Security and Backup Policy

Latest revision: June 2 2020

Overview

This document is designed to provide an overview of the procedures that we have in place to ensure that your systems are available and your confidential information is secure. It covers the information about your business and processes that you share with us, and your data that you maintain on our hosted environment. **You should read this policy in conjunction with our** Service Level Agreement **and service** Terms & Conditions**.**

Responsibility matrix

Regarding security, some areas are the responsibility of our hosting provider, which rents us a data center. Some areas are the responsibility of Valiantys, which handles the service installation and support. Finally, some areas are the responsibility of Atlassian, that develop & maintain the software.

RASCI Matrix		
R	Responsible – who is responsible for carrying out the entrusted task ?	They will be included into contracts and during customer onboarding. **They SHOULD be SIGNED by all the stakeholders.**
A	Accountable (also Approver) – who is responsible for the whole task and who is responsible for what has been done ?	
S	Support – who provides support during the implementation of the activity / process / service ?	
C	Consulted – who can provide valuable advice or consultation for the task ?	
I	Consulted – who can provide valuable advice or consultation for the task ?	

Valiantys Support

		Client	Valiantys	Atlassian	App vendors
Tools administration	**Upgrade apps (formerly add-ons)** (upgrade an app/add-on directly on Jira, Confluence, …)	R	A	–	S
	Customize the configuration items (workflows, notifications, …) (manage configuration items inside Jira, Confluence, …)	R	A	S	S
	Projects/Spaces/Repositories/User Directories creation/modification/deletion (manage projects, spaces, repositories, user directories on Jira, Confluence, …)	R	A	S	S
	User access management (manage users on Jira, Confluence, …)	R	A	–	–

		Client	Valiantys	Atlassian	App vendors
	Functional issues management (global issue and project management)	R	A	S	S
	Products & apps' licenses management (license installation of Atlassian's products and add-ons)	R	A	–	–
Infrastructure	**Install environment** (Installation of Jira, Confluence, …)	R	I	–	–
	Patch management on Atlassian application and apps/add-ons (Deployment of patches related to apps/add-ons, Jira, Confluence, …)	R (Silver/ Gold) A (Platinum)	S (Silver /Gold) R (Platinum)	S	S
	Patch management OS and middleware (RProxy, database, mail server, …) (Deployment of OS and middleware patches)	R	I	–	–

	Client	Valiantys	Atlassian	App vendors
Upgrade OS and middleware (RProxy, database, mail server, …) (Upgrade of OS and middleware)	R	I	–	–
Upgrade Atlassian's products (Upgrade of Jira, Confluence, …)	R	S	S	S
System & applicative performances (Application's performances)	R	S	S	S
System & applicative monitoring (Monitoring health and performances of the system and application)	R	I	–	–
Application down (Application or system unresponsive/unavailable)	R	S	S	S
Security incident	R	S	S	S

		Client	Valiantys	Atlassian	App vendors
	(Events that could impact you security/control of your data (CVE release, publications, …)				
	Outgoing mail issues (Issues impacting outgoing emails from your Jira, Confluence, …)	R	S	S	S
	Incoming mail issues (Issues impacting incoming emails to your Jira, Confluence, …)	R	S	S	S
Offering management	**Manage support contacts** (Management of your declared contacts on our support portal)	R	A	–	–
	Support offering reporting (Transmission of a monthly report)	I	R	–	–

		Client	Valiantys	Atlassian	App vendors
	Support offering follow-up meetings (Meeting follow-up scheduled by Valiantys)	**I for Platinum only**	**R for Platinum only**	–	–

Valiantys Cloud

		Client	Valiantys	Cloud provider	Monitoring provider	Atlassian	App vendors
Infrastructure	**Install environment** (Installation of Atlassian applications – Jira, Confluence, … – in Server, Cold Standby or Data Center architectures)	I	R	S	–	S	–
	VPN configuration (Secure connection between the VPC dedicated to the customer and the customer's network via a VPN)	R	R	S	–	–	–
	Patch management on Atlassian application (Deployment of patches related to the applications – Jira,	I	R	–	–	S	S

	Client	Valiantys	Cloud provider	Monitoring provider	Atlassian	App vendors
Confluence, … – or the apps/add-ons)						
Atlassian application upgrades (Jira, Confluence, …)	R	A	–	–	S	S
System & applicative performances (Management of the resources – CPU, RAM – of the environments running the Atlassian applications)	A	R	S	S	S	S
System & applicative monitoring (Monitoring health and performances of the System and the applications)	I	R	S	S	–	–
Patch management on System / Middleware (Deployment of patches of Ubuntu OS and/or middleware applications – Apache, PostgreSQL, Postfix, …)	I	R	S	–	–	–
System / Middleware upgrades (Ubuntu OS, Apache, DB, Postfix, …)	I	R	–	–	–	–

		Client	Valiantys	Cloud provider	Monitoring provider	Atlassian	App vendors
	Application down (System or Atlassian application unresponsive / unavailable)	I	R	A	–	S	S
	Security incident (Events that could impact the Security or the control of your data (CVE release, publications, …)	I	R	S	–	S	S
	Outgoing mail issues (Issues impacting outgoing emails from the Atlassian applications – Jira, Confluence, …)	I	R	S	–	S	S
	Incoming mail issues (Issues impacting incoming emails to the Atlassian applications – Jira, Confluence, …)	I	R	–	–	S	S
Offering management	**Technical support contacts management** (Management of the declared contacts on Valiantys Support portal)	R	A	–	–	–	–

	Client	Valiantys	Cloud provider	Monitoring provider	Atlassian	App vendors
Cloud offering reporting (Monitoring / Uptime reports)	I	R	–	–	–	–

Definitions

Term	Description
AWS regions & availability zones	Amazon EC2 is hosted in multiple locations world-wide. These locations are composed of regions and Availability Zones. Each **region** is a separate geographic area. Each region has multiple, isolated locations known as **Availability Zones**.
VPC & Subnet	A **Virtual Private Cloud (VPC)** is an on-demand configurable pool of shared computing resources allocated within a **public** cloud environment, providing a certain level of isolation between the different organizations (denoted as **users** hereafter) using the resources. The isolation between one VPC user and all other users of the same cloud (other VPC users as well as other public cloud users) is achieved normally through allocation of a Private IP Subnet and a virtual communication construct (such as a VLAN or a set of encrypted communication channels) per user. In a VPC solution, the previously described mechanism, providing isolation within the cloud, is accompanied with a VPN function (again, allocated per VPC user) that secures, by means of authentication and encryption, the remote access of the organization to its VPC cloud resources. With the introduction of the described isolation levels, an organization using this service is in effect working on a **'virtually private'** cloud (that is, as if the cloud infrastructure is not shared with other users), and hence the name VPC.
Security group	A **security group** acts as a virtual firewall that controls the traffic for one or more instances. When you launch an instance, you associate one or more security groups with the instance. You add rules to each security group that allow traffic to or from its associated instances. You can modify the rules for a security group at any time; the new rules are automatically applied to all instances that are associated with the security group. When we decide whether to allow traffic to reach an instance, we evaluate all the rules from all the security groups that are associated with the instance.
Elastic IP	An **Elastic IP address** is a static IP address designed for dynamic cloud computing. An Elastic IP address is associated with your AWS account. With an Elastic IP address, you can mask the failure of an instance or software by rapidly remapping the address to another instance in your account.

Data Center location

Amazon Web Services	Orange Flexible Engine
US East N. Virginia, Ohio	**France** Pantin, Aubervilliers et Saint Denis
US West N. California, Oregon	
Canada Central	
Europe Frankfurt, Dublin, London, Paris, Stockholm	

Standards

Valiantys aligns with ISO 27001, supported by strong processes, documentation and culture. Data centres used by us have the following accreditations:

Amazon Web Services	Orange Flexible Engine
· SOC 1/SSAE16/ISAE 3402 (formerly SAS 70) – report available after signing a NDA · SOC 2 · **SOC 3** · FISMA, DIACAP, and FedRAMP · ISO 9001 / ISO 27001 · ITAR · FIPS 140-2 · MTCS Level 3 · GDPL (details here: **https://aws.amazon.com/blogs/security/aws-and-the-general-data-protection-regulation/**) · See **https://aws.amazon.com/fr/compliance/** for more details.	· ISAE 3402 Type II (ex SAS 70) · SOC 1 Type II · SOC 2 Type II · ISO 9001 / ISO 27001 · ISO 14001:2004 · ISO 22301 · OHSAS 18001 · HDA/HADS · PCI-DSS · See **https://cloud.orange-business.com/en/3rd-az-in-paris-region/**

Physical and environmental security

Fire detection and suppression

Amazon Web Services	Orange Flexible Engine
Automatic fire detection and suppression equipment has been installed to reduce risk. The fire detection system utilizes smoke detection sensors in all data center environments, mechanical and electrical infrastructure spaces, chiller rooms and generator equipment rooms. These areas are protected by either wet-pipe, double-interlocked pre-action, orgaseous sprinkler systems.	Fire protection is provided by two complementary systems: high sensitivity smoke detection and a water mist fire-fighting system. Supervision center

Power

Amazon Web Services	Orange Flexible Engine
The data center electrical power systems are designed to be fully redundant and maintainable without impact to operations, 24 hours a day, and seven days a week. Uninterruptible Power Supply (UPS) units provide back-up power in the event of an electrical failure for critical and essential loads in the facility. Data centers use generators to provide back-up power for the entire facility.	Energy management: in order to reduce infrastructures ecological footprint, Orange has selected eco-friendly data center. To avoid any power outage, the power supply is doubled and the sites have autonomous generators. Redundant supply systems Redundant general table of low-voltage distribution Generator – with 72-hour minimal autonomy Inverters

Climate and Temperature

Amazon Web Services	Orange Flexible Engine
Climate control is required to maintain a constant operating temperature for servers and other hardware, which prevents overheating and reduces the possibility of service outages. Data centers are conditioned to maintain atmospheric conditions at optimal levels. Personnel and systems monitor and control temperature and humidity at appropriate levels.	N+1 redundancy, meaning that the failure of one air conditioning tool has no impact on the cold production chain

Storage device decomissioning

Amazon Web Services	Orange Flexible Engine
When a storage device has reached the end of its useful life, AWS procedures include a decommissioning process that is designed to prevent customer data from being exposed to unauthorized individuals. AWS uses the techniques detailed in DoD 5220.22-M (“National Industrial Security Program Operating Manual“) or NIST 800-88 (“Guidelines for Media Sanitization”) to destroy data as part of the decommissioning process. All decommissioned magnetic storage devices are degaussed and physically destroyed in accordance with industry-standard practices.	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.

Physical access to the Data Center

Amazon Web Services	Orange Flexible Engine
AWS's data centers are state of the art, utilizing innovative architectural and engineering approaches. Amazon has many years of experience in designing, constructing, and operating large-scale data centers. This experience has been applied to the AWS platform and infrastructure. AWS data centers are housed in nondescript facilities. Physical access is strictly controlled both at the perimeter and at building ingress points by professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized staff must pass two-factor authentication a minimum of two times to access data center floors. All visitors and contractors are required to present identification and are signed in and continually escorted by authorized staff. AWS only provides data center access and information to employees and contractors who have a legitimate business need for such privileges. When an employee no longer has a business need for these privileges, his or her access is immediately revoked, even if they continue to be an employee of Amazon or Amazon Web Services. All physical access to data centers by AWS employees is logged and audited routinely.	Physical security: rigorous access control of data centers, including identifying persons after authorization is granted to enter data rooms. The sites are protected and supervised 24×7 with triple access control and video surveillance systems. 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.

Business Continuity Management

Availability

Amazon Web Services	Orange Flexible Engine
Data centers are built in clusters in various global regions. All data centers are online and serving customers; no data center is "cold." In case of failure, automated processes move customer data traffic away from the affected area. Core applications are deployed in an N+1 configuration, so that in the event of a data center failure, there is sufficient capacity to enable traffic to be load-balanced to the remaining sites.	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

Amazon Web Services	Orange Flexible Engine
	option. Cloud Orange services are operated from several operating platforms. Thus, if an operation site becomes totally unavailable, administration continuity is ensured for the Cloud services. Operational continuity is regularly tested through simulations.

Company-wide executive review

Amazon Web Services	Orange Flexible Engine
Amazon's Internal Audit group has recently reviewed the AWS services resiliency plans, which are also periodically reviewed by members of the Senior Executive management team and the Audit Committee of the Board of Directors.	A Security Committee meeting – including the Security Managers (Engineering and Operations) and the Product Manager – monthly reviews: • The risks • The security indicators • The security action plan • The security strategy

Network security

Secure network architecture

Amazon Web Services	Orange Flexible Engine
Network devices, including firewall and other boundary devices, are in place to monitor and control communications at the external boundary of the network and at key internal boundaries within the network. These boundary devices employ rule sets, access control lists (ACL), and configurations to enforce the flow of information to specific information system services. ACLs, or traffic flow policies, are established on each managed interface, which manage and enforce the flow of traffic. ACL policies are approved by Amazon Information Security. These policies are automatically pushed using AWS's ACL-	The architectures apply the same model of separating the trust areas, including the Back-End (internal Orange) and the Front-End – which supports the Cloud services exposed to the customers. The Back-End – Front-End partitioning is physical, which means that some servers are dedicated to the Back-End and other ones are dedicated to the Front-End. Within each trust area, security zones provide a logical partitioning by the implementation of certain features such as: Virtualization (virtual machines, virtual firewall, load-balancer virtual, virtual routers / VRF) VLANs (802.1Q)

Amazon Web Services	Orange Flexible Engine
Manage tool, to help ensure these managed interfaces enforce the most up-to-date ACLs.	Virtual networks (VPN IPSec, SSL) Virtual storage (virtual drives) This logical partitioning guarantees the insulation of the different customer environments. The communications (network flow) between the different security zones are systematically controlled by firewalls (stateful filtering). The local configuration of the various components is also constructed to reinforce the partitioning and security (ex: ACL in routers).

Secure access points

Amazon Web Services	Orange Flexible Engine
AWS has strategically placed a limited number of access points to the cloud to allow for a more comprehensive monitoring of inbound and outbound communications and network traffic. These customer access points are called API endpoints, and they allow secure HTTP access (HTTPS), which allows you to establish a secure communication session with your storage or compute instances within AWS. To support customers with FIPS cryptographic requirements, the SSL-terminating load balancers in AWS GovCloud (US) are FIPS 140-2-compliant. In addition, AWS has implemented network devices that are dedicated to managing interfacing communications with Internet service providers (ISPs). AWS employs a redundant connection to more than one communication service at each Internet-facing edge of the AWS network. These connections each have dedicated network devices.	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. The customers connect to Cloud environments for purposes of administration or access to applicative services. The environments can be accessed via the Internet and / or a private network (VPN client), based on the features and options selected by the customer. The customers' access are always secured using SSL / TLS (HTTPS or SSH flows): • The authentication to a server (administration portals, application servers) is routinely performed via a "server" X.509 certificate delivered by a recognized certification authority. • The authentication of a customer is performed using a login / password combination previously transmitted via a secured process. For some offers, it is possible to implement strong authentication (token generating a

Amazon Web Services	Orange Flexible Engine
	password for single use, “client” X.509 certificate). • The network flows are systematically encrypted (using AES algorithm).

Transmission protection

Amazon Web Services	Orange Flexible Engine
You can connect to an AWS access point via HTTPS using Secure Sockets Layer (SSL), a cryptographic protocol that is designed to protect against eavesdropping, tampering, and message forgery. For customers who require additional layers of network security, AWS offers the Amazon Virtual Private Cloud (VPC), which provides a private subnet within the AWS cloud, and the ability to use an IPsec Virtual Private Network (VPN) device to provide an encrypted tunnel between the Amazon VPC and your data center.	The customers connect to Cloud environments for purposes of administration or access to services. The customers’ administration flows are systematically secured by protocols to ensure the authentication, the confidentiality and the integrity (SSLv3/TLS, AES256, …). The access methods vary depending on the features and options selected by the customer. The safety of service flows depends on the offer, but in general the exchanges are secured with SSL/TLS connections. 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.

Fault-tolerant design

Amazon Web Services	Orange Flexible Engine
AWS has designed its systems to tolerate system or hardware failures with minimal customer impact. Data centers are built in clusters in various global regions. All data centers are online and serving customers; no data center is “cold.” In case of failure, automated processes move customer data traffic away from the affected area. Core applications are deployed in an N+1 configuration, so that in the event of a data center failure, there is sufficient capacity to enable traffic to be load-balanced to the remaining sites.	TIER classification – The data centers for Flexible Engine are built and operated to be compatible with TIER III requirements. 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.

Network monitoring and protection

Amazon Web Services	Orange Flexible Engine
AWS utilizes a wide variety of automated monitoring systems to provide a high level of service performance and availability. AWS monitoring tools are designed to detect unusual or unauthorized activities and conditions at ingress and egress communication points. These tools monitor server and network usage, port scanning activities, application usage, and unauthorized intrusion attempts. The tools have the ability to set custom performance metrics thresholds for unusual activity. AWS security monitoring tools help identify several types of denial of service (DoS) attacks, including distributed, flooding, and software/logic attacks. When DoS attacks are identified, the AWS incident response process is initiated. In addition to the DoS prevention tools, redundant telecommunication providers at each region as well as additional capacity protect against the possibility of DoS attacks. The AWS network provides significant protection against traditional network security issues, and you can implement further protection. The following are a few examples: • **Distributed Denial Of Service (DDoS) Attacks**. AWS API endpoints are hosted on large, Internet-scale, world-class infrastructure that benefits from the same engineering expertise that has built Amazon into the world’s largest online retailer. Proprietary DDoS mitigation techniques are used. Additionally, AWS’s networks are multi-	AWS uses a wide variety of automated monitoring systems to provide a high level of performance and service availability. • **DDoS** Orange Business Services 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. 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. • **ECS Security Architecture**

homed across a number of providers to achieve Internet access diversity.

- **Man in the Middle (MITM) Attacks**. All of the AWS APIs are available via SSL-protected endpoints which provide server authentication. Amazon EC2 AMIs automatically generate new SSH host certificates on first boot and log them to the instance's console. You can then use the secure APIs to call the console and access the host certificates before logging into the instance for the first time. We encourage you to use SSL for all of your interactions with AWS.
- **IP Spoofing**. Amazon EC2 instances cannot send spoofed network traffic. The AWS-controlled, host-based firewall infrastructure will not permit an instance to send traffic with a source IP or MAC address other than its own.
- **Port Scanning**. Unauthorized port scans by Amazon EC2 customers are a violation of the AWS Acceptable Use Policy. Violations of the AWS Acceptable Use Policy are taken seriously, and every reported violation is investigated. Customers can report suspected abuse via the contacts available on our website at: **http://aws.amazon.com/contact**-us/report-abuse/. When unauthorized port scanning is detected by AWS, it is stopped and blocked. Port scans of Amazon EC2 instances are generally ineffective because, by default, all inbound ports on Amazon EC2 instances are closed and are only opened by you. Your strict management of security groups can further mitigate the threat of port scans. If you configure the security group to allow traffic from any source to a specific port, then that specific port will be vulnerable to a port scan. In these cases, you must use appropriate security measures to protect listening services that may be essential to their application from being discovered by an unauthorized port scan. For example, a web server must clearly have port 80 (HTTP) open to the world, and the administrator of this server is responsible for the security of the HTTP server software, such as Apache. You may request permission to conduct vulnerability scans as required to meet your specific compliance requirements. These scans must be limited to your own instances and must not violate the AWS Acceptable Use Policy. Advanced approval for these types of scans can be initiated by submitting a request via the website at: **https://aws**–**portal.amazon.com/gp/aws/html**-forms-controller/contactus/AWSSecurityPenTestRequest
- **Packet sniffing by other tenants**. It is not possible for a virtual instance running in promiscuous mode to receive or "sniff" traffic that is intended for a different virtual instance. While you can place your interfaces into promiscuous mode, the hypervisor will not deliver any traffic to them that is not addressed to them. Even two virtual instances that are owned by the same customer located on the same physical host cannot listen to each other's traffic. Attacks such as ARP cache poisoning do not work within Amazon EC2 and Amazon VPC. While Amazon EC2 does provide ample protection against one customer inadvertently or maliciously attempting to view

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.

- **VPC**

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.

- **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

- **Blackholing probes deployed in the heart of Orange Group network**

This system protects against massive attacks by denial of service from the Internet, and is triggered as soon as the attack reaches 1 Gb of traffic. It is deployed in the heart of the Orange Group's operator network and it uses safety equipment specifically designed for this purpose. The solution offers two levels of service, one standard and one optional:
Activation of a function of type "black-hole" / "blackholing": this solution makes the attacked IP addresses unreachable from the Internet and protects the other customers of the Public Cloud against potential edge effects, such as performance degradation of Internet bandwidth. This is included by default in the service.
Re-routing of flows to a real-time cleaning center and then redirection to the healthy traffic. This is a paid optional service.

- **Firewalls located at the entrance of the Orange cloud platforms**

Firewalls protect the Orange Cloud platforms against attacks with weaker volumes (less than 2 GB) but in a finer manner, filtering only the packets deemed suspicious by analyzing certain protocols, such as:
IP protocols: SYN Flood, UDP Flood, ICMP Flood, Land Attack, ICMP fragment, ICMP Large packets, SYN fragments.
Applicative flows: firewalls are also able to detect

another's data, as a standard practice you should encrypt sensitive traffic.

In addition to monitoring, regular vulnerability scans are performed on the host operating system, web application, and databases in the AWS environment using a variety of tools. Also, AWS Security teams subscribe to newsfeeds for applicable vendor flaws and proactively monitor vendors' websites and other relevant outlets for new patches. AWS customers also have the ability to report issues to AWS via the AWS Vulnerability Reporting website at: **http://aws.amazon.com/fr/security/vulnerability-reporting/.**

specific attacks on some applicative protocols: DNS, FTP, HTTP, REAL, RSH, RTSP, TALK, TFTP, XING, …

- **Detection probes for intrusion attempts on shared Cloud services, located on the Orange Cloud platforms**

IDS probes are used to monitor Cloud administration services (administration portal, administration services, …). The virtual machines dedicated to a customer are not monitored by default because the analysis of the alerts requires knowing the customer's context. The implementation of the monitoring of the customer's servers can be proposed as an option.

More generally, the monitoring implemented on all the components of the infrastructure ensures a global intrusion detection

Account review and audit

Amazon Web Services	Orange Flexible Engine
Accounts are reviewed every 90 days; explicit re-approval is required or access to the resource is automatically revoked. Access is also automatically revoked when an employee's record is terminated in Amazon's Human Resources system. Windows and UNIX accounts are disabled and Amazon's permission management system removes the user from all systems. Requests for changes in access are captured in the Amazon per missions management tool audit log. When changes in an employee's job function occur, continued access must be explicitly approved to the resource or it will be automatically revoked.	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.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.

Background checks

Amazon Web Services	Orange Flexible Engine
AWS has established formal policies and procedures to delineate the minimum standards for logical access to AWS platform and infrastructure hosts. AWS conducts criminal background checks, as permitted by law, as part of pre-employment screening practices for employees and commensurate with the employee's position and level of access. The policies also identify functional responsibilities for the administration of logical access and security.	Orange has formal policies in place to ensure background checks are executed within their team during the hiring process, conforming with local regulations.

VPN

Amazon Web Services	Orange Flexible Engine
A VPN connection can be set up as an option between your network and your VPC on AWS, in order to integrate tools from your infrastructure and tools hosted on our Cloud offer, e.g. LDAP/Active Directory integration, other Atlassian tools, remote data source…	VPN as a Service is available to connect your network to the Orange Flexible Engine VPC: **https://cloud.orange-business.com/en/offers/infrastructure-iaas/public-cloud/features/vpn-as-a-service/**

Workstations security

All members from our Managed Services team are equipped with Cisco AMP, protecting us from viruses, malware and data leakage.

Separated private clouds per customer

Each customer has its own VPC (Virtual Private Cloud), segregating its network and data from other customers'.

Firewalls

Amazon Web Services	Orange Flexible Engine
Each server is connected to a **Security Group** that serves as a Firewall for inbound and outbound connections.	Each server is connected to a Security Group that serves as a Firewall for inbound and outbound connections.

The default configuration is as follows:

Inbound		Outbound	
Port	**Source**	**Port**	**Destination**
22	Only bastion servers which are accessible only from Valiantys offices IP addresses	All	Anywhere
80	Anywhere*		
443	Anywhere		

***please note that all requests sent on port 80 are rewritten to port 443 automatically with HTTPS using an Apache HTTPD rewrite rule.**
That default firewall setup can be customized depending on the customer's requirements.

Credentials

Individual user accounts

Each member of our Managed Services team has access to administration console & servers using an individual account.

All actions are logged and accessible at any time from our team. Access logs are kept for a period of 4 weeks.

Password policy

AWS & OFE Console access requires a password that contains at least 8 characters, including an upper-case letter, a lower-case letter, a number and a special character. The password must be changed every 40 days, and a similarity check is done with the 5 last passwords.

Multi-Factor Authentication

In addition to login/password protection, access to the AWS console requires a code that is generated by Google Authenticator (or Authy) mobile app associated to that user's account. Codes are generated every 30 seconds.

Server access

Server access through SSH is only possible through bastion servers from the Valiantys different locations' static IP due to security group and VPC peering configuration. Access is secured using an individual passphrase-protected private key.

Private keys are securely stored on each person's computer.

Data security

Data access

Our providers (AWS or OFE) do not have access to any of the tools that we install for our customers. They also do not have SSH access to the servers that we deploy on their infrastructure.

As data is encrypted on volumes we deploy, they cannot have access to any stored data.

Data deletion at the end of the contract

Valiantys is committed to deleting all customer data at the end of the contract.

If a customer wishes not to renew its contract, we follow this process:

- We do a full backup of the customer environment (XML backup + attachments for all Atlassian tools)

- We provide the customer with this backup using a secure method (provided by the customer or by default using an AWS S3 bucket with dedicated login/password)
- Once the customer has successfully retrieved all data, we remove this data from the repository
- We shut down the server and delete it completely
- We wait for 1 month before deleting the snapshots of the servers in case the customer would have issues importing the data on his end. This delay could be reduced upon customer's request.

Once all data has been deleted on our end, we send the customer an email stating that Valiantys does not hold customer data anymore.

SSL

Access to Atlassian tools is only possible through SSL (TLS v1.2 protocol) using our default *.**valiantys.net** certificate (2048 bits key), or certificates provided by the customer and associated to its own domain name.

Data encryption

Amazon Web Services	Orange Flexible Engine
File system can be encrypted with a unique 256-bit key, all snapshots from that server would then be encrypted using the same key. Amazon's overall key management infrastructure uses Federal Information Processing Standards (FIPS) 140-2 approved cryptographic algorithms and is consistent with National Institute of Standards and Technology (NIST) 800-57 recommendations.	File system can be encrypted with a unique 256-bit key, all snapshots from that server would then be encrypted using the same key. LUKS encryption is used to encrypt the whole partition where Atlassian tools and the database is installed.

Security logs

All sorts of logs are recorded on the servers we manage:

- SSH access logs that log all SSH accesses to the server
- Apache HTTPD access logs that log all HTTP requests that enters the server
- Applications access logs that log all HTTP requests that enters the tools

Those log files are accessible upon request.

Backups

Full VM snapshots are done automatically every night around midnight and kept for 7 days. Snapshots are encrypted.

In addition to those backups, XML backups are automatically generated every day on Atlassian tools. They allow for a quicker data retrieval in some cases (i.e. data deleted by error).

These XML backups can be disabled by our team on larger environments to avoid performance issues.

Backups retention

VM snapshots & XML exports are automatically deleted after 7 days.

Restore procedure

Restore procedure depends on the type of failure that is encountered. Our Managed Services team can choose from the two methods after analyzing your issue.

VM snapshot restore:

- We use the snapshot to create a new VM of the same size
- We redirect the server's IP to the newly created server (usually takes ~10mn)

XML export restore:

- We copy the XML export to the restore folder
- We start the XML import process (can take from a few minutes to a few hours depending on your environment size)

Restore procedure check

In order to make sure our backups and backup procedures are working as expected, we run a restore test for all our servers once per year.

A report of the results can be provided to the customer on request.

Pen tests

We run black box penetration tests at least annually on all our servers.

We run white box penetration tests at least annually on a sandbox environment which we create using the standard setup process we use for creating client environment. This helps us continuously improving our security.

We also allow customers to run their own penetration tests (by themselves or through a third party).

If a customer wants to organize penetration tests, he needs to reach out to us at least 4 weeks in advance as we need to get approval from our providers.

Vulnerability Management

Detection process
We carry out a weekly review of our reference base using specialized tools in addition to any alerts provided by Atlassian.

Detection times
The detection period is 7 days (except for an Atlassian notification during this period)

Treatment delays
Treatment is immediate. As soon as the vulnerability is known, an incident ticket is automatically created on behalf of the customer, in order to organize the remediation with his agreement and in order to impact at least the production service.

Without a response from the customer within 7 days, the vulnerability will be automatically addressed without waiting for consent on unworked slots.

Operations

Security organization

The person accountable for security in the company is currently Walé Robert, IT Manager.

We do not have a dedicated security team. All members of our Managed Services team are aware of security rules applied at Valiantys and are able to deal with security incidents.

A corporate security policy is in place and is part of our company rules and regulations that needs to be accepted by any member of Valiantys. It includes the following sections:

Information Security Policy organization

- INTRODUCTION
    - Purpose

    - Scope
    - History
    - Responsibilities
    - General Policy Definition
    - Glossary
- NETWORK ACCESS POLICY
    - Purpose
    - Scope
    - Policy Definition
- DATA, INFORMATION AND SECURITY POLICY
    - General Use and Ownership
    - Security and Proprietary Information
    - Unacceptable Use
- POLICY COMPLIANCE
    - Compliance Measurement
    - Exceptions
    - Non-Compliance
- IT ASSETS POLICY
    - Purpose
    - Scope
    - Policy Definition
- ACCESS CONTROL POLICY
    - Purpose
    - Scope
    - Policy Definition
- PASSWORD CONTROL POLICY
    - Purpose
    - Scope
    - Policy Definition
    - Statement of guidelines
- EMAIL POLICY
    - Purpose
    - Scope
    - Policy Definition
- INTERNET POLICY
    - Purpose
    - Scope
    - Policy Definition
- ANTIVIRUS POLICY
    - Purpose
    - Scope
    - Policy Definition
- INFORMATION CLASSIFICATION POLICY
    - Purpose
    - Scope
    - Policy Definition
- REMOTE ACCESS POLICY
    - Purpose
    - Scope
    - Policy Definition
- OUTSOURCING POLICY
    - Purpose

        - Scope
        - Policy Definition
        - Mandatory
    - PATCH MANAGEMENT POLICY
        - Purpose
        - Scope
        - Policy Definition
    - CLOUD APPLICATION INFORMATION POLICY
        - Purpose
        - Scope
        - Policy Definition

Before hiring new employees, we run a background check that depends on the country where they are hired:

- A basic Criminal Disclosure
- Employment History & Gap Analysis
- ID Check
- Right to Work check

Change management process

Change submission

When in a Support ticket a change is requested by a customer or identified by the Support engineer, a workflow transition allows to reach a specific status. From that status it will be possible to create a Change ticket in the dedicated Jira project via a dedicated button. The created Change ticket includes the following information:

- Description – ex: **Jira uprgade to 8.5.3 on MyCustomer PROD instance**
- Priority – ex: **Critical**
- Origin – ex: **Notification from Atlassian**
- Reason for Change – ex: **Security vulnerability on Jira version currently installed**
- Risk analysis – ex:
    - **Same parameters in setenv.sh and server.xml configuration files**
    - **Compatibility of installed apps – upgrade the impacted ones if needed**
    - **Apps requiring paid license – client must be informed as well as the account manager**
    - **No unforeseen security risk, as the data is not transferred anywhere**
- Change procedure – ex: **<link_to_the_corresponding_internal_documentation>**
- Rollback procedure – ex: **Restore of the snapshot generated before the operation**
- Estimated duration – ex: **1h**
- Link to the related Support ticket – Automatically defined by the Jira app creating the Change ticket

Change approval

After creation and depending on what type of Change is expected, the Change ticket is assigned to a Support Team Leader or a Cloud Architect. The designated assignee reviews the information provided by the reporter and decides to approve or reject the Change request.

There are some reasons to refuse a Change: technical risk, security risk, unclear or incomplete procedure, etc. In case the Change is refused, the Change ticket is closed and the workflow of the related Support ticket is automatically set back to the previous status in consequence.

If the Change is approved, a form allows to schedule the corresponding operation via the following fields:

- Start Date – i.e. date and time the Change operation will start
- End Date – i.e. date and time the Change operation will end
- Assignee – i.e. the resource in charge of the execution of that operation.

In consequence, the Change ticket is updated with those information and the workflow of the related Support ticket is automatically set to a new status indicating that a corrective operation is planned.

Change application

Once the Change has been applied, the resource assigned to the operation updates the workflow status of the Change ticket, which automatically triggers an update of the workflow of the related Support ticket so that the processing of that ticket resumes.

Emergency Change

In case of urgent issues – ex: blocking issue raised by a customer or security vulnerability raised by Atlassian impacting some application versions, an emergency Change process is applied, with the following deviations in regard of the regular process:

- A Support Team Leader or Cloud Architect handles the corresponding Support ticket, creates and approves the related Change ticket
- The Change is planned to the shortest delay possible, depending on the availability of the Support resources and the customer's agreement

Change committee

Every Monday afternoon, a 30-minute **Change committee** including at least the Head Of Managed Services, the Support Offerings and Cloud Offerings Managers, and the Support Team Leaders, review the Change tickets processed over the last period using a dedidcated Jira Dashboard.

Incident management process

Incidents are created by customers using emails, our support portal or by the Valiantys Managed Services team.

An incident can be declared as severity 1, 2, 3 or 4. Definition for our severity levels can be found on https://valiantys.com/sla.

Our Managed Services team handles tickets in the following priority order:

- Service down (incidents are automatically created from our monitoring tool) – with a 30mn recovery time objective
- Security incidents

- S1 incidents
- SLA breached tickets
- All other tickets ordered by SLA

Here is the incident workflow.

Workflow

- When the ticket is created, it arrives with “Submitted” status
- As soon as our team starts working on it, the ticket moves to “In progress” status
- If we need more information from you, the ticket is moved to “Waiting for customer” status, and a comment is added on the ticket detailing what information we need to solve the issue.
- If we need a third party to help solve the issue, the ticket is moved to “Waiting for customer” status. Third parties could be Atlassian, an add-on vendor or a hosting provider.
- When providing you with an answer, the ticket moves to “Closed” status. A comment is added providing you with all necessary information.
- If the answer does not fully meet your expectations, you have the ability to reopen it, which moves the issue to “In progress” status.
- If the ticket requires a planned operation in order to be solved (e.g. upgrade, planned restart), it will move to “Operation planned” status, letting you know in a comment when the operation will take place and who will be responsible for it.
- If the ticket is identified as a consulting task, it will be closed with resolution “Consulting”.

SLA

- Response SLA is measured from the “Submitted” status to the “Operation planned”, “Waiting for customer”, “Resolved” or “Closed” status
- Response SLA only counts the first hit of the “Operation planned”/”Waiting for customer”/”Resolved”/”Closed” status – this means that if an issue is reopened or gets back to “In progress” status, the SLA does not restart.
- Response SLA is paused on the following status: “Waiting for third party”.
- Response SLA mesured according to the coverage selected for the support contract: during business hours for clients who chose business hours coverage and 24/7 for clients who chose 24/7 coverage.

Notifications

- After ticket creation, the creator receives an email notification.
- At any moment, the creator is able to reply to this notification in order to add a comment to the ticket.
- The creator receives an email notification as soon as the ticket enters “In progress” status, meaning the issue has been taken into account.
- The creator receives an email notification when the ticket enters “Waiting for customer”, “Waiting for third party” or “Operation planned” status.
- Our support team receives a notification when the issues is reopened from “Waiting for customer” or “Closed” status.
- The creator receives an email notification when the ticket enters “Closed” status.
- The creator receives an email notification each time the ticket is commented by our support team.

Automations

- Tickets that have “Waiting for customer” status and have not been modified for 3 and 6 days are automatically commented, which sends a notification to the creator.
- Tickets with “Waiting for customer” status that have not been modified for a period of 12 days are closed automatically, but you are able to reopen them if required.

Escalation

- At any point in time, you have the ability to escalate the issue. Escalating the issue sends a notification to the Support Team lead and the Head of Managed Services so that they can review the ticket and take action if necessary.
- We recommend using this escalation button if
    - You are not getting the required level of expertise for dealing with an issue
    - You are not getting appropriate response time after you provided all information for the support agent to work on your ticket
- If you have a very urgent issue, we recommend you to call us directly rather than escalating the issue (as a reminder, our phone numbers can be found here: **Raise a support request**)

Periodic access rights review

Every month, a check is done on all applications used at the managed services level to check that access is granted only to authorized personnel.

A report is created and shared with the Managed services team manager. In case of erroneous access rights, immediate action is taken for fixing the issue by creating a security incident if necessary.

Security incident management

Our security incident process is covered through our incident response plan.

STEP 1 – Discovery & Categorization
Anyone who discovers the incident will contact Valiantys Support team by creating a Security Incident in the dedicated Jira Service Desk instance:

- Name of caller or source of incident alert (AWS or Software notifications) – ex: **Email from Atlassian or ticket created by a customer on Valiantys Support portal**
- Time of first report
- Nature of the incident – ex: **Vulnerability on an Atlassian product**
- What system(s) or persons were involved? – ex: **impacted Atlassian applications (i.e. versions affected by the vulnerablilty)**

- Location of equipment or persons involved – ex: **Valiantys Cloud infrastructure**
- How incident was detected – ex: **Pen tests executed by a customer on their applications**

All logs information from the server must be copied to the ticket for future investigation.

STEP 2 – Escalation & Exploration

Depending on the type of the Security Incident, the created ticket is assigned to the Support or Cloud Offerings Manager. The designated assignee:

- Reviews the contents of the ticket
- Assess the list of impacted environments. Those environments must be backed up (unless there is already a recent backup available) in order to protect the information for identification, collection, acquisition and preservation, which can serve as evidence
- Based on that assessment, the assignee decides the type of remediation to be applied so that the Incident is solved as quickly and efficiently as possible. Depending on the cases, the type of remediation will be:
    -
        - Deployment of a workaround, if any
        - System or applicative upgrade – ex: OS, Apache, PostgreSQL, Atlassian tool or app
        - Migration to a new environment

For each environment impacted by the Incident, and based on the **Change Management** process, Change tickets are created in the dedicated Jira project, with the following information:

Change Management process

- Description – ex: **Jira uprgade to 8.5.3 on MyCustomer PROD instance**
- Priority – ex: **Critical**
- Origin – ex: **Notification from Atlassian**
- Reason for Change – ex: **Security vulnerability on Jira version currently installed**
- Risk analysis – ex:
    - **Same parameters in setenv.sh and server.xml configuration files**
    - **Compatibility of installed apps – upgrade the impacted ones if needed**
    - **Apps requiring paid license – client must be informed as well as the account manager**
    - **No unforeseen security risk, as the data is not transferred anywhere**
- Change procedure – ex: **<link_to_the_corresponding_internal_documentation>**
- Rollback procedure – ex: **Restore of the snapshot generated before the operation**
- Estimated duration – ex: **1h**
- Link to the related Support ticket – Automatically defined by the Jira app creating the Change ticket

Security breach definition

A security breach is defined as unauthorized acquisition of data that compromises the security, confidentiality or integrity of personal information maintained by Valiantys.

Good faith acquisition of personal information by an employee or agent of our company for business purposes is not a breach, provided that the personal information is not used or subject to further unauthorized disclosure.

Incident alert caller

Following are some examples of values to be indicated in the **Origin** field of the Change ticket:

- Customer
- Atlassian consultant for a customer
- Third party such as CVE database
- Valiantys
- Atlassian
- Unknown

Nature

Following are some examples of values to be indicated in the **Reason for Change** field of the Change ticket:

- Breach of personal information
- Denial of service/Distributed denial of service
- Excessive port scan
- Firewall breach
- Virus outbreak
- Exploitation of server resources

Classification / Criticality

Following are some examples of values to be indicated in the **Priority** field of the Change ticket:

CRITICAL

- Definition: Incidents that have a critical impact on the firm's business or the service to Valiantys MS customers
- Example: Unauthorized system access

MEDIUM

- Definition: Incidents that has a significant or has the potential to have a critical impact on the firm's business or the service to Valiantys MS customers
- Example: Password cracking attempt

LOW

- Definitions: Incidents that has the potential to have a significant or critical impact on the firm's business or the service to Valiantys MS customers
- Example: Firewall scanning

STEP 3 – Remediation

Each concerned System is scanned for the open vulnerability before application of the remediation.

The Security incident is fixed via the application of the Change procedure – provided in a dedicated field of the Change ticket.

After validation of the corrective action, the concerned Systems are scanned again to verify that the vulnerability has been eliminated.

STEP 4 – Communication

- Valiantys MS is committed to ensure correct communication to customers or partners that could be impacted with the Security issue within a delay of maximum 24 hours, by sending an email (or creating a support ticket on behalf of the client) to all the concerned customers including the following information:
    - Detail of the vulnerability (CVE number if available)
    - How the vulnerability impacts their environment
    - Description of the remediation plan
    - Description of the communication plan
    - Details of the investigation if the vulnerability was exploited already
- Valiantys Managed Services Management must be involved in the process as early as possible from the escalation step, and depending on the severity and the scope of the impacted environments:
    - Valiantys Head of managed services must be involved for any security incident of Critical criticality
    - Valiantys CTO must be involved if the scope of the Security Incident is Valiantys Cloud infrastructure and the criticality Critical.
    - Valiantys CISO must be involved if the scope of the Security Incident is Valiantys infrastructure and the criticality Critical.
    - Valiantys CEO must be involved if the Security Incident commits Valiantys from a legal or financial point of view and the criticality Critical.

STEP 5 – Post-mortem review of the response and update policies

In order to take preventive steps so the incident doesn't happen again, we asked ourselves several questions:

- Would an additional policy have prevented the incident?
- Has a procedure or policy not been followed which allowed the incident? What could be changed to ensure that the procedure or policy is followed in the future?
- Was the incident response appropriate? How could it be improved?
- Was every appropriate party informed in a timely manner?
- Were the incident-response procedures detailed and did they cover the entire situation? How can they be improved?

- Have changes been made to prevent another incident? Have all systems been patched, systems locked down, passwords changed, anti-virus updated, email policies set, etc.?
- Should any security policies be updated?
- What lessons have been learned from this experience?

This discussion can take place in the weekly **Change committee** including at least the Head Of Managed Services, the Support Offerings and Cloud Offerings Managers, and the Support Team Leaders.

Periodic Testing

Valiantys Managed Services team has the responsibility to test and review the Security Incident Response Plan on a quarterly basis, at the following periods:

- Mid-March
- Mid-June
- Mid-September
- Mid-December

Disaster recovery

All our virtual servers (except Valiantys Cloud Starter offerings) are backed up every night. The backup is stored in the same region than the servers. For disaster recovery purposes, all snapshots are automatically copied to another region.

In case of an important failure on the main region, we are able to restart the service from the latest backup very quickly, limiting impact on the customers as much as possible:

Our recovery time objective and recovery point objectives are available at https://valiantys.com/sla.

We test this process once every year on all our servers in order to be ready and efficient when the time comes.