VMT tool and for confidentiality reasons is not sent to customers.

In the case of a critical vulnerability likely to seriously affect its customers, Orange Business may communicate directly with its customers to inform them.

4.8.6. Security of administrative workstations within Orange

All Orange operator workstations used to administrate Cloud infrastructures and services are based on generic and secure configuration models:

- The configuration of workstations (operating system, middleware and applications, user rights) is stringent from a security point of view. (master "E-buro" on OS Microsoft Windows base);
- There is systematically an antivirus which is updated automatically and regularly;
- Policy of system updates with Orange patches validated internally before circulation.

All workstations hard disks are encrypted using McAfee Endpoint Encryption for PC (eePC) solution. Sensitive data can also be encrypted using ZoneCentral encryption tool. ZoneCentral is a security product for encrypting the data on hard drives, emails, and containers with access reserved for authorized and identified users only. The ZoneCentral tool has been certified EAL3+ by ANSSI (ANSSI-CC-2012/07) issued on February 13th, 2012.

All work to be done or executed are performed via jump servers that are under stricter and more secured environments. These jump boxes are then filtered and monitored by token based access platforms such as CyberArk that trace each activity and can be replayed for auditing.

4.8.7. Secure administration portal

An administration portal is available for each customer for managing the various services and users; it is accessible from the Internet and/or the Orange Business customer VPN. The administration portal is accessible in HTTPS (TLS) via a login/password. The administrators' user IDs (login/password) are previously sent by Orange Business to the customer following a secure procedure. Server authentication is performed with an X.509 certificate.

All these actions performed on the portal are logged.

4.9. Communications security

This chapter corresponds to section 13 in ISO 27002 version 2013.

4.9.1. Orange cloud architecture security

General principles

The architectures of the various Cloud services or templates of architectures for Flexible Engine are validated by an engineering security manager under the governance of the Chief Security, Risk and Compliance Officer. The latter guarantees the coherence of the various architectures from a security point of view.

This approach based on a predefined architecture model offers several advantages:

- Simplifies deployment of new Cloud services;
- Makes practices and architectures homogeneous across the various services;
- Enables sharing of some equipment;
- Guarantees the security of services: the models are validated then their deployment is overseen by the Cloud Security Competence Center. The security level of each new service is also assessed by audits and intrusion tests before production launch. This also applies to all major changes. The results of internal audits and intrusion tests are not sent to customers for confidentiality reasons.
- Makes it easier to maintain secure conditions.

For each new service, the security validation of the architecture is completed with a systematic risk analysis. These 2 actions form the basis of the "SecurityInTTM" approach to guarantee that security is taken into account in the design of services.

Trusted domains and Security zones

The architectures of the various Cloud services all apply the same model consisting of separating trusted domains, including the Back-end (Orange internal) and Front End supporting the Cloud services seen by customers. Back-end / Front End partitioning is physical, i.e. there are servers dedicated to the back-end and others dedicated to the front-end.

Within each trusted domain, security zones provide virtual partitioning by using functionalities such as:

- Virtualization (virtual machines, virtual firewalls, virtual load balancers, virtual routers/VRF);
- VLANs (802.1Q);
- Virtual networks VPN (IPSec, SSL);
- Virtual storage (virtual drives).

Virtual partitioning guarantees the isolation of the various customer environments.

Communications (network flows) between the various security zones are systematically controlled by firewalls (stateful type filtering). The local configuration of the various components is also designed to strengthen partitioning and security (e.g.: ACL in the routers).

Orange infrastructure

Infrastructure servers are clustered into security zones according to their function (administration server, front-end server used by customers), their nature (database,

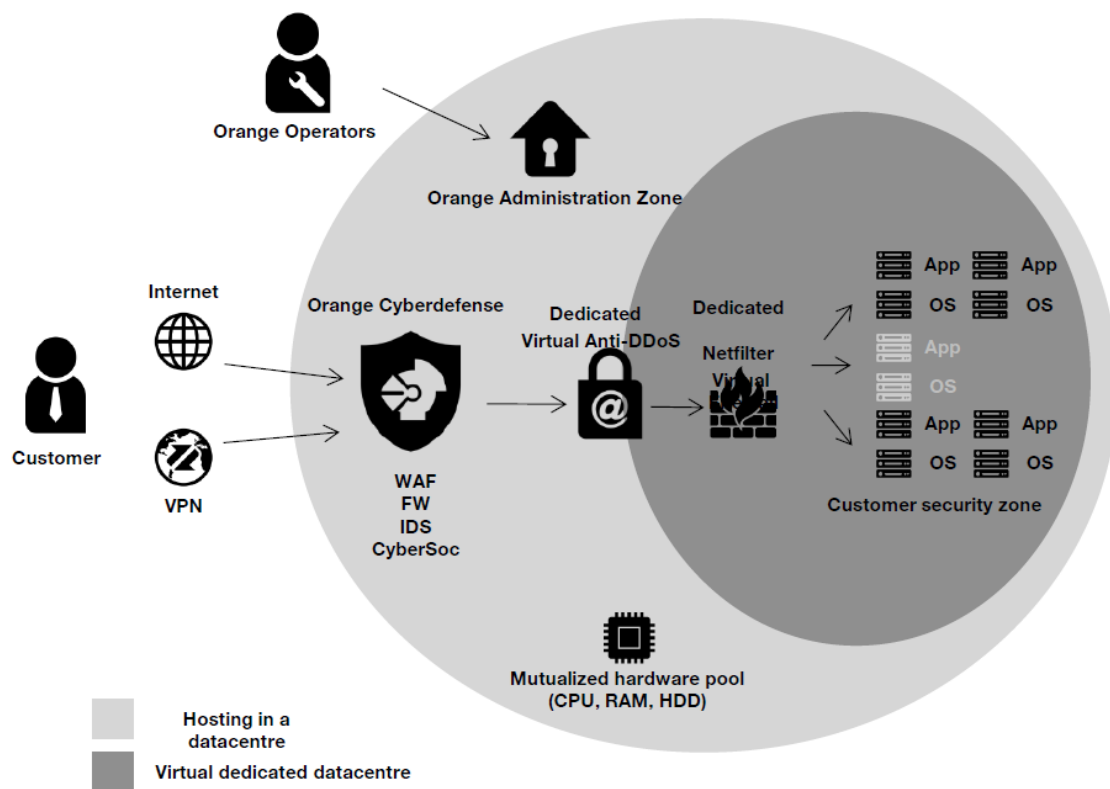
webservers) and their level of exposure (for example, whether or not they are accessible by customers). Thus, security zones are dedicated to the Orange operating tools.

In general, all the servers have an interface dedicated to data workflows and a technical interface dedicated to service workflows (including administration).

All cloud infrastructure components (servers, firewall, routers, storage bays, etc.) are configured based on secure configuration guidelines established by the various engineering departments with the assistance of Orange group security experts.

Customer environments

Isolation between customer environments is based on the virtualization functions described above. Within an environment, the customer can manage security functions (including virtual firewalls). The functions supplied to customers vary according to the chosen services and options.



4.9.2. Security of exchanges

Security of Orange internal workflows

Orange internal exchanges are rendered secure by the following means:

- Networks: The interconnection networks between Cloud platforms and operating platforms are protected by one or more of the following solutions:
 - VPN MPLS;
 - Encrypted tunnels (IPSec, TLSv1.2) ;
 - Dedicated fiber network.
- Messaging: At the customers' request, sensitive data exchanged by email between customers and Orange Business can be encrypted by a third party tool agreed with

the customer (e.g.: Zed containers). Orange Business leverage PKI based email security across all departments and locations using smart cards or USB dongles.

- Administration flows: Administration workflows are given state of the art protection in accordance with SSL-type connections (SSH, HTTPS).

Exchanges between the various security zones are systematically controlled by firewalls applying the elementary principle, "anything not explicitly authorized is forbidden". Filtering is stateful (source/destination/protocol control) and flows are logged.


Customer flow security

Customers log in to cloud environments for the purpose of administration or access to services. Customers' administration flows are systematically rendered secure by protocols guaranteeing their authentication, confidentiality and integrity (TLSv1.2, AES256, etc.). Access methods vary according to the services and options chosen by the customer.

The security of service workflows depends on the service in question, but in general, exchanges are secured with SSL/TLS type connections.

4.9.3. Protection against denial of service and intrusions

Orange Business has means of protecting its customers against denial of service attacks and network intrusion attacks in Flexible Engine. To simplify the complex defense-in-depth security design and implementation of security solutions, we have defined security zones and holistic network segregation strategy to minimize the impact of attacks based on business functions and security risks.



- DMZ zone offers the first line of detection and protection mechanisms to host public facing services from the Internet or the customer's VPN access. This includes infrastructure components (for example load balancer, proxy) and application

components (for example API gateway, service portal). By directing untrusted traffic from external network to dedicated DMZ zone, it can be treated and filtered with the combination of Anti-DDoS, web application firewall (WAF), next-generation firewall (NGFW) and intrusion detection system (IDS) thus ensuring only trusted and necessary access to other zones is permitted.

- Management zone contains O&M components of Flexible Engine where authorised O&M personnel can access via mandatory virtual private network (VPN) to designated systems. From this zone, they can perform administrative tasks through secure jump hosts with privileged access management to managed nodes.
- Point of Delivery (POD) zone represents the resource pool that comprises of managed compute, storage, database and network clusters. These resources are further protected and isolated through relevant security measures as described in chapter 5 to ensure one customer cannot access another's resources.
- Net(work) Service Zone hosts firewalls to provide network address translation (NAT), bandwidth management, VPN gateway functions. This zone will also be responsible for tenant's ELB, EIP and virtual firewall management.
- Universal Distributed Storage (UDS) zone provides the object storage service for tenants which can be accessed from Internet with S3 (Simple Storage Service) interface. This zone is protected by port filtering, NAT and load balancing firewalls.

Logs collected from security equipments will be sent to a SIEM for real-time analysis and remediation by OBS's CyberSOC.

In addition to advanced perimeter protection measures of Flexible Engine, the tenants can also set fine-grained policies on Anti-DDoS for their EIPs, security groups in VPC and access controls via IAM to strengthen overall protection against external threat actors and malicious insiders.

4.10. Acquisition, development and maintenance of information systems

This chapter corresponds to section 14 in ISO 27002 version 2013.

4.10.1. Risk analysis

All changes (acquisitions, developments, etc.) are subject to security risk analysis carried out by the Orange Business Cloud Competence Center. Risk analyses enable security measures to be put into place to limit the risks identified. For confidentiality reasons these documents are not sent to customers.

4.10.2. Good development and integration practices

Good development practices

The specific developments carried out by Orange or by a subcontractor comply with the internal Orange guide on best secure development practices (comprising 38 measures).

Good integration practices / strengthening of configuration security

Systems and software under Orange's responsibility are configured with a high level of security by applying a "hardening guide". For example, for OpenStack based environments

like Flexible Engine, Orange relies on the security guides established by OpenStack (<https://docs.openstack.org/security-guide/index.html>).

Security certification of critical components

The Orange Cloud Security Competence Center makes sure that critical components have recognized security certification within an appropriate scope before being integrated into our Cloud services. Thus, the Cloud Security Competence Center generally relies on "Common Criteria" certifications (ISO 15408) and verifies that they are relevant by checking the following parameters:

- Level of assurance: EAL 3 as a minimum
- Security target covering sufficiently wide functional domains: cryptographic support, identification and authentication, imperviousness of virtual machines, etc.

As an example, the following critical components are common criteria-certified:

- VMWare Vsphere 5.0: EAL4.
- Juniper firewall: EAL4.
- Huawei FusionSphere V5: EAL3+.
- Huawei NGFW V1 (Firewall): EAL4+.
- Huawei FusionServer Management Software: EAL3+.

4.10.3. Security acceptance

Security validation of components

The most critical components on the Cloud platform are subject to security tests (configuration/code review and/or intrusion tests) to validate their security level.

As an example, the administration portals have already been subjected to several intrusion tests and a configuration/code review).

All the templates (OS or applications) supplied by Orange are also subject to validation and are regularly tested and updated to take account of the latest vulnerabilities.

Non-production platform

Development/qualification/pre-production platforms are available, in particular to validate the security of changes or to conduct complete intrusion audits.

Absence of customer data on non-production platforms

There are no customer data on non-production platforms, which are completely separate from production platforms.

4.11. Subcontractor management

This chapter corresponds to section 15 in ISO 27002 version 2013.

4.11.1. Security in contracts with our subcontractors

In our contracts, our service providers/suppliers give an undertaking on the confidentiality and integrity of data to which they have access during the service.

Subcontractors integrated into internal Orange teams use the same tools and processes as Orange staff and therefore, by default, comply with best practice as set out in this document: awareness, physical and virtual access control, etc.

4.11.2. Security monitoring of services provided by our subcontractors

As part of sourcing process, Orange subcontractors are due to describe their security measures in a security assurance plan reviewed annually. Orange may also audit its subcontractors to verify that they are complying with the contractual security undertakings. These audits result in security action plans for improving the security level of subcontractors.

4.12. Management of security incidents

This chapter corresponds to section 16 in ISO 27002 version 2013.

The policy on managing security incidents is described in chapter 5 and annex A of "OCB security policy". It is set out in chapter "15.1 – Security Incident Management" in the incident management process. The present chapter explains the principles described in these documents.

Preparation of operational teams

The purpose of this phase is to prepare each person for processing incidents: operational excellence (awareness, regular training in managing various types of incident, etc.) and quick access to up-to-date documentation (asset inventories, network diagrams, technical documentation, logs, etc.).

Detection of security incidents

The means of detection deployed for detecting a security incident are:

- Supervision tools (Shinken) ;
- SIEM⁹ (netForensics type) in charge of supervising logs (IDS, log FW, systems log, application logs, etc.);
- Watch cell;
- Security team and other personnel;
- Alerts from customers;
- Complaints from other providers via Abuse Desk.

This supervision is performed on components that are the direct responsibility of Orange (portals, hypervisors, customer environments with security managed by Orange, etc.).

Recording and qualifying security incidents

Orange Business have a security incident management tool.

Two categories of staff enter alerts into the tool in the form of an incident ticket:

- Technical staff in charge of service operation / supervision;

⁹ SIEM: Security Information Event Management, a tool for correlating logs, i.e. relating different events to a single cause.

- Support center staff, alerted by a customer's call or email.

The alerts are qualified according to their nature and seriousness.

Nature of the incident - The incident management tool distinguishes four categories of incident:

- Intrusion;
- Malfunction;
- Vulnerability;
- Legal (regulatory type incident involving particular actors).

Seriousness of the incident - The table below summarizes the levels of seriousness of incidents and the actions to be carried out:

Level	Description
High	Corrective action immediately
Medium	Action can be delayed, but a security maintenance operation must be scheduled now
Low	Action can be delayed until the next scheduled maintenance operation

The incident ticket is directed to the appropriate team including the Security Department.

Response to security incidents

The following responses may be given:

- Emergency measures (quarantine, etc.);
- Crisis cell activation;
- Communication with customers, partners, operators, etc.
- Patch application;
- System recovery;
- etc.

Post-incident review and actions

Once the security incident has been dealt with, the Cloud Security Competence Center analyzes the nature of the incident and the quality of Orange's response. If necessary, the Cloud Security Competence Center updates the procedures for managing security incidents as part of the continuous improvement approach.

4.13. Security in the management of business continuity

This chapter corresponds to section 17 in ISO 27002 version 2013.

The availability rates of each service are given in the service descriptions.

All our Cloud services have complete infrastructure redundancy over several sites (or, as a temporary exception, over several computer rooms), with high availability mechanisms so that local failures are made transparent: network (Internet access, VPN access, firewall), administration portal, virtualization, storage, etc.

Backups (configurations, customer backups) are in priority replicated on remote sites to guarantee the availability of data in case of a major incident on the production site. Depending on the service, remote customer backups may only be proposed as an option.

Cloud Orange services are operated from several operating platforms, as explained in part 4.7. Thus, if an operation site becomes totally unavailable, administration continuity is ensured for the Cloud services.

Operational continuity is regularly tested through simulations.

There are rollback clauses in the Cloud services contracts. If the Service is terminated, unless the termination is through a fault of the Customer, the Customer can ask Orange Business to trigger rollback. Orange Business will then deploy the means reasonably necessary to ensure continuity of the Service, so that at the end of the Rollback Period, the customer is afforded the capacities to continue to satisfy their requirements (cf. contract for more details).

4.14. Conformity

This chapter corresponds to section 18 in ISO 27002 version 2013.

4.14.1. Compliance with legal and contractual requirements

As part of the risk analysis attached to all Orange projects, a review of the legal and contractual obligations is conducted and an action plan is proposed. The resulting deliverable is called the LOA (Legal Obligation Assessment). The points covered are:

- compliance with contractual clauses (licenses, industrial property, specific undertakings in the description of services, etc.).
- compliance with specific regulations arising from Orange's field of activity (LCEN¹⁰ and CPCE¹¹);
- keeping a register of the processing of data of a personal nature, in conformity with the law of 6 January 1978 and General Data Protection Regulation (GDPR). Orange Business has nominated a Data Protection Officer (DPO) who keeps the register. The DPO is the contact¹⁸ specialized in the protection of personal data, both for the manager processing these data and in the relations between the latter and the CNIL (French National Commission for Data Protection and Freedom), the independent administrative authority charged with overseeing compliance with the freedom of information act. The DPO thus has a central role in the secure development of new information and communication technologies, and within the company disseminates the culture of freedom of information and the control of risks to personal data.


4.14.2. Monitoring compliance with security policies

The Cloud Security Competence Center is in charge of monitoring application of the Security Policy for each Cloud service. Thus it performs continuous monitoring and runs organizational and technical audits at least every 6 months on a revolving basis.

- Example of organizational audit: control of rights management
- Examples of technical audits: intrusion tests, Nessus or Qualys scans, configuration audit, etc.

¹⁰ LCEN: Law on Confidence in the Digital Economy

¹¹ CPCE: Post and Electronic Communications Code



Example of ISO 27001 maturity report (values given as examples)

In addition to this, Orange Business undergoes regular audits (AFNOR, ISAE 3402) conducted by external bodies and leading to security level checks.

4.14.3. Security-related certifications

Orange19 has the following certifications pertaining to the scope of Cloud services:

- ISAE 3402 type II
 - All datacenters are ISAE 3402 Type II certified (ex-SAS70)
- ISO 27001 security certification
 - Orange Business has ISO 27001 security certification within the scope of the "deployment, supply and support of service management and communications solutions" for the sites in Cesson-Sévigné, Egypt, Mauritius and India20.
- ISO 27017 security certification
 - Orange Business is one of the first French companies to be ISO 27017 certified, taking into account the entirety of the security recommendations specifically dedicated to cloud security (the highest security level).
- ISO 27018 security certification
 - Orange Business is one of the first French companies to be ISO 27018 certified, taking into account the entirety of the security recommendations specifically dedicated to the protection of personal data (the highest security level).
- Multi-Tier Cloud Computing Security (MTCS SS584:2015) Level 3: Aims to encourage the adoption of sound risk management and security practices for cloud computing, by providing relevant cloud computing security practices and controls for cloud users, auditors and certifiers to understand cloud security requirements,

and for public cloud service providers to strengthen and demonstrate the cloud security controls in place, in their cloud environments.

- Common criteria certification EAL 2+ (ISO 15408) for the security of the international IP-VPN network in 2008
 - The certification report is available here:
https://www.ssi.gouv.fr/uploads/IMG/certificat/dcssi_2008-20fr.pdf
 - The certification target is available here:
https://www.ssi.gouv.fr/uploads/IMG/certificat/dcssi-cible_2008-20fr.pdf
- Common criteria certification (ISO 15408) or equivalent independent certification of security equipment and virtualization equipment used in our services:
 - Huawei FusionSphere: EAL3+
 - Huawei NG Firewall: EAL4+
 - FortiGate Firewall: EAL4+
 - McAfee IPS: FIPS 140-2, DoD UC-APL
 - F5 Big-IP WAF: ICSA Labs

5. Specific security measures

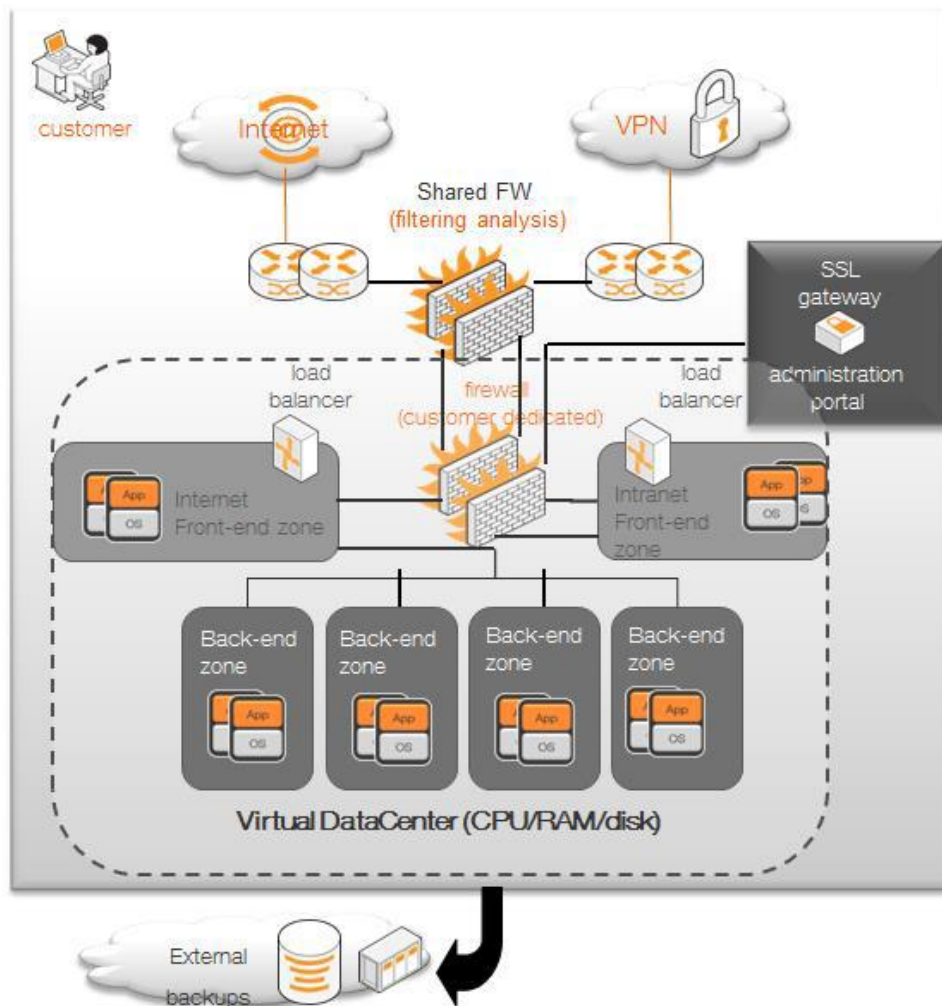
The purpose of this chapter is to present the security measures specific to each service. They are complementary to the common security measures presented in chapter 4.

5.1. Compute services

Orange Flexible Engine provides wide range of cloud-based computing services with varying instance configurations that can scale up and down automatically depending on customer business requirements.

The following security measures are deployed in this service:

- Virtual firewall security zones;
- Virtual Load Balancing;
- Strong authentication by software authenticator.



5.1.1. Elastic Cloud Server

An Elastic Cloud Server (ECS) is a computing server that consists of CPU, memory, image, and Elastic Volume Service (EVS) disks and allows on-demand allocation and elastic scaling.

ECS integrates virtual private cloud (VPC), security group, and multi-data-copy capabilities to build up an efficient, secure, and reliable computing environment for users to ensure stable, continuous running of their services.

ECS is a self-service. Tenant only needs to specify the amount of CPU/Memory and template. ECS instance will be generated in a short time.

ECS Security Architecture

ECS uses resource isolation, network isolation, security group rules, anti-DDoS, and brute-force attack prevention to provide a secure environment. ECS supports 99.95% availability and 99.99995% data durability.

Image hardening

Flexible Engine's professional security team performs security hardening on public images and patches any system vulnerabilities that may occur. Secure, updated public images are created with the help of an image factory and provided to users through Image

Management Service (IMS). Pertinent hardening and patch information is also provided to tenants for reference during image testing, troubleshooting, and other O&M activities. When creating VMs, tenants can decide based on their applications and security policies whether to use an up-to-date public image or create a private image that has the required security patches installed.

Security zones and Virtual firewall

Within their Virtual Datacenter environment, each customer can create several security zones (8 maximum) and control the workflows between these zones via a dedicated firewall. The security zones can also be of the Front-End or Back-End type depending on the role of the servers.

We support security group in public cloud service. Security group functions to provide ACL for different VMs belong to one tenant. It supports five-tuple ACL rule, which includes source IP, source port, target IP, target port, and protocol. This provides tenants the ability to specify the type of traffic and direction (ingress/egress) that is allowed to pass through a port.

When a port is created in an OpenStack network, it is associated with a security group. If a security group is not specified, the port will be associated with a default security group. This group will drop all ingress traffic and allow all egress by default. Rules can be added to this security group to change the behavior. A security group supports the five-tuple rule, which includes source IP, source port, target IP, target port, and protocol.

Security group features:

- Each VM has an independent access control list (ACL) that can be automatically refreshed after VM migration.
- Security group policies can further isolate VMs in the same VLAN, thereby reducing the number of VLANs needed in the system for VM isolation.
- VMs are controlled by distributed security policies so that packets do not need to be sent to a centralized policy control node for processing, avoiding system performance deterioration.
- Security groups (horizontal protection) and boundary firewalls (vertical protection) jointly enhance system security protection capabilities.

The security zones and virtual firewall (security rules) are managed via the administration portal by the customer. Virtual firewall configuration is therefore the customer's responsibility.

Virtual firewall logs are available to the customer via the administration portal.

Virtual Load Balancing

The customer can deploy a "Load Balancing" function within each security zone to increase the availability of their services. Thus they have a virtual instance dedicated to "Load Balancing" whose configuration is managed via the administration portal.

Traffic management functionalities include:

- intelligent load balancing (choice of 2 algorithms: "Round Robin" and "Least Connection");
- source IP persistence.

Strong authentication by software authenticator

By default, customers' remote administration access (portal, virtual machines) is secure with standard authentication (login/password) on an SSL gateway. Optionally, the customer can receive a strong authentication service through the use of a software token specific to each user.

Network isolation

We provide VPC service to tenants. Virtual Private Cloud (VPC) is an isolated section in the public cloud where a tenant can launch resources in a virtual network defined by the tenant. Each VPC is in an isolated Virtual Extensible LAN (VXLAN), which is allocated by cloud platform and is not visible to tenants. So tenant network is isolated in L2 network.

VXLAN solution provides the Virtual Tunnel Endpoint (VTEP) function to original switches to make them compatible with both VXLANs and VLANs. This solution uses VXLANs to provide virtual networks and still use VLAN mode for virtual switch management, maintenance, and monitoring, thereby simplifying user operations.

There is a virtual switch in each compute node; there are VLANs connection between VMs and the virtual switch in the same host. While VM traffic pass through the virtual switch, there is a VTEP function will transfer VLAN to VXLAN by encapsulation VXLAN into VLAN packet. The network connection will be VXLANs to other physical machines or network nodes.

IP/MAC address spoofing protection

To avoid network issues that may occur if users change their IP or MAC addresses at will, IP and MAC addresses are bound together using DHCP snooping. Spoofing is further prevented by using IP Source Guard and dynamic ARP inspection (DAI) to filter out packets from unbound addresses.

Remote access control

Tenants can log in to their VMs over SSH to perform system maintenance. However, leaving the SSH port open is a relatively high security risk. For security purposes, tenants can enable access authentication by username/password or crypto key (public and private key pair). It is recommended that crypto key-based authentication be selected.

Resource management

Tenants can manage ECS computing resources through API. API access requests must be authenticated and authorized through IAM before resources can be managed.

5.1.2. Auto-Scaling

Auto Scaling (AS) uses preset AS policies to automatically scale service resources up and down based on service requirements. AS ensures that resource usage satisfies current service requirements without any manual intervention, scaling application systems out as service utilization grows with tenant business and scaling in as it declines.

Customers can configure scheduled and periodic scaling tasks, monitoring policies, and AS group capacity thresholds to enable AS to automatically increase or decrease the number of ECS instances. This helps reduce resource and labor costs and ensures that services operate in a stable and sound manner. By automating the allocation of computing resources and the enforcement of management and control policies, AS helps avoid the impact of resource exhaustion-type attacks as well as security risks resulting from human errors during manual allocation of resources.

Auto Scaling Service will invoke ECS API to create and destroy resources. The API requests are authenticated and only authorized users can access and manage Auto Scaling. This is tightly integrated with Identity and Access Management to enforce role based access control (RBAC).

AS can automatically add managed instances to an ELB listener, which delivers access traffic to all instances in a scaling group. This offers a higher level of protection against DDoS attacks than the traditional method of directly accessing a single backend server and service. AS can detect instance status in real time and launch new instances to replace those that are not operating properly. AS can also deploy the instances in a scaling group evenly across multiple availability zones (AZs) to improve system availability and support disaster recovery of applications deployed in the scaling group.

5.1.3. Image Management Service

An image is used to create ECSs and consists of a common operating system (OS), preinstalled public applications, and the user's private applications. Image Management Service (IMS) transforms the image services previously provided by IT administrators into custom services that allow for self-service. IMS allows the User to create, edit, upload, and delete images.

Flexible Engine classifies images into public, private, and shared:

- Public images are standard operating system images provided by Flexible Engine.
- Private images are created by users for their own use.
- Shared images are custom images created by any user, maintained on a voluntary basis by the user community and provided for all users to use.

Image Management Service (IMS) provides simple and convenient self-service management functions for images. Tenants can manage their images through the IMS API or the management console. Flexible Engine staff update and maintain public images, which includes performing security hardening and applying security patches on them as required. The staff also provide security-related information for users to reference in deployment testing, troubleshooting, and other O&M activities. Users can deploy their ECS servers by selecting one of the public images provided, creating a private image from an existing cloud server deployment or an external image file, or using a shared image and participating in its development and maintenance.

Considering that attacks on the IMS API could have severe consequences, such as the disclosure of many tenants' data or the interruption of management services, IMS offers a wide range of security measures to protect the IMS management system from such attacks.

Tenants must be authenticated with IAM and receive a token before they can use IMS. The service uses a multi-tenancy-based permissions management model and secure communications protocols, strictly verifies parameters, and provides measures for protecting sensitive information and auditing logs.

IMS supports encryption and integrity verification for the transmission and storage of images. All data is stored in an image database on a trusted subnet, and public and private images are stored in different buckets using Object-based Storage (OBS). IMS comes with secure cryptographic algorithms and functions that enable users to encrypt their image files and sensitive data in both transmission and storage.

IMS requires tenants to have sufficient permissions to perform any operation, and keeps audit logs of major operations. Audit logs are retained indefinitely so that tenants can accurately trace operations performed over a long period of time.

5.1.4. Cloud Container Engine

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

The CCE service needs to work with Elastic Cloud Server (ECS), Virtual Private Cloud (VPC), Object Storage Service (OBS), Identity and Access Management (IAM), Elastic Load Balance (ELB), and Elastic Volume Service (EVS).

CCE leverages several open source software components to provide features which have established security measures in place (CIS¹² Benchmarks, NIST). The major open source components contain:

- Kubernetes (k8s): k8s is the container management system provided by Cloud Native Computing Foundation. K8s is a distributed system with master-slave mode. The major components are k8s master and k8s slave. See more details of k8s at <https://kubernetes.io/>
- Docker Engine and Registry: Docker engine is the software that provides support of running Docker container in an operation system. And Docker registry is the repository of saving Docker images. See more details of Docker at <https://www.docker.com/>
- Cloudify: Cloudify is an engine of orchestrating and deploying applications. It provides a graphical interface that can let user to define complicated topology with drag and drop style. See more details of Cloudify at <https://cloudify.co/>
- Kafka: Kafka is a high throughput distributed messaging system. CCE uses it to send messages between components. See more details of Kafka at <https://kafka.apache.org/>
- Etcd: Etcd is a distributed key-value store that ensures reliability. CCE uses it to synchronize management and configuration data between components. See more details of etcd at <https://coreos.com/etcd/>

¹² Center for Internet Security

- ELK: ELK is a solution for log management. CCE leverage ELK to collect and save logs. ELK is composed of Elasticsearch, Logstash and Kibana. See more details of ELK at <https://www.elastic.co/>
- HAProxy: HAProxy is high performance load balancer software. CCE uses HAProxy to manage network access traffic. It's working with active-active mode. See more details of HAProxy at <https://www.haproxy.org/>
- Nginx: Nginx is a reverse proxy server. CCE uses Nginx to translate external traffic. See more details of Nginx at <https://nginx.org/>
- SaltStack: SaltStack is an engine to perform automated deployment and configuration. CCE leverages SaltStack to perform deployment, configuration and upgrade of the service. See more details of SaltStack at <https://saltstack.com/>

In addition to the open source components, CCE has some components that perform integration and communication functions. The major components include:

- PaaS Console: the component provides all the user interaction features. It's integrated with OCB console.
- API Servers: that provides API of each feature, such as deployment, template management, and operations management.
- PaaS IAM: the component provides centralized user/authentication/authorization management function for CCE. It integrates with the IAM service of OCB.
- Cluster Management: the component calls OCB IaaS APIs to perform lifecycle management of clusters, which is composed of ECS instances. As CCE calls ECS API as a regular user, the ECS instances created by CCE can be seen by user from ECS portal.

5.1.5. Dedicated Cloud

Dedicated Cloud (DEC) provides an isolated virtual computing resource pool in the Public Cloud, in which a Tenant can apply for dedicated physical servers to build their own computing resources pool and use highly reliable distributed storage resources and isolated networks.

The ECSs created by the DEC Tenant are deployed in the dedicated virtual resource pool. The DEC resources, like a private cloud, can be managed and monitored through a management console.

Besides physical servers are dedicated and isolated physically, the network in DEC is provided by VPC service, which is isolated from other VPCs and other End Users. DEC service is very suitable for scenarios which customers have very critical requirements about security isolation, high performance and high reliability, such as financial and government organizations etc.

5.1.6. Bare Metal Server

Bare Metal Server (BMS) provides compute resources at the physical layer that tenants can lease on demand in a self-service manner. BMS instances, each of which is a physical server, are physical computing environments that include basic server components: CPUs, memory, an operating system, hard disks, and bandwidth. Tenants have administrator

permissions for the instances that they create and can turn their servers on or off, mount hard disks, deploy environments, and perform other basic operations.

Like ECS, BMS also provides multiple layers of protection, including host and network security, remote access control, and other security management controls. For details, see section 4.1.1 ECS. More importantly, BMS has the unique security advantage of physical isolation. With its comprehensive security design covering hosts and networks, it ensures tenant security and offers a reliable, flexible, and high-performing application environment running on an isolated physical-layer compute environment.

5.2. Storage, Backup and Disaster Recovery Services

5.2.1. Elastic Volume Service

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

Different from traditional disk arrays, EVS evenly stores two or three identical copies of the same set of data on several storage nodes, ensuring data availability at highest level.

Security protection contains Anti-DDoS, IDS and WAF. Cloud service frontend provides service management portal, cloud service backend deployed as combination API. Console and API communicate with IAM, which provides the capabilities of authentication and authorization.

The lower layer platform is implemented by EIP¹³ firewall, VPN firewall and FusionSphere OpenStack (FSO), which contains cascading OpenStack and cascaded OpenStack.

5.2.2. Volume Backup Service

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

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

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

5.2.3. Object Storage Service

Object Storage Service (OBS) provides customers with massive, secure, reliable, and cost-effective data storage capabilities, including bucket creation, modification, and deletion as well as object upload, download, replication, modification, and deletion. It applies to the storage of any type of file and is suitable for ordinary subscribers, website, enterprises, and developers.

OBS has the following functions:

¹³ Elastic IP

- Creates and deletes buckets in specific regions.
- Manages objects, including uploading (such as multipart upload), replicating, downloading, and deleting objects.
- Manages bucket permissions, including bucket policies, access control lists (ACLs), and cross-origin resource sharing (CORS).
- Provides the access key and secret key for data access, controls ACL file permissions, and implements fine-grained access control on buckets.
- Supports data transmission over HTTP or SSL and allows users to encrypt data before storing it.
- Supports versioning to protect tenant data from both logical and physical failures, guarding against data loss from unintended user actions.
- Supports Server encryption, After server encryption of the OBS is enabled, and static data will be automatically encrypted by the symmetric AES256 key and decrypted when the data is downloaded or looked up.

Access controls

Requests to access OBS can be controlled through ACLs, bucket policies, and user signature verification.

- **Access control list (ACL):** OBS access permissions can be assigned to accounts by using an ACL. The ACL can grant all or certain accounts read, write, or full permissions on a per bucket or per object basis. Other access policies can also be configured, such as public access to a specified object (allowing all users read permissions only). By default, a bucket and the object(s) in the bucket can be accessed only by the creator of the bucket.
- **Bucket policy:** The owner of a bucket can create a bucket policy to restrict access to the bucket. Bucket policies restrict access in a centralized fashion based on many conditions: OBS operation, applicant, resource, and other request information (such as IP address). Permissions can be assigned for specific buckets and specific accounts.

Unlike ACLs, which only control permissions for single objects, bucket policies can affect multiple or all objects within buckets. Permissions for any number of objects in a bucket can be configured with a single request. Multiple objects can be specified by using wildcard characters in resource names and other fields, similar to regular expression operators, which allows the configuration of permissions for groups of objects.

OBS determines whether to accept or deny requests to access a bucket based on the policy configured for that bucket.
- **User signature verification:** To access OBS, users must provide an access key ID (AK) and secret access key (SK), which are authenticated by IAM. Therefore, OBS authenticates and authorizes user accounts with the AK and SK to ensure that OBS resources cannot be accessed without proper authorization. The headers of access requests sent to OBS contain authentication information generated based on the SK, request timestamp, and request type. OBS also independently performs URL encoding on bucket and object names before generating authorization information.

Only accounts that pass crypto based authentication and authorization can access OBS resources.

The OBS API is fully compatible with Amazon's Simple Storage Service (S3) interface so that tenants can conveniently and seamlessly migrate data from AWS to Flexible Engine. Tenants can securely and reliably migrate their AWS data from AWS locations, specified by Amazon Resource Name (ARN), to Flexible Engine using the AWS interface and Amazon Signature Version 2 or Version 4 Signing Process¹⁴.

Cross-Origin Resource Sharing (CORS)

For security reasons, browsers restrict cross-origin HTTP requests initiated from within scripts. For example, XMLHttpRequest and the Fetch API follow the same-origin policy. This means that a web application using those APIs can only request HTTP resources from the same domain the application was loaded from unless CORS headers are used. With the Cross-Origin Resource Sharing (CORS) policy enabled, assets such as web fonts and images stored in an OBS bucket can be safely referenced by external web pages, style sheets, and HTML5 applications.

Data Durability

Under double AZ, OBS designed for durability of 99.99999999% of objects, if customers store 1011 objects in OBS, after one year there maybe 1 object lost.

In one AZ, OBS Use 6+3 redundant way to storage data, that while three hard disk failures without any loss of data. OBS uses synchronous way to back up data between two AZ, Even if a datacenter suffers an accident was completely destroyed. Backup data is also in mirrored AZ. There will be no loss of data objects.

Data Availability

All OCB internal components are redundant carrier-class configuration, including load balancers, switches, servers, software modules, etc. All components of the cluster are designed for availability of 99.999%, with failure less than 5 minutes one year. However general datacenter availability is at 99.9%, Therefore, taking into account the availability of the datacenter, and external service availability is limited by the availability of the datacenter.

Access logs

OBS can log bucket access requests for use in analysis or auditing. These access logs allow the owner of a bucket to comprehensively analyze the nature and type of requests to access the bucket and identify trends. Once logging is enabled for a bucket, OBS automatically records all access requests into a log file that is written to a user-specified bucket. Note that because these logs occupy tenants' OBS space and may cause additional storage fees to be incurred, logging is disabled by default. Tenants can enable it manually if required for analysis or auditing purposes.

¹⁴ Version 4 uses the more secure HMAC-SHA256 algorithm and includes user data in signature computation. The header used during computation can be user-defined, greatly improving the security of authentication requests. It is therefore recommended that Amazon Signature Version 4 be used during migration.

URL anti-spoofing and validation

To prevent URL spoofing for OBS tenants, URL validation based on HTTP header and referer as well as access whitelists and blacklists are supported. The source website from which a user is linked to a destination website can be determined based on the header of the HTTP request. Requests that originate from an external website can then be denied or redirected to a specified web page. URL anti-spoofing and validation mechanism can also check requests against a blacklist or whitelist; access is granted when a match with a whitelist entry, otherwise denied or redirected to a specified web page.

5.2.4. Scalable File Service

Scalable File Service (SFS) provides an on-demand, scalable, and high-performance shared file system for Elastic Cloud Servers (ECSs). An elastic cloud server (ECS) of an AZ can access a file system of another AZ in the same virtual private cloud (VPC).

The service provides the standard file access protocol NFSv3. Users can integrate their existing applications and tools with SFS without any modifications on their applications and tools.

SFS consists of three components: Console, OceanStor DJ, and FusionStorage File. When tenants access SFS Console, the Load Balancer (LB) balances the load and sends access requests to SFS Console. After SFS Console receives these requests, it initiates authentication requests to IAM to authenticate the requesters. If requesters pass the authentication, their requests are sent to OceanStor DJ.

Network isolation

SFS isolates shared file systems of different tenants using VPCs. When creating a shared file system, a tenant must specify the VPC ID. In this way, the created shared file system can be accessed only by VMs of the specified VPC.

5.2.5. Dedicated Storage Service

Dedicated Storage Service (DSS) provides customization, flexibility and scalability within multiple dedicated storage pools. It offers multiple types of storage needs, including hybrid workload, high performance computing and Online analytical processing (OLAP) Applications. Dedicated Storage Service helps ensure that you're the only one accessing data on your DSS pool. It can also help satisfy certain compliance requirements. The storage nodes belongs to DSS Pools are physically separated from the storage nodes belongs to shared pools.

When all tenants access DSS console, their authentication and resource management are implemented through the token provided by IAM. Tenants can obtain tokens using their user names and passwords from IAM. In this way, tenants can pass the identity authentication when accessing DSS and then obtain related resource information.

Tenant resources provided by DSS include DSS pool and DSS disks. DSS disks created by a user belong to the user's account. By default, only users of that account can access the DSS volumes. The DSS pool belong to a user will only store the data of the user.

Platform Logging is integrated to ELK via ELK agent installed on cinder nodes which will collect all the operations of DSS disks to the ELK server.

5.2.6. Storage Disaster Recovery Service

Storage Disaster Recovery Service (SDRS) is a Disaster Recovery as a Service (DRaaS) across Flexible Engine AZs which is based on the block storage replication and enables an RPO of 0 and RTO < 4h. No computing of target VM resources are occupied as well as very little performance impact and does not require any agent inside the virtual machine. DRaaS can be configured via API or Console.

5.2.7. Cloud Server Backup Service

Cloud Server Backup Service (CSBS) offers the backup protection service for Elastic Cloud Servers (ECSs). It works based on the consistent snapshot technology for Elastic Volume Service (EVS) disks. It enhances data security with crash-consistent backups and applies to some restoration scenarios from malware infection, mis-deletion, application update errors or system breakdown.

Backups of all the EVS disks on an ECS are generated at the same point in time; however, applications and file systems on the ECS are not suspended before backup, and memory data is not backed up.

CSBS enhances data integrity and service continuity. For example, if an ECS is faulty or a misoperation causes data loss, you can use data backups to restore data quickly. CSBS combines ECS and Object Storage Service (OBS) to back up ECS data to object storage durability of 99.999999999% by design.

5.2.8. Data Express Service

Data Express Service (DES) is a transmission service oriented to TB- or PB-level data. It uses a physical storage medium (Teleport) to transmit a large amount of data to the Flexible Engine. DES has the following functions:

DES has a comprehensive security mechanism to protect user's data from being accessed and tampered with.

Security assurance of the migration medium: Teleport is dust and water proof and resistant to vibration and crush. The user can have logical protection of AES 256 software encryption for the data by using the KMS key with the DESClient software as well as physical protection of password locker.

- Security assurance during data transmission: After copying data to the Teleport, the user mails the device to Flexible Engine data center. The administrator receives and mounts the Teleport to the server. Then an SMS message is sent to notify the user of inputting the access keys (AK/SK) and the encryption key. After the verification is successful, data upload is triggered. In this manner, no others have any access to the user's keys or data, which ensures data security during transmission.
- Secure data transmission: Key Management Service (KMS) is a hosting service that allows you to easily create and control encryption keys for encrypting data. KMS uses hardware security modules (HSMs) to secure your keys. Specifically, the KMS ID of a Teleport task order is associated with a KMS key.
- Data integrity: The user can prepare the MD5 hash file for every file. DES Client can automatically load the MD5 file and check if the uploaded file is correct.

- Data sanitization: Data erasure can be performed on a Teleport using the SmartErase technology that follows the National Institute of Standards and Technology (NIST) guidelines for media sanitization.


5.3. Network Services

5.3.1. Virtual Private Cloud^

Virtual Private Cloud (VPC) enables customers to provision a logically isolated, configurable, and manageable virtual network environment, improving security of resources in a public cloud and simplifying network deployment.

Customers have complete control over their virtual network environment. They can use security groups to improve security of their network environments.

Customers can apply for a public IP address for a VPC to connect the VPC to the public network. They can also connect a VPC to their physical datacenter using a VPN, implementing smooth application migration to the cloud. IPsec VPN is supported.



- Private Subnet: Private subnet is owned and managed by a tenant himself. The tenant can manage the private (RFC 1918) addresses of each subnet. Each subnet is L2 isolated by means of VXLAN network.
- Virtual Router: Virtual Router provides L3 routing service for the VMs located in the VPC. Virtual router works as the gateway of the private subnet.
- Public IP: When a VM is going to communicate with the Internet, the tenant needs to apply a public IP and bind it to that VM.
- Firewall ACL: traffic control feature is provided to tenants to control access from Internet to his own VPC.
- Security group provides filtering for VMs and virtual firewall in a port granule. Security group can be used to isolate VM groups in a VPC.
- Source Net Address Translation (SNAT) provides local services for VPCs. For example, if a VM is to be connected with the NTP server located in local service area, SNAT is implemented automatically.
- Public IP implemented with bandwidth control.

As firewalls and security groups are both major factors in enhancing the cybersecurity of Flexible Engine VPCs, understanding the differences between them is important for creating effective network security policies for VPCs. These differences are summarized below for easy reference.

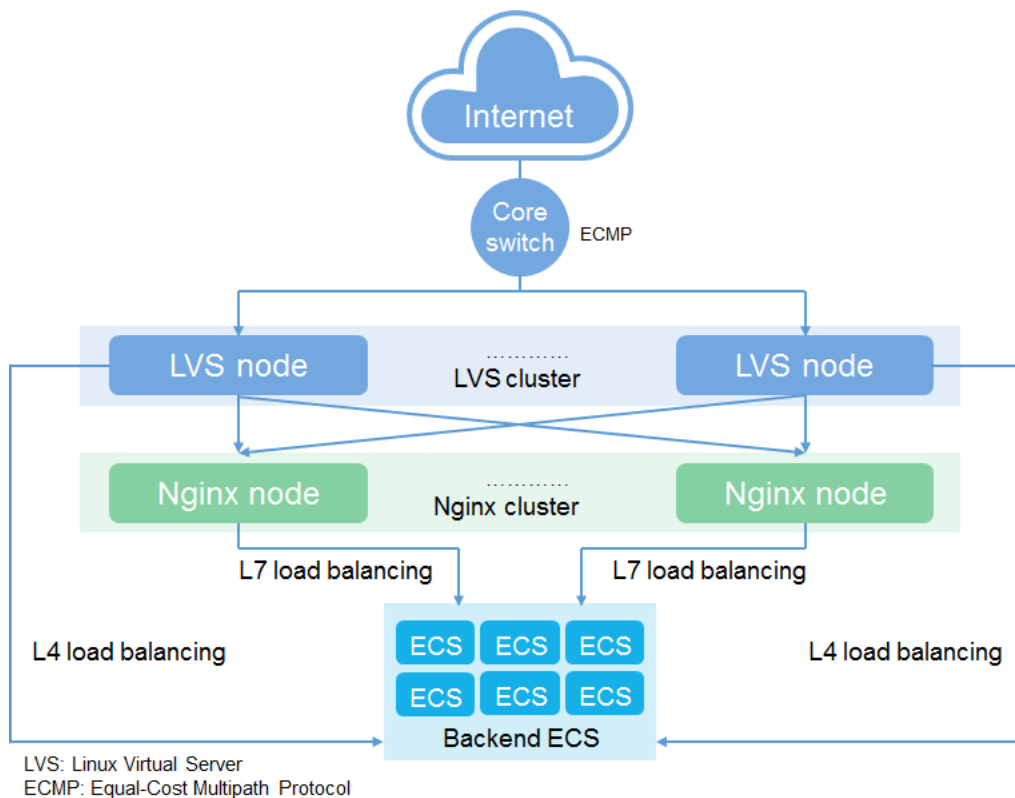
Security Groups	Firewalls
Work on the Elastic Cloud Server instance level (first layer of protection).	Work on the subnet level (second layer of protection).
Support permit policies.	Support permit, deny, and reject policies.
If rules conflict with each other, only the parts in agreement take effect.	If rules conflict with each other, only the first of the conflicting rules takes effect
Must be selected when an Elastic Cloud Server instance is created; take effect automatically on Elastic Cloud Server instances.	Cannot be selected during subnet creation. At a tenant's discretion if to set up a firewall, add associated subnets and inbound/outbound rules, and then enable the firewall in order for the associated subnets and the Elastic Cloud Server instances on those subnets to be protected.
Support packet filtering by 3-tuple (protocol, port, and destination IP address).	Support packet filtering by 5-tuple (protocol, source port, destination port, source IP address, and destination IP address).

5.3.2. Elastic Load Balancer

Elastic Load Balance (ELB) is a service that automatically distributes access traffic to multiple ECS to balance their service load. It enables tenants to achieve higher levels of fault tolerance in their applications and expand application service capabilities.

ELB allows the User to control load balancing actions by the User. A web-based self-service management console is provided to allow the User to configure the ELB service with ease and to quickly add services to implement load balancing among service resources.

- Supports automatically route traffic across multiple instances.
- Elastic, automatic scaling based on traffic demands.
- Supports automatically scales its request handling capacity in response to incoming application traffic.
- Supports linear expansion, no single point of failure.
- Supports ELB backend nodes in both public and private networks.
- Supports Layer 4 (over TCP) and layer 7 (over HTTP and HTTPS) load balancing.
- Working with AS group to implement automatic scaling based on service workload.



ELB provides the following security functions:

- **IP address and port masking:** ELB allows external networks to see only a single IP address and service port(s); the actual IP address(es) and port(s) used by the backend are not exposed. This prevents the disclosure of network information and minimizes the attack surface.
- **Automatic scaling based on traffic status:** ELB can work with AS to provide more flexible scaling and better DDoS mitigation than the traditional method of directly connecting to single backend server(s) and service(s).
- **ELB Intranet security groups:** ELB security groups can be created on a tenant's intranet to ensure that tenant instances only receive traffic from the load balancer. Tenants can define allowed ports and protocols to ensure that traffic in both directions is sent through ELB. For a detailed description of security groups, see section 4.3.1 VPC.
- **Source IP address transparency:** ELB can transparently transmit source IP addresses when listening for HTTP and HTTPS services. This enables tenants to perform source tracing, collect connection or traffic statistics, enforce blacklisting or whitelisting for source IP addresses, and perform other tasks needed to meet enhanced security requirements. By implementing these through their applications, tenants can more quickly detect and respond to attacks.
- **SSL/TLS offloading and certificate management:** With SSL/TLS offloading, the SSL/TLS encryption and decryption of packets is performed on ELB, reducing the processing burden that these tasks place on the tenant's backend server. In this process, encrypted traffic is sent to ELB for decryption and then delivered to the tenant's backend server; likewise, outbound traffic is sent to ELB for encryption and

then on to its destination. To use SSL/TLS offloading, tenants must upload any SSL/TLS certificates and private keys required to ELB for management.

- Support for encryption protocols and cipher suites: Tenants communicating with ELB over HTTPS can select an encryption protocol and configuration as required. TLS 1.2 is used by default. The default cipher suite was chosen to allow access from legacy browsers such as Internet Explorer 8. Tenants requiring higher security and more robust encryption algorithms can select them from ELB's extended list of cipher suites.

5.3.3. Domain Name Service

Domain Name Service (DNS) provides a highly available and scalable DNS administration service running authoritative domain name servers. DNS translates human-friendly domain names and application resources into IP addresses, which are used to establish the network connections that give end users access to desired resources.

DNS can resolve domain names into ECS, OBS, RDS, and other service addresses for ease of access to service resources. Users can resolve their internal domain names using the DNS service. Domain names and their associated addresses can be customized using VPC. By resolving the challenge of registering and managing domain names for tenants' internal services, DNS reduces the complexity of service deployment and maintenance while also making high-availability design a possibility. DNS is built on Flexible Engine's highly available and reliable infrastructure. The distributed nature of DNS servers helps to improve service availability and ensures that end users are routed to their desired applications. If a fault occurs on a service node, tenants can ensure service availability by modifying the domain name record to fail over to an operational node.

DNS provides the following key security functions:

Reverse lookup records (IP address to domain name) can be used to reduce spam emails.

- The DNS cache is protected against viruses and attacks by regular updates, shortened time-to-live (TTL), and frequent cleansing.
- DNS includes DDoS mitigation to ensure that services can operate in a stable and secure manner. Behavioral profiling is performed on inbound traffic, attack traffic is scrubbed, and access from malicious IP addresses is restricted or blocked. The Layer 7 protection algorithm used by DNS scrubs and filters traffic layer by layer, offering

complete protection against transport-layer and application-layer attacks. This anti-DDoS functionality also blocks DNS amplification attacks.

Tenants can use IAM to assign DNS service and permissions to their members. With access keys, resources can be accessed through API.

5.3.4. NAT Gateway

The NAT Gateway service offers the Network Address Translation (NAT) function for Elastic Cloud Servers (ECSs) in a Virtual Private Cloud (VPC), allowing these ECSs to access the

Internet using elastic IP addresses (EIPs) or to provide services for external networks.

The NAT Gateway service provides different types for different application scenarios.

- The NAT gateway type determines two elements of the Source Network Address Translation (SNAT) function, the maximum number of connections and the number of new connections per second. The data throughput is determined by the bandwidth of EIPs. SNAT allows resources that are not assigned EIPs in a VPC to access the public network directly and supports a huge number of concurrent connections
- DNAT supports port mappings. After DNAT rules are configured, packets are forwarded based on the rules. DNAT maps a public IP address (outside port) with a specified protocol to a private IP address (inside port). In this way, data from the Internet toward the public IP address will be forwarded to the configured private IP address. Users can control bandwidth resources more precisely.

The NAT gateway supports automatic disaster recovery through hot standby and provides the Cloud Eye service and alarm reporting for users, thereby reducing risks and improving availability.

The NAT gateway uses three network planes hence traffic between each plane is filtered for authorized access. The NAT gateway traffic from the ECS first passes through NAT gateway node by VXLAN, which is inside DC. Then the traffic passes the firewall to outside. The backward traffic from internet first passed through the firewall and the softNAT node. And then, it passed NAT gateway node to ECS.

5.4. Security and Identity Services

5.4.1. Anti-DDoS

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

Anti-DDoS provides the following functions:

- Self-service protection policy: Users can select the defense template that best meets the needs of their bandwidth and business model.
- Traffic inspection and scrubbing: Anti-DDoS checks traffic in real time and performs scrubbing on attack traffic when it reaches pre-defined threshold(s).
- Ease of administration: Users can view traffic trends and reports in real time on a management platform that is flexible and easy to use. The platform makes it simple to configure the service, set up stringent controls, and monitor service resources.
- Report monitoring: Users can query DDoS protection-related information about specific public IP addresses. This information includes current protection status, protection parameters, and the last 24 hours of information about traffic, scrubbing operations, and black holes. Security reports are available for review, displaying DDoS protection information generated by week. Users can query the past four

weeks of DDoS protection information, including but not limited to scrubbed traffic, number of intercepted DDoS attacks, and top 10 frequently attacked Elastic Cloud Servers.

- **Log analysis:** Anti-DDoS receives and analyzes logs reported by DDoS mitigation devices and displays the results on the management console.

5.4.2. Key Management Service

Key Management Service (KMS) is a centralized key hosting service built on a framework that ensures isolation between tenants. It uses a simple and convenient key management interface and provides easy-to-use and highly secure cloud-based encryption and key management functions.

KMS enables users to manage their keys conveniently and ensures the security of critical business data by supporting data encryption using a data encryption key (DEK) at any time. The DEK is encrypted using the customer master key (CMK) that is stored in KMS. The CMK, in turn, is encrypted by the root key, which is stored in the hardware security module (HSM), and saved as ciphertext on the key storage node. The chain of trust in KMS is rooted in the HSM, which is FIPS 140-2 (Level 2 and Level 3) certified to meet users' data security compliance requirements.

The KMS interconnects with cloud storage services such as EVS and OBS so that users can encrypt data stored in Flexible Engine simply by selecting the required CMK. It provides a convenient way for other cloud services to support full encryption of user data, especially sensitive data. Users can focus on their business and rest assured that the encryption keys for their data in Flexible Engine are properly managed.

To ensure the security and reliability of tenants' keys, KMS provides the following security features:

- **Random number generation:** All keys in the KMS are generated by the HSM's hardware true random number generator (TRNG) to ensure key randomization.
- **Secure key storage:** Key disclosure is prevented by storing the root key of the KMS in the HSM. The root key at no time appears outside the HSM. In addition, at least two HSM devices are deployed as a pair to ensure reliability and availability. The CMKs are encrypted using the root key and saved as ciphertext on the key storage nodes, which store them in a security-hardened MySQL database. The MySQL database is deployed in a cluster of two nodes, one active and the other standby. When a user key is saved to the active MySQL database, it is also backed up to the standby MySQL database so that the standby database can take over and continue to provide service in the event of a failure on the active node. The HSM is the root of trust in the KMS, and in combination with the other keys forms a complete chain of trust.
- **Postponed key deletion:** KMS is used to manage CMKs throughout their lifecycle, enabling, disabling, and deleting them. KMS also offers postponed key deletion. This function allows tenants to set a time (7 days to 3 years) during which CMK deletions can be canceled, avoiding CMK deletion by mistake.

- **KMS disaster recovery:** KMS provides a full range of backup and disaster recovery functions to ensure that keys are available and not lost. In the event of a disaster, KMS is switched over to a server at a secondary location to ensure service continuity. The root key stored in the HSM is backed up through its dedicated backup tool. The key storage node that stores CMKs backs up keys (incremental and full) on a regular basis to the specified storage device. If an error causes user keys to be lost, the KMS can recover the keys using the backup data.
- **Trusted links:** To ensure the security of KMS data, KMS hosts use a standard encrypted transmission mode to establish secure communication with the KMS service node.
- **Access control:** KMS performs centralized RBAC based on IAM roles. Operations on the CMKs stored in KMS can be performed only by users who have been authenticated by IAM and KMS and have the appropriate permissions. Users with read-only permissions can query information about CMKs but cannot perform other operations. In addition, KMS isolates the CMKs of different tenants so that tenants can access and manage their own CMKs only. Although system administrators have permissions to manage devices, they cannot access CMKs.
- **Operations log auditing:** Logs are generated for all major operations (such as creating a CMK or encrypting a DEK) and recorded to CTS so that CMK operations can be audited.

KMS also leverages other Flexible Engine technologies to enhance its security capabilities: namely, the secure infrastructure and platform, secure network architecture, perimeter protection, zone division, virtual network isolation, tenant KMS instance isolation, and API security.

5.5. Management tools and portals

5.5.1. Cloud Eye Service

Cloud Eye Service (CES) is one of the major IaaS services that helps customers monitor resources in Public Cloud. It collects statistics derived using monitoring metrics from different dimensions for resources to help them monitor resource changes in real time. CES monitors metrics related to Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Elastic Load Balance (ELB), Virtual Private Cloud (VPC), Auto Scaling (AS), and Relational Database Service (RDS).

Customers can also create, modify, delete, enable, and disable alarm rules for monitoring metrics on the management console. By configuring alarm thresholds and alerts, Users are informed of every resource status change, allowing them to respond quickly to resource exceptions and preventing service interruption.

CES generates alarms if certain thresholds are crossed that can trigger scaling actions implemented by the AS service. The alarm rules configured during AS group creation are synchronized to the CES. Then, the CES generates alarms and sends the alarms to AS for implementing scaling actions.

Like all FE services, CES requires that every request made to its control API be authenticated so only authenticated users can access and manage CES. Console and API communicate with IAM, which provides the capabilities of authentication and authorization.

5.5.2. Cloud Trace Service

Cloud Trace Service (CTS) records operations on cloud service resources so that they can be queried, audited, and traced. It records operations performed on the management console, executed through an API, and internally triggered on the Flexible Engine system. CTS is an essential support system for tenant-specific industry certification and IT compliance certification. It provides the following functions:

- **Resource change auditability:** Changes to Flexible Engine resource and system configurations performed by all users are recorded systematically and in real time. This is superior to the traditional method in enterprise IT environments of manually auditing each change.
- **Access security monitoring and auditability:** All management console operations and API calls are recorded systematically and in real time to help query, analyze, and locate issues closer to real time or after fact.
- **Data auditability:** Users can verify whether data has been disclosed by collecting activity data about OBS objects and object-level API events recorded by CTS for audit purposes.
- **Low cost:** CTS can merge records into event files on a regular basis and move these to an OBS bucket for storage, making logs highly available over a long period of time and at a low cost.

The security design for CTS is based on the Flexible Engine security framework. The security of the cloud computing services provided to tenants is ensured through secure network architecture and through the implementation of network perimeter, application, and data protection. Application and data security are described as follows.

- **Application security:** Valid requests for compliance event queries and tracker operations sent by legitimate users and also valid compliance events from interconnected services are accepted and processed by CTS. All requests must be transmitted over HTTPS. Sensitive data is encrypted, and a number of measures are taken to ensure security when interacting with external services: interface control, whitelist control, requestor authentication, and multiple rounds of verification. Furthermore, the web security of CTS control nodes has been hardened to defend against a wide range of attacks.
- **Data security:** The security requirements for user log data processed by CTS differ as the data is generated, transmitted, and stored. When generated, log data must be desensitized within each service and verified to contain no sensitive data. When transmitted, the accuracy and completeness of log data transmission and storage must be ensured through identity authentication, format validation, whitelist inspection, and unidirectional reception. When stored, log data must have multiple backup copies stored in a distributed manner, and databases must be hardened in accordance with Orange security requirements to prevent data security threats such

as spoofing, repudiation, tampering, and leakage. For additional security, CTS can be configured for encryption of log data when saved in an OBS bucket.

5.5.3. Tag Management Service

Tag Management Service (TMS) provides cross-region and cross-service tag management and resource classification. It also provides the predefined tag, export and import functions to help users migrate tags, and it can be accessed through the TMS console or using APIs.

Tags identify cloud resources so that they can be categorized easily and searched quickly. A cloud resource must have a unique key. the following services are supported by TMS: ECS, OBS, VPC, VBS, EVS, IMS, and AS. The creation/deletion request is dispatched by global LB and goes to TMS API service.

TMS calls service APIs in regions where the resources reside to modify resource tags. TMS invokes interfaces of the IAM service, OperationCenter, and Nova to complete interaction. Users obtain token from the IAM service and access predefined tag data on the TMS console or by invoking APIs.

TMS provides a GUI interface for tenant administrators and end users, a web console for tenant administrator and end users, and open APIs for other cloud service. The communication is encrypted using HTTPS protocol channel transmission.

5.5.4. Identity & Access Management

Identity and access management (IAM) enables enterprises to self-manage cloud resources and offers user roles, such as enterprise, department, employee, and business staff accounts, as well as role-based security governance functions.

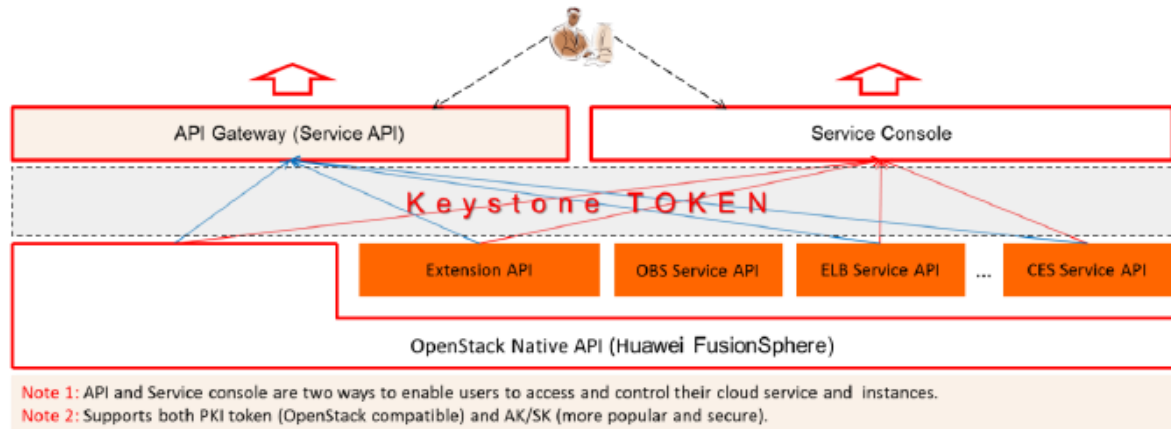
IAM is a core service of the cloud solution. Other services in the cloud solution are dependent on IAM for authentication. Only users explicitly authorized are allowed to access cloud resources. The IAM service restricts the cloud service access permissions of users based on their responsibilities. When a user accesses a website, the User are able to easily update the cloud service access permissions of the user to reflect the change in the role.

- Centrally controls user security certificates such as access key of each user (incl. Open API authentication).
- Centrally controls the user access policies such as ACL.
- Provides permissions based on user groups (administrator, engineer, etc.).
- Supports identity federation: A root account can create an enterprise identity provider which provides identity authentication for enterprise employees during daily operation.
- Enterprise employees can switch to the Public Cloud system on the enterprise network by using the SAML 2.0 SSO protocol.

Flexible Engine is based on OpenStack, and IAM is based the Keystone extension implementation.

- IAM Frontend: Used to provide authentication interface and account management UI (for example Login/Logout/Register).

- IAM Cached Proxy: Routed to core services, can provide local acceleration of Core Service.
- IAM Core Service: Based Keystone implementation, and used keystone patch to complete new features.



Service API includes two types:

- Native OpenStack API: providing OpenStack native API, including Nova, Cinder, Glance, KeyStone and Neutron;
- Extended Service API: extension for Nova, Cinder, Neutron etc. to provide enhanced features, and the extended new service APIs, such as ELB, CES, AS and etc.

5.5.5. VPC Flow Logs

Flow Log is a feature of VPC that enables user to capture information about the IP traffic going to and from network interfaces. Flow log data is published to LTS and can be view & search on LTS. Flow log agent on host is responsible for capturing the traffic log and sending to LTS. Monitor & audit the traffic that is reaching your instance using flow logs as a security tool.

5.5.6. Log Tank Service

Log Tank Service (LTS) collects and stores logs, allowing you to query them in real time. It simplifies decision making, helps you perform routine O&M, and improves log processing efficiency.

LTS uses permission controls and HTTPS encryption to ensure data reliability, supporting transfer of all logs to OBS buckets for storage. The collected logs are retained in LTS for 7 days and can be periodically transferred to OBS. LTS calls the IAM and OBS APIs for authentication and log storage. With Agent, LTS collects logs in real time, preventing your data from being mistakenly deleted or maliciously deleted by hackers.

Logs on the LTS console can be transferred to OBS buckets for long-time storage to meet compliance requirements. Log data of different tenant is saved to different index files of ES cluster and grouped by tenant's project id.

After interconnecting with Cloud Eye Service, you can monitor the number of keywords and set an alarm rule to monitor keywords of logs in real time. The supported services include VPC, IMS and OBS.

5.6. Analytics Services

5.6.1. Map Reduce Service

MapReduce Service (MRS) provides storage and analysis capabilities for massive data and builds a reliable, secure, and easy-to-use operation and maintenance (O&M) platform. Users can apply for use Hadoop, Spark, HBase, and Hive services to quickly create clusters and provide storage and computing capabilities for massive data analysis or real-time processing. After being processed and analyzed, data is encrypted by using Secure Sockets Layer (SSL) and transmitted to the Object Storage Service (OBS) system or Hadoop Distributed File System (HDFS). After data storage and computing are fulfilled, the cluster service can be terminated, or run them permanently.

There are two ways for users to create, operate and destroy of big data cluster: REST API and Console. After the cluster has been successfully created by user using default configuration, the user can submit MapReduce or spark jobs with the parameters specified by user, view the information of the cluster and the execution status of the tasks and modify the cluster configuration.

The major components include:

- **MRS Console:** the component provides all the user interaction features.
- **API Servers:** that provides API of each feature, such as deployment, template management, and operations management.
- **Cluster Management:** the component calls IaaS APIs to perform lifecycle management of clusters, which is composed of ECS instances. As MRS calls ECS API as a regular user, the ECS instances created by MRS can be seen by user from ECS portal.

Users can log in to MRS using its client or a web browser. The MRS supports single sign-on (SSO) based on central authentication service (CAS) so that users can conveniently access the web pages of other Big Data platform components without being prompted for authentication again.

- **User password management:** MRS uses IAM (Kerberos and LDAP) to manage user passwords. Kerberos encrypts user passwords and saves them to the LDAP database.
- **Permissions control:** MRS uses RBAC. Assigning a role to a user grants that user the permissions associated with the role. The permissions of each role can be configured based on the component resources that the role is required to access.
- **Data encryption:** The HBase and Hive functions of MRS support column-based storage encryption. When importing data, users can choose which data to store under encryption.
- **Data integrity:** MRS user data is stored in Hadoop Distributed File System (HDFS), which uses CRC32C to verify data integrity. Note that the default setting of CRC32C

can be changed to the slower CRC32 if desired. Verification data is stored on the HDFS DataNode (DN). If the DN detects that data transmitted from a client is abnormal (i.e. incomplete), the DN reports an exception to the client and asks it to write the data again. Data integrity is also verified by the client when reading data from the DN. If the client detects that the data is incomplete, it attempts to read the data from another DN.

- **Disaster recovery:** MRS provides geographically redundant disaster recovery for user data stored in its cluster along with a basic O&M tool for external systems that can set active and standby nodes, rebuild and verify data, and check the progress of data synchronization. To implement disaster recovery, MRS backs up data from the HBase cluster to another cluster, and clusters and data tables requiring synchronization are configured to trust each other. Note that HBase, HDFS, ZooKeeper, Kerberos, and LDAP Server must be installed on the cluster for this process to succeed. Once disaster recovery is configured, a standby cluster can immediately take over services in the event that data on the active cluster is damaged.

5.6.2. Data Pipeline Service

Data Pipeline Service (DPS) is a web service running on the public cloud. It enables the users to easily automate the movement and transformation of data between different services. DPS has successfully interconnected with OBS, RDS, and MRS, and will be able to interconnect with DWS, CDM, DIS, MLS, CloudStream, and CloudTable in the near future.

DPS is deployed in a high-availability and fault-tolerant environment to execute and monitor users' data processing flows, record the execution status, and handle exceptions (if any).

DPS management components are deployed in the dedicated subnet of the public service zone. DPS uses iptables to create a whitelist for access control. This prevents other cloud services from directly accessing the DPS Manager, Invoke Node, and MySQL database in the public service zone. Communications between DPS and all other services (such as IAM, CTS, CDR, and OBS.) are encrypted using SSL/TLS.

All calls to DPS APIs will be authenticated and all API parameters will be verified before being processed to prevent malicious attacks, such as XSS, CSRF, SQL inject, and Shell inject. Only pipeline owners (tenants and sub-accounts) can manage and run their pipelines. Accessing pipelines created by others is not allowed.

DPS just schedules activities to run in tenant's resources, for example, in MRS, RDS, DWS, MLS clusters. These resources are managed by corresponding cloud services, and DPS does not involve in managing these resources.

5.6.3. Data Warehouse Service

Data Warehouse Service (DWS) is a cloud-based online MPP (massively parallel processing)-based database. It is fast, stable, reliable, secure, scalable, easy to administer, and cost-effective.

DWS makes it easy to control network access to tenants' databases. DWS also lets tenants run their DWS clusters in a virtual private cloud (VPC), which enables tenants to isolate their DWS clusters and to connect their existing IT infrastructure to the DWS clusters through an Elastic IP (EIP). In addition, DWS supports the use of SSL to make data transmission more secure.

DWS depends on other cloud services to implement administrative functions, effectively benefit from each service's security features, including ECS, EVS, VPC, IAM, OBS, and CES. DWS system is composed of 3 planes: global, region and Point of Delivery (POD). The connections between planes are secured by HTTPS, SSH or other secure protocols. DWS Service accesses instances over SSH to execute management jobs.

5.6.4. Data Ingestion Service

Data Ingestion Service (DIS) is a highly scalable real-time streaming service which is capable of capturing and processing large amount of streaming data for specific needs. Data send to DIS can be stored for offline processing and analytics.

DIS synchronously replicates data across availability zones (AZ), providing high availability and data durability. DIS management components are deployed in dedicated subnet in Public Service zone. Communications between DIS and all other services (e.g. IAM, CTS, CDR, OBS, etc.) are encrypted using SSL/TLS.

Each user can create one or more DIS Stream, each Stream has one or more partition(s), and a partition has a corresponding folder in Kafka. Different users' partitions will not write to a same file in Kafka. Users can access Streams and partitions under the same tenant, but cannot access other tenants' Streams and partitions.

5.6.5. Machine Learning Service

Machine Learning Service (MLS) is an analysis platform service that helps users to identify patterns in data to construct machine learning models. Users can use this model to process new data and make predictions on service applications.

MLS depends on other cloud services to implement administrative functions, effectively benefit from each service's security features, including ECS, EVS, VPC, IAM, OBS, CES, MRS and CTS. MLS components are deployed in different network zones to reduce security risk. Communications between network zones are secured by HTTPS, SSH, or other secure protocols.

MLS instance VMs and storage resources are allocated under a resource tenant account, and final tenants are not allowed to log in to VMs directly.

MLS backs up each final tenant's instance metadata data to a different bucket, which belongs to a resource tenant. Therefore, each final tenant's backup data is logically isolated from the backup data of any other final tenant.

5.6.6. Cloud Stream Service

Cloud Stream Service (CS) is a real-time big data stream analysis service running on Flexible Engine. Cloud Stream Service focus on the Real-time analysis and the IoT scenario, is applicable to business scenarios requiring high throughput and low latency. Mainly used

in the Internet industry SME / IoT / financial anti-fraud and other industries applications scenarios, such as Internet cars, online log analysis, online machine learning, online graph processing, online recommendation algorithm applications.

Cloud Stream is fully managed, user cannot perceive the physical resources. The Cloud Stream internal resource tenant will apply ECS and VPC automatically for running clusters in POD zone. User cannot directly handle the ECS.

Cloud Stream components are deployed in different network zones to reduce security risk. Communications between network zones are secured by HTTPS, SSH or other secured protocols.

Cloud Stream is a web-based cloud service and consists of four parts: web console, Cloud Stream server, Cloud Stream backend and Cloud Stream cluster.

Cloud Stream manage two kinds of cluster:

- **Multi-user shared cluster:** Shared cluster is used for all user, they share the same VPC and VM resources. In order not to affect other users, it can only run Stream SQL jobs,
- **Tenant exclusive cluster:** Exclusive cluster is belong to tenant, it's physically isolated from other tenants including shared cluster, the VPS and ECS are isolated from each other to ensure access security. User can run user defined jar job and Stream SQL job in exclusive cluster.

All ECS and VPC resources are allocated by Cloud Stream resource tenant account, and final tenants are not allowed to log in to VMs directly.

The final domain tenant can manage their exclusive clusters logically. They can control how much Stream Process Unit (SPU) the cluster can use most. They also can allocate the exclusive cluster to their sub-users. Sub-users can run job on the clusters that assigned to them.

For security, the final user job's intermediate states and log will upload to user's OBS bucket. User should authorize to Cloud Stream service to read and write data to the bucket they choose, thus user can access the state and check log through OBS. In addition, Cloud Stream sends audit logs to CTS.

5.7. Database Services

5.7.1. Relational Database Service

Relational Database Service (RDS) is a cloud-based online relational database service. It is fast, stable, reliable, secure, scalable, easy to administer, and cost-effective. RDS based on the MySQL database engine provides standalone and primary-standby deployment. The database installation and deployment are performed automatically by RDS in minutes.

- Provides the access key and secret key for data access, controls ACL file permissions, and implements fine-grained access control (IP whitelist) on RDS instances.

- Supports data transmission over SSL and allows customers to encrypt data with SHA-2 before storing it.
- Supports Security Group and Subnet Group in VPC, so the DB instances can be isolated by the network for security purpose and to connect to them from the existing IT infrastructure through an industry-standard encrypted IPsec VPN.
- Supports multiple Replica instances and Multi-AZ deployment, backup data segments on different disks in redundancy mode.

Network isolation

RDS instances run in independent tenant VPCs and can also be deployed in subnet groups that span multiple AZs to provide high availability. After an RDS instance is created, the tenant is allocated an IP address in the subnet group for that instance to enable connection to the database. To control access to their databases, tenants can configure a range of IP addresses that are allowed to access their VPC(s) designated for database instance(s). After deploying an RDS instance on a VPC, tenants can configure a VPN to allow other VPCs to access it. Alternatively, tenants can deploy an Elastic Cloud Server in a VPC and connect to the database through a private IP address. Subnet groups and security groups can be configured in combination to isolate RDS instances and enhance instance security.

Access control

Creating an RDS instance also creates a primary account for the instance that its creator can use to perform operations on it. The password of this account can be set by the creator. The primary account can be used to connect to the instance, create sub-accounts, and assign database objects to those sub-accounts based on service planning. This provides a certain degree of security isolation. Furthermore, during database instance creation, a security group can be selected in which to deploy the NICs for the instance. VPC can be used to set inbound and outbound rules for the RDS instance and thereby control the scope of access to it. Only the database listening port is allowed to accept connections. Once configured, a security group immediately takes effect without the need to restart an RDS instance.

Transmission encryption

The connections between database clients and servers can be encrypted with TLS. A specified certificate authority generates a unique service certificate for each RDS instance upon provisioning. Database clients can download a root certificate from the management console and provide this certificate when connecting to the database to authenticate the server and enable encrypted transmissions.

Storage encryption

RDS can encrypt data before storage. Encryption keys are managed by KMS.

Automatic backup and snapshot

These features help recover RDS databases in the event of a fault. Automatic backup is enabled by default, and backups can be stored for a maximum of 35 days. Automatic backup allows tenants to perform point-in-time recovery (PITR) on their databases. Automatic backup performs a complete backup of all data and then incremental backups of transaction logs every 5 minutes so that a tenant can restore data to its status at any

second before the previous incremental backup. Tenants can also manually create a complete backup, known as a snapshot. Database snapshots are stored in OBS buckets and removed upon deletion of the corresponding database instance. New instances can be created based on existing snapshots.

Data replication

RDS instances can be deployed in a single AZ or across multiple AZs for high availability. When the latter option is chosen, RDS initiates and maintains data replication for database synchronization. High availability is achieved by having a secondary instance take over in the event that a failure occurs on the primary instance. It is also possible to create read-only MySQL database instances when operations are read-heavy. RDS maintains data synchronization between those read-only instances and primary instances, and tenants can connect to either type of instances as required by business to isolate read and write operations.

Data deletion

Removing an RDS instance will delete all data stored in that instance. No one can view or restore data once deleted.

Customer can easily perform the O&M of databases, including connecting their application and database instance, as well as data migration, backup and recovery and also monitoring. Customer can use the CES (Cloud Eye Service) Management Console to view key operational metrics, including CPU/memory/storage capacity utilization, I/O activity, and instance connections.

5.7.2. Distributed Cache Service

Distributed Cache Service (DCS) is reliable and scalable Redis cache service that enhances security and performance. DCS supports large-scale, highly-concurrent access in memory storage and can be used as a database, cache, or simple message queue.

DCS reports all changes to cache instances within the lifecycle to the Cloud Trace Service (CTS). Redis data commands may only be used by tenants, and Redis management commands may only be used in the DCS management plane.

DCS consists of three parts: the DCS Web Console, DCS Server and DCS Data Plane. Users can access the DCS Web console through a browser after they are authenticated by IAM over TLS encrypted channels. The DCS Server provides RESTful APIs to the DCS Web Console. These APIs are authenticated with the user's token to prevent unauthorized users from using these APIs. The DCS Data Plane component provides the cache service. Redis commands are authenticated with a password within the tenant VPC, only VMs in the tenant VPC can access the DCS Data Plane. In addition, security group rules can be used to further protect cache service and limit access users. The data plane of each cache instance is deployed in the respective tenant's VPC and is physically isolated from the cache instances of other tenants.

5.7.3. Document Database Service

Document Database Service (DDS) is a cloud-based NoSQL database service. It features high-performance storage, high-availability architecture, and primary/standby switchover. It

has a mature performance monitoring system, a multi-level security protection mechanism, and a professional database management platform.

DDS is compatible with MongoDB protocol and inherits same strong security access and SSL connection. Customers can also isolate their DDS instances further using VPC, subnet, security groups. User access is managed by IAM and CES will provide monitoring functions to view key operational metrics, including CPU/memory/storage capacity utilization, I/O activity, and database connections.

5.8. Application Services and PaaS

5.8.1. Simple Message Notification

Simple Message Notification is a simple, flexible, and scalable push message notification service. It enables users to send messages through emails, phones, HTTP/HTTPS applications in an efficient and inexpensive way. In addition, users can send messages to a group of subscribers in batches. SMN can be easily integrated with other cloud services to receive event notifications from them.

Access Control

Tenants use a username and password to log in to the console or use an AK/SK to call the SMN APIs to manage the topic resources.

By default, a tenant is not allowed to access topic resources of another tenant. The tenant can configure a topic policy to grant the permission to other tenants.

The service features a wide range of security measures to protect the management system from attack. It employs a tenant-based permissions model, strict parameter verification, secure communications protocols, and measures for protecting sensitive information and auditing logs.

Console accounts as well as users created by IAM and assigned SMN administrator permissions can perform all SMN operations. Users created by IAM but assigned tenant permissions can only perform query operations.

Data backup and availability

All SMN data is saved in GaussDB. Therefore, only data in GaussDB needs to be backed up. The GaussDB backup mechanism is as follows:

- GaussDB is fully backed up once a day, and the backup data is uploaded to the FTP server.
- GaussDB is incrementally backed up once every 4 hours, and the backup data is uploaded to the FTP server.

Data encryption

The SMN API can be accessed only through HTTPS, with TLS 1.2 and Perfect Forward Secrecy (PFS) enabled by default. Mobile phone numbers, email addresses, and other sensitive tenant data are stored under encryption using reliable algorithms. Audit logs can be stored over a long period of time to support tracing operations that may be necessary.

5.8.2. Workspace

Workspace provides virtual application services and Windows clients using Virtual Desktop Infrastructure (VDI). It enables users to access their cloud desktop over thin client hardware from anywhere at any time. Workspace provides a higher level of security than traditional PCs by isolating users' interfaces and data and thereby prevents data leakage by centrally storing and processing that data.

No data is stored on the thin clients used with Workspace; they are used for running the Workspace client program only. Desktop user interface is regenerated as graphics on the client side but the actual business data is not transmitted. Information transmitted to and from Workspace is sent over the highly secure Huawei Desktop Protocol (HDP), and input from local peripherals (USB devices, multimedia devices, flash storage, keyboards, and mice) is recalibrated for security.

Users can log in to their Workspace accounts at any time from thin clients in a tenant space or over Direct Connect. With these capabilities, Workspace offers greater efficiency and flexibility than the laptop computers and external storage devices of the past.

Workspace greatly improves maintenance efficiency by centrally managing password policies, session timeout, desktop delivery, peripheral devices, patches, and upgrades.

With virtualized management of all hardware, virtualized resources can be allocated to users as needed. This lengthens the service life of user desktops and reduces costs associated with replacing and upgrading hardware.

Workspace offers the following security functions:

- **User identification:**
 - Administrators and end users are assigned unique identifiers, which are linked to all auditable events.
 - All users must be authenticated using a password that meets predefined complexity requirements (such as password length or types of characters included).
 - A default timeout value is provided. If a user does not perform any action within a specified time, Workspace will automatically terminate the session or require reauthentication of the user.
 - Any user who unsuccessfully attempts to log in too many times within a specified interval will be locked out. This ensures user security and has the added benefit of limiting the number of authentication requests sent.
- **Access control:** The subjects (such as users and services) and objects involved in resource access along with any operations between subject and object are all within the scope of access control. Authorized users' permissions to access content and perform operations on protected resources are not allowed to exceed the preset scope. For added security, data related to user authentication is stored under encryption.
- **Transport security:** Desktop access is performed over HDP to ensure the security and integrity of data transmissions. The number of sessions to a single desktop can

be limited, and TLS 1.2 can be used to establish encrypted communication channels.

- **Image security:** The integrity and confidentiality of VM image files are protected, and residual data from VM images and snapshots is completely erased.
- **Backup and restoration:** A backup management mechanism for VDI ensures that backed-up data can be restored.
- **Security monitoring:** Security event data is processed and classified into different levels of alarms. Workspace can monitor in real time the online status and usage status of users, the operating status of VMs, and the online status of terminals, and can form the information obtained into security events.
- **Security auditing:** All user activities, operations, and commands that affect the system can be logged to support subsequent auditing. Log data includes the following: login type, operation type, log level, event time, event subject, IP address, event description, and event outcome. To ensure that audit logs are not lost, they are stored on non-volatile storage media and can be migrated if storage space becomes insufficient. Only authorized users can access and review system logs; logs cannot be accessed, modified, or damaged by unauthorized users. Audit logs can be queried based on one or more of the following: event type, event time, event subject, IP address, event outcome, and keywords. To ensure the confidentiality and integrity of transmissions, secure access is required to view logs.

5.8.3. Terraform

Terraform is an open-source resource management tool which manages OpenStack by compiling configuration file based on API call processes. Flexible Engine is certified by HashiCorp and joint provider list of Terraform.

Terraform is used to create, manage, and update infrastructure resources such as physical machines, VMs, network switches, containers, and more. Almost any infrastructure type can be represented as a resource in Terraform.

The infrastructure Terraform can manage includes low-level components such as compute instances, storage, and networking, as well as high-level components such as DNS entries, SaaS features, etc. Currently Supported resources in Flexible Engine are RDS, VPC, OBS, ECS, DNS, ELB, SMN, NAT, and DRS.

5.8.4. Resource Template Service

Resource Template Service (RTS) is an API which gives developers and systems administrators an easy way to create and manage a collection of related Flexible Engine's resources, provisioning and updating them in an orderly and predictable fashion.

RTS provides a native REST API (heat-api) and a CloudFormation compatible API (heat-api-cfn), in a similar way to many other OpenStack projects

- REST API uses keystone auth_token middleware
- REST API can optionally be configured to allow user/password authentication when heat is used in "standalone" mode (use local heat against a remote cloud tenants don't control)

- The CFN compatible API uses heat-specific ec2token middleware, which relies on the keystone ec2tokens extension for signature validation

It uses AES algorithm for encrypting sensitive data in the RTS DB (e.g. credentials). RTS uses SHA1 hash to identify the implementation signature of a template defined resource and SHA256 for signature validation of ec2 signed requests (via the ec2tokens keystone extension).

Resource Template Service (RTS) calls Flexible Engine's API.

5.8.5. ServiceStage

ServiceStage is a one-stop application platform service that provides development, deployment, governance and maintenance features for enterprise microservice applications. It embeds multi-tenant container cluster management, application orchestration, microservice DevOps, application performance management and software repository features.

ServiceStage components can be split into a couple of groups and they will be deployed into different zones. The Portal is deployed in the Global Zone, it provides web-based GUI. ServiceStage Base is deployed in the OM Zone, It provides installation, operation & maintenance functions. ServiceStage Core is deployed in the Region Zone, It provides core functions of ServiceStage, such as application scheduling, software repository, microservice framework, DevOps pipeline, cluster management and so on. In the Pod Zone, there is a CFE Cluster Manager, which manages the tenant's application and resources directly and is operated by Cluster Manager of Region Zone.

ServiceStage use IAM service to support user identity management and access control. After authenticate the user, IAM will return a token to user, which contains the tenant and role information.

ServiceStage itself and its subsystems record logs of operations on both OM and Tenant Management Layer via audit log SDK (Software Development Kit). This ensures the audit log has a unified standard output. ServiceStage will build Audit Anomaly Detection (AAD) capability based on the standard log output. There are two main tasks:

- big data analysis capabilities to process priority analysis logs include OS operation log, OS event log, application log and platform log
- interface with the SIEM system to provide abnormal behavior detection capabilities.

5.8.6. Distributed Message Service

Distributed Message Service (DMS) is a cloud message service supporting internet application development for users in public cloud environment. With the use of distributed cluster technology, DMS supports large-scale and high-concurrency access.

DMS Web console will interact with IAM to verify the tenant's account and password. API Gateway receives access API request from the tenant's application, interact with IAM to verify the tenant's AK/SK. API Service verifies the legality of token for all DMS API call requests, and also verifies the token that come from web console. Kafka Proxy receives access API request from LVS (Kafka API) and interacts with IAM to verify the tenant's

AK/SK. Message Broker receives access API request from LVS (Kafka API) and interacts with IAM to verify the tenant's AK/SK.

DMS does not depend on other services. Because DMS is located in public service areas, all tenants share one or more DMS clusters; tenant's resource is isolated by tenant name and queue name.

DMS is integrated with CTS to send change logs of message queues and consumption groups by REST API. Also, it is integrated with CES to send monitor information about message production and consumption to CES by REST API periodically.

5.9. Developer Tools and APIs

5.9.1. Open APIs

Flexible Engine is based on OpenStack™ technology leveraged and enriched by Huawei™. All Flexible Engine Public cloud services provide RESTful APIs. Using those APIs, developers can benefit from Infrastructure as Code (IaC) and can automate the use of the cloud to enable an incredibly rich set of possibilities for SaaS applications, for automated Continuous Integration, for interaction of the cloud with other environments.

Since the automation takes the important role in IaC, it is necessary to build a solid, deep, and automated defense system, an IAM system is used to control the authentication and authorization of users; WAF is deployed to strengthen the service console and API gateway security, to prevent hacking, to ensure the availability, confidentiality and integrity of the portal.

WAF is capable of detecting the unknown web attack based on abnormality detection, including:

- DDoS web attack
- Injection
- XSS and CSRF
- Cross directory
- Component vulnerability attack
- Identity forgery
- Web tamper-proofing

A. Tier Classification of a datacenter

Datacenter classification by the Uptime Institute is currently the only world-recognized reference base: it defines performance in terms of the continuity of services and technical infrastructure of a datacenter. This classification is deliberately simplified into five main levels (called tiers), shown in the following table, corresponding to the main infrastructure categories, the initial deployment phases and the handling capacities in terms of electrical load (W/m²). Each level is assigned a statistical availability index based on the functioning history of dozens of major datacenters and on reliability studies carried out using expert software.

	Tier I	Tier II	Tier III	Tier III+	Tier IV
Source of replacement	1 generator	1 generator	1 generator	1 generator	1 generator
Primary energy	1 active line	1 active line	1 active line and 1 passive line	Rated power after major fault	Rated power after major fault
High quality energy	1 active line	1 active line	1 active line and 1 passive line	2 active lines	2 active lines
High Quality redundancy	N	N+1	N+1	2N or 2(N+1)	2N or 2(N+1)
Maintainable without operation shutdown	No	No	Yes	Yes	Yes
Maximum charge level (Watt/m ²)	250 W/m ²	500 W/m ²	1000 W/m ²	>1000 W/m ²	>1000 W/m ²
Theoretical availability rate	99.671%	99.749%	99.962%	99.990%	99.995%
Maximum number of hours of service down time per year averaged over several years and datacenters	28.82 h/year	21.98 h/year	1.58 h/year	0.87 h/year	0.44 h/year

The infrastructures set up by Orange in its generation 2004-2010 datacenters are classified at TIER III+ level. The latest generation datacenters, built after 2010, are in tier IV. In some of them, the air conditioning installations are equipped with the free-cooling function.