

Version: 1.0

Alibaba Cloud Security Whitepaper - International Edition

April. 2018

Contents

1. Introduction	4
2. Shared Security Responsibilities	5
2.1 Security Responsibilities of Alibaba Cloud	6
2.2 Security Responsibilities of Customers	7
3. Security Compliance and Privacy	8
3.1 Compliance	11
3.2 Privacy & Data Protection	13
3.3 Transparency	14
4. Alibaba Cloud Infrastructure	15
5. Alibaba Cloud Security Architecture	16
5.1 Cloud Platform Security Architecture	17
5.1.1 Physical Security	17
5.1.2 Hardware Security	20
5.1.3 Virtualization Security	20
5.1.4 Cloud Product Security	23
5.2 User-Side Security Architecture	25
5.2.1 Account Security	25
5.2.2 VM Security	28
5.2.3 Application Security	29
5.2.4 Network Security	29
5.2.5 Data Security	31
5.2.6 Operation Security	33
6 Cloud Product Security	34
6.1 Elastic Computing	34
6.1.1 Elastic Compute Service (ECS)	34
6.1.2 Auto Scaling	42
6.1.3 Resource Orchestration Service	43
6.2 Networking	43
6.2.1 Server Load Balancer (SLB)	43
6.2.2 Virtual Private Cloud (VPC)	45
6.3 Database	49
6.3.1 ApsaraDB for RDS	49
6.3.2 ApsaraDB for Redis	53
6.3.3 ApsaraDB for Memcache	54
6.4 Storage and CDN	56
6.4.1 Object Storage Service (OSS)	56
6.4.2 Table Store	60
6.4.3 Network Attached Storage	61
6.4.4 Alibaba Cloud CDN	62
6.5 Analytic and Big Data	64

6.5.1 MaxCompute	64
6.6 Application Service	66
6.6.1 Log Service	66
6.7 Management and Monitoring	66
6.7.1 Identity and Access Management (RAM&STS)	66
6.7.2 Key Management Service	77
6.7.3 ActionTrail	79
6.7.4 CloudMonitor	81
7. Alibaba Cloud Security	82
7.1 Basic Protection	82
7.1.1 Anti-DDoS Basic	82
7.2 Advanced Protection	83
7.2.1 Anti-DDoS Pro	83
7.2.2 Mobile Security	84
7.2.3 Web Application Firewall (WAF)	86
7.2.4 Server Guard	88
7.2.5 Alibaba Cloud SSL Certificates Service	89
7.2.6 Managed Security Service	90
8. Alibaba Cloud Security Ecosystem	93
9. Version History	93

1. Introduction

Data security and user privacy are the top most priorities of Alibaba Cloud. Alibaba Cloud strives to provide customers with consistent, reliable, secure, and regulation-compliant cloud computing services, helping customers ensure the availability, confidentiality, and integrity of their systems and data.

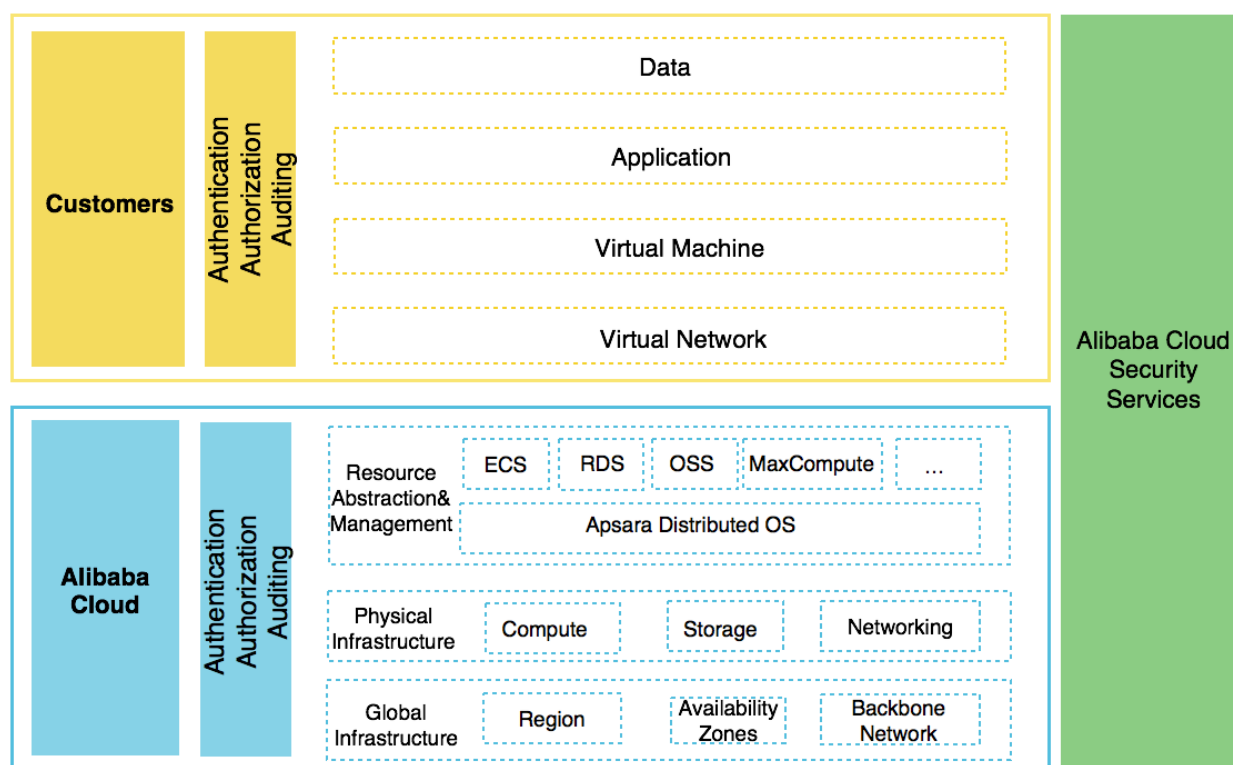
This white paper discusses various security aspects of Alibaba Cloud in the following parts:

- Shared security responsibilities
- Security compliance and privacy
- Alibaba Cloud infrastructure
- Alibaba Cloud security architecture
- Security features of Alibaba Cloud products
- Security services provided by Alibaba Cloud
- Alibaba Cloud security ecosystem

This white paper also provides the best practices about how to securely use Alibaba Cloud products and security services, thus allowing customers to make better uses of the Alibaba Cloud platform and to get insights into the overall cloud security environment.

2. Shared Security Responsibilities

Alibaba Cloud and its customers are jointly responsible for the security of customers' applications built on Alibaba Cloud. Alibaba Cloud is responsible for the security of the underlying cloud service platform and infrastructure, and customers are responsible for the security of applications built on top of or connected to the cloud. The shared security responsibility model is somewhat different than the typical security model a customer would see in an on-premises data center. Customers are able to leverage the underlying security assurance and capabilities that Alibaba Cloud provides, thus getting an overall better security return by using Alibaba Cloud.



Alibaba Cloud must ensure a securely managed and operated infrastructure (including but not limited to data centers deployed across regions and zones, and Alibaba backbone networks), physical devices (including computing, storage, and network devices), distributed cloud OS named Apsara, and various cloud services and products running on top of the Apsara OS.

By leveraging its years of expertise in attack prevention technologies, Alibaba Cloud offers various security features and services to help protect customers' applications and systems. In turn, customers must, in a secure manner, configure and use cloud products (such as the Elastic Compute Service (ECS), Relational Database Service (RDS) instances, etc.) , and build applications based on such securely configured cloud products. Customers can choose to use the Alibaba Cloud security services or any third-party security products in the Alibaba Cloud security ecosystem to protect their applications and assets.

With security responsibilities shared between Alibaba Cloud and its customers, Alibaba Cloud provides a secure infrastructure to help mitigate the security needs of customers, thus relieving much of the underlying security burdens while allowing customers to focus more on their core business needs.

2.1 Security Responsibilities of Alibaba Cloud

Alibaba Cloud is responsible for the security of its infrastructure, physical devices, Apsara OS, and cloud services/products, and provides customers with the technical means necessary to protect their cloud applications and data.

Alibaba Cloud ensures the cloud platform security by:

- Protecting the physical security of cloud data centers;
- Protecting the security of hardware, software, and network of the cloud platform by means of OS- and database-patch management, network access control, Anti-DDoS, and disaster recovery, etc.;
- Identifying and fixing security vulnerabilities of the cloud platform in a timely manner without affecting customers' service availability;
- Cooperating with independent third-party security regulation and audit agencies to audit and evaluate security and compliance of Alibaba Cloud.

Alibaba Cloud provides customers with the following technical security measures:

- Providing multihomed BGP access networks and cloud data centers distributed across multiple regions and zones, thus enable customers to build high availability cloud applications;
- Providing Alibaba Cloud account authentication and authorization that support two level account credentials (Alibaba Cloud account and individual RAM user accounts)

for easy segregation of duties, multi-factor authentication (MFA), grouped authorization policies, fine-grained authorization control, and temporary authorization token;

- Providing security audit support;
- Providing data encryption support;
- Providing various Alibaba Cloud security services (in-house & third party);
- Introducing third-party security vendors to offer customers security solutions tailored for their needs.

2.2 Security Responsibilities of Customers

Customers who build cloud applications based on Alibaba Cloud services is responsible for protecting their own systems by using the security features of Alibaba Cloud products, Alibaba Cloud Security services, and the third-party security products provided by the Alibaba Cloud security ecosystem.

Customers must protect their Alibaba Cloud account credentials by allocating an independent RAM (Resource Access Management) user account for each maintenance personnel, granting only the minimum permissions required, and ensuring a separation of duties by means of assigning authorization by groups. We recommend that the customers enable the multi-factor authentication (MFA) for their accounts. Furthermore, customers could use the Alibaba Cloud ActionTrail to record OpenAPI call logs and operations performed on the management and control console, and use encryption at rest and in motion capabilities in various Alibaba Cloud products to protect sensitive data.

Customers have full control of the ECS and Virtual Private Cloud (VPC) instances provided by Alibaba Cloud, and are responsible for managing these instances and performing the necessary security configurations. For example, customers could perform security hardening to their ECS Operating Systems, install security patches in a timely fashion, and configure firewalls (security groups) for network access control enforcement.

For other Alibaba Cloud services, such as RDS and MaxCompute, customers do not need to maintain the underlying computing instances, such as keeping the OS and database updated, hardened, and patched. Instead, customers are only responsible for managing the service account credentials and resource authorization, and using the build-in security features, such as configuring a source IP address whitelist for RDS, etc.

3. Security Compliance and Privacy

Alibaba Cloud adheres to domestic and international information security standards, as well as industry requirements. We integrate compliance requirements and standards into our internal control framework, and implement such requirements and standards by design in our cloud platform and products. Alibaba Cloud is involved in the development of multiple standards for the cloud industry and contributed to the best practices. We also engage with independent third parties to verify the compliance of Alibaba Cloud according to various requirements. Certified by more than 10 agencies across the globe, Alibaba Cloud is a cloud service provider with the most complete range of certifications in Asia. Our certifications and compliance credentials are shown in the following:

Certification	Description
ISO 27001	An international certification for Information Security Management System (ISMS). Alibaba Cloud is certified in such aspects as data security, network security, communication security, and operational security according to this standard.
CSA STAR	An international certification for cloud security. Alibaba Cloud won the world's first gold medal for CSA STAR certification.
ISO 20000	A standard for IT service management system. This certifies that Alibaba Cloud has established and strictly implemented a standard service process. Complied cloud services can improve IT efficiency and reduce the overall IT risk.
ISO 22301	A standard for Business Continuity Management (BCM) system. This certifies Alibaba Cloud at meeting the requirements for business continuity planning, disaster recovery and regular drills to enhance the stability of cloud platform.
The Multi-Level Protection Scheme (MLPS)	The Multi-Level Protection Scheme (MLPS) tiered protection system is a basic information security system in China.

	<p>Alibaba Cloud has obtained a Level III certification, which means our systems ensure the security and recovery capabilities of level three information in response to security threats.</p> <p>Alibaba Cloud's Finance Cloud is China's first cloud platform that passed the Level 4 certification of The Multi-Level Protection Scheme (MLPS), a requirement for critical infrastructures in China.</p>
Cloud service network security audit by Cyberspace Affairs Leading Group	Alibaba Cloud is the only cloud service provider that passed an enhanced auditing level (more than 500 checkpoints) among the first group of cloud service providers that passed the security audit of Cyberspace Affairs Leading Group.
Cloud Service Capability Standard Test by the Ministry of Industry	The only cloud service capability certification test in China that is based on national standards for public clouds and private clouds.
CNAS	China National Accreditation Service for Conformity Assessment (CNAS) certification for cloud products is the only product-level classified certification based on national standards in China.
PCI DSS (Payment Card Industry Data Security Standard)	PCI DSS focuses on the management and control of payment card information within the organization throughout its life cycle, from generation/entry, transfer, storage, processing, and to destruction. Alibaba Cloud is dedicated to payment security and is strictly compliant with PCI Data Security Standards.
MTCS T3	The Multi-Tier Cloud Security (MTCS) level T3 is Singapore's highest-level certification for security of cloud service providers. This qualifies Alibaba Cloud to be used in Singapore public sector's projects.
Service Organization Control (SOC) audit	Service Organization Control (SOC) Reports are internal control reports on the services provided by a service organization. The reports provide valuable information that users need to assess and address the risks associated with

	an outsourced service. Alibaba Cloud has obtained the SOC1, SOC2 and SOC3 Type I, Type II reports.
German Cloud Computing Compliance Controls Catalog (C5)	Alibaba Cloud is the first cloud provider that attested the additional requirements of the C5, which is recognized in financial sector and mandatory for public sector customers.
TRUSTe	Alibaba Cloud is certified by TRUSTe Enterprise Privacy Certification. This marks the compliance of Alibaba Cloud in collecting and processing personal information, and meeting the standard for responsible data collection practices.
HIPAA	Alibaba Cloud complies with the US Health Insurance Portability and Accountability Act (HIPAA) to protect the security of personal health information (PHI), and support covered entity with Business Associate Agreement (BAA) to meet customer needs.
MPAA	Alibaba Cloud complies with the best practices and common guidelines of the Motion Picture Association of America (MPAA).
PDPA	Alibaba Cloud complies with the requirements of Singapore's Personal Data Protection Act (PDPA).
Trusted Cloud Label Certification	Alibaba Cloud is a member of Trusted Cloud promoted by German Federal Ministry of Economics and Energy, and also certified and obtained their quality label of reliable cloud services.
Founding member of EU Cloud Code of Conduct	As a founding member and a member of the General Assembly, Alibaba Cloud is actively engaged in creating the EU Cloud Code of Conduct. Alibaba Cloud supports improved cloud computing industry transparency and helping cloud customers understand how data protection issues are addressed by cloud service providers.
Proposer of "Data Protection Initiative"	This is the first "Data Protection Initiative" that Alibaba Cloud launched for China's cloud service providers, promoting data protection and cloud service providers' responsibilities and obligations in protecting customers' data.

3.1 Compliance

Alibaba Cloud is constantly improving by following standards and industry best practices, and has adopted a series of certifications. At the same time, we undergo third-party audits and self-assessments, in order to better demonstrate our compliance practices to the customers.

Small and medium-sized enterprises (SMEs) have their limitations in experience and resources to meet the compliance requirements. Alibaba Cloud hopes to assist SMEs in their efforts to compliance in many ways, such as providing audit reports, compliance solutions, and compliance architecture consulting, in order to maximize the value of Alibaba Cloud's compliance practices.

Given the compliance requirements are different in contexts, industries and regions, Alibaba Cloud's overall compliance framework is divided as follows:

1) **Management system compliance:** To demonstrate Alibaba Cloud's mature management mechanism and industry's best practices it complies with.

- ISO 27001: Information Security Management System
- ISO 20000: IT Service Management System
- ISO 22301: Business Continuity Management System
- CSA STAR: Maturity Model of Cloud Service Security
- The Multi-Level Protection Scheme (MLPS): level 4
- CNAS: Test for Cloud Computing Standards

2) **Systematized compliance reports:** To demonstrate Alibaba Cloud's integrity and effectiveness of management and controls. For example, the effectiveness of system control, the accuracy of separation of duties, and the maintenance audits.

- PCI-DSS: Payment Card Industry Data Security Standard
- MPAA: Best Practices – Common Guidelines of Motion Picture Association of America (MPAA)
- TRUSTe: TRUSTe Enterprise Privacy Certification
- SOC 1/2/3 TYPE II: The Service Organization Control (SOC) reports are a series of audit reports from independent third-party auditors to indicate the effectiveness of Alibaba Cloud's control objectives and activities. These reports are designed to help customers

and their auditors to get a picture of the control measures behind operation and compliance. Alibaba Cloud SOC reports are categorized into three types:

- SOC 1 TYPE II: Internal control report on financial reporting
- SOC 2 TYPE II: Report on security, availability, and confidentiality
- SOC 3: Report on security, availability, and confidentiality
- German Cloud Computing Compliance Controls Catalog (C5): The cloud requirements catalog in Germany for assessing the information security of cloud services, which defines a baseline for cloud security. Proofs of a cloud provider meets the requirements of the catalog is provided by an SOC 2 report. This is based on the internationally recognized testing regime of ISAE 3000 used by accountants.
- Trusted Cloud label: Issued by the Trusted Cloud Competence Network, it is awarded to trustworthy cloud services which meet the minimum requirements with regard to transparency, security, quality and legal compliance.

3) **Legal compliance:** Compliance with local laws and regulations is the primary condition for cloud services to be implemented in different regions. However, legal compliance cannot be reflected in the form of certificates or audit reports due to its unique nature.

- HIPAA – Alibaba Cloud complies with US Health Insurance Portability and Accountability Act (HIPAA) to protect the security of personal health information (PHI), and support covered entity with Business Associate Agreement (BAA) to meet customer needs.
- GDPR – Alibaba Cloud strives to provide supports for customers and partners while committed to the EU data protection regulations.

4) **Others:** Some compliance efforts cannot be demonstrated in preceding three forms. Alibaba Cloud has been assisting regulators in different regions in establishing and improving standards by sharing our best practices.

- MTCS – Multi-Tier Cloud Security is the cloud security standard proposed by the Info-communications Development Authority of Singapore and released by Singapore Standards, Productivity and Innovation Board. MTCS security certification has three levels. Alibaba Cloud obtained Level-3 certification - the highest security level.

3.2 Privacy & Data Protection

Alibaba Cloud personal information handling principle: customers have ownership and control of all the personal information provided to Alibaba Cloud.

Alibaba Cloud is committed to protect customers' personal information and guarantees that such information is only used for the purposes agreed by customers. Alibaba Cloud's privacy policy is completely transparent to the public, and can be found on our official website. At the same time, Alibaba Cloud takes various technical measures to ensure that the customers' personal information is well protected.

Please see Alibaba Cloud International Website Privacy Policy at:

<https://www.alibabacloud.com/help/faq-detail/42425.htm>

Alibaba Cloud provides comprehensive compliance information in our online Security and Compliance Center to allow customers to better understand our compliance practices. The goal is to provide guidance for customers to improve their compliance capability together with us. The upcoming GDPR requirements enhanced the existing individual rights by creating a new Right to Data Portability. To help the cloud industry and the customers better understand the GDPR requirements, especially on how to translate the regulations to the actual implementations in operation, Alibaba Cloud is partnered with the Carnegie Mellon University for a research project that supporting data portability in the cloud. Furthermore, Alibaba Cloud also works with our partner, TrustArc, to provide the privacy compliance services for customers.

To download our research paper on Supporting Data Portability in the Cloud Under the GDPR:

<https://www.alibabacloud.com/trust-center/gdpr#resources>

To connect with our privacy partner TrustArc:

<https://www.alibabacloud.com/trust-center/gdpr#trustarc>

Again, we would like to assure customers that Alibaba Cloud is committed to protecting the personal information of customers worldwide, and to complying with applicable laws of countries where our business is operated.

3.3 Transparency

Similar to other large internet companies around the world, at times Alibaba Cloud is required by law to provide records to government authorities during an investigation or in the course of litigation. Such scenarios include cooperating with legal authorities to collect evidence for investigations into criminal cases, assisting regulators with security checks according to legal requirements, and assisting relevant authorities in anti-counterfeiting and anti-piracy activities.

We have established procedures to support litigation, court order, discovery and other legal matters that may require disclosure of data. Each request is carefully reviewed and analyzed by our legal team to ensure the validity of the request with consideration on means to minimize disclosure of personal data.

4. Alibaba Cloud Infrastructure

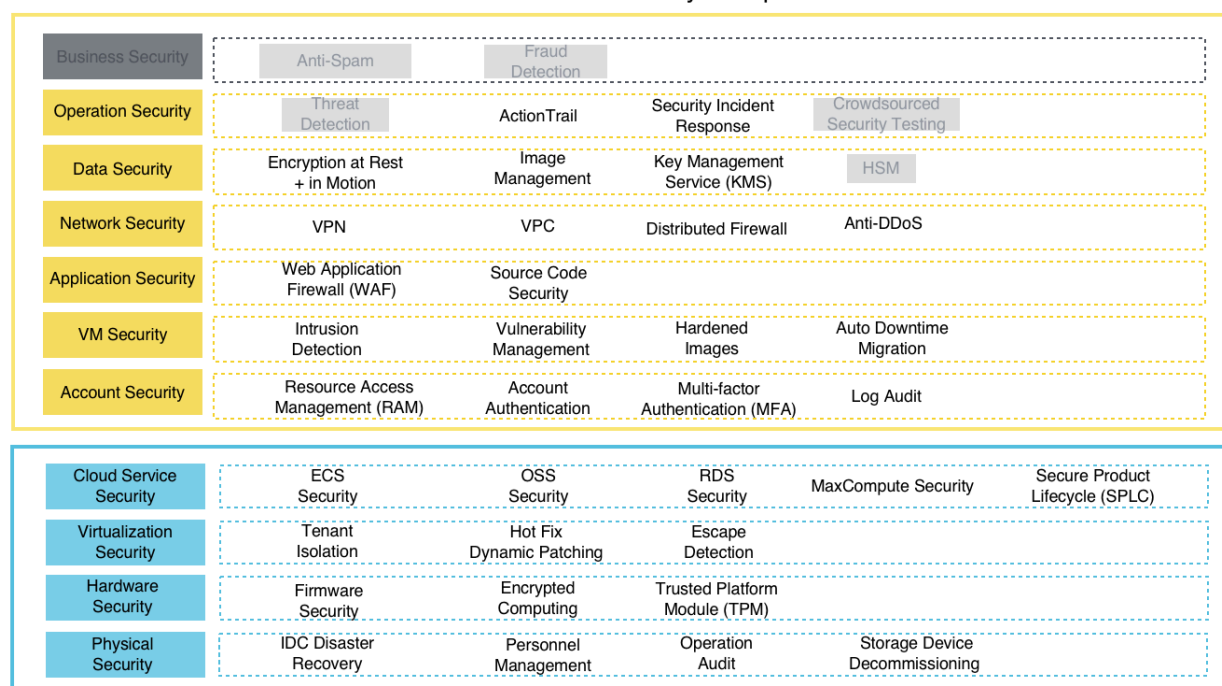
Alibaba Cloud offers high availability, secure, and reliable cloud computing infrastructure by: setting up cloud data centers across multiple regions and zones globally; delivering a better network access experience with multihomed BGP networks; providing cloud products with high availability infrastructure and multi-replica data redundancy based on the Apsara distributed cloud OS; upgrading products and fixing vulnerabilities using the hotfix dynamic patching technology; and ensuring operation security while achieving national-leading compliance.

Alibaba Cloud's data centers are deployed across multiple regions worldwide, with each region supporting multiple zones. Customer businesses can be deployed across regions and zones to implement a high availability architecture, such as same-city active-active architecture, remote data recovery, remote multi-active architecture, and geo-redundant disaster recovery architecture featuring same-city recovery and an additional remote recovery.

Country	Region	Number of zones
China	China North 1	2
	China North 2	6
	China North 3	2
	China North 5	1
	China East 1	7
	China East 2	5
	China South 1	4
	Hong Kong	2
Overseas	Asia Pacific SE 1	2
	Asia Pacific SE 2	1
	Asia Pacific SE 3	1
	Asia Pacific SE 5	1
	Asia Pacific NE 1	1
	Asia Pacific SOU 1	1
	US East 1	2
	US West 1	2
	EU Central 1	2
	Middle East 1	1

5. Alibaba Cloud Security Architecture

Alibaba Cloud Security Compass



As shown in the figure above, Alibaba Cloud provides an 11-layer security architecture, including four layers oriented at cloud platform (physical security, hardware security, virtualization security, and cloud service security), and seven layers oriented at cloud users (account security, VM security, application security, network security, data security, operation security, and business security). **Please note that there are several items being greyed out in the above diagram. The greyed-out items are the ones currently only available in Mainland China, though they would be made available globally in the future.** The rest of this whitepaper will not cover the greyed-out items in detail.

This chapter gives an overview of the overall security architecture and describes the key features of each architecture layer by briefly covering various Alibaba Cloud products. For details of various products, please see the relevant sections in Chapter 6 and 7 of the white paper.

5.1 Cloud Platform Security Architecture

Cloud Service Security	ECS Security	OSS Security	RDS Security	MaxCompute Security	Secure Product Lifecycle (SPLC)
Virtualization Security	Tenant Isolation	Hot Fix Dynamic Patching	Escape Detection		
Hardware Security	Firmware Security	Encrypted Computing	Trusted Platform Module (TPM)		
Physical Security	IDC Disaster Recovery	Personnel Management	Operation Audit	Storage Device Decommissioning	

As shown in the figure above, Alibaba Cloud platform security architecture includes four important layers: physical security, hardware security, virtualization security, and cloud service security.

5.1.1 Physical Security

All of the Alibaba Cloud data center and office areas are configured with access control, with visitor areas marked out separately. Visitors are required to carry entry pass and be escorted by Alibaba Cloud staff when visiting Alibaba Cloud premises. Alibaba Cloud's data centers are all in compliance with the requirements for Class A in the GB 50174 Code for Design of Electronic Information System Room and the T3+ standards in the TIA-942 Telecommunications Infrastructure Standard for Data Centers, including the following requirements for physical and environmental security control:

5.1.1.1 IDC Disaster Recovery

- **Fire detection and handling**

Alibaba Cloud data centers are equipped with fire detection systems using thermal and smoke sensors. The sensors, fitted to the ceiling and floor, would give audible and visual alarms when triggered. Each data center comes with an integrated gas extinguishing system and fire extinguishers. Trainings and drills on how to detect and respond to fires are organized regularly.

- **Power**

To achieve a 24/7 uninterrupted service, Alibaba Cloud data centers are powered by dual main supplies and redundant power systems. In case of a power failure, redundant battery packs and diesel generators are enabled to power data center devices, thus allowing the data center to run continuously for a certain period of time.

- **Temperature and humidity**

Alibaba Cloud data centers are fitted with precision air conditioners to ensure a constant temperature and humidity level, which are electronically monitored. In case of any fluctuation of temperature or humidity outside of the normal range, an alarm is triggered and actions are taken immediately. All air conditioning units work in hot standby mode.

5.1.1.2 Personnel Management

- **Access management**

At each Alibaba Cloud data center, the long-term access permissions are assigned only to the corresponding maintenance personnel. If an employee is transferred to another position or leaves the company, his/her access permissions are cleared immediately. If it is necessary for any other person to enter the data center, he/she must submit an application in advance, and is granted the temporary permission only upon the approval of the corresponding department heads. For each entry to or exit from the data center, such person must display ID to check in, and be under the escort of the data center's maintenance personnel for the entire duration of the visit.

An Alibaba Cloud data center consists of equipment rooms, electrical measurement areas, warehouses, and other areas, with each area equipped with an independent access control system. Two-factor authentication (such as using fingerprint) is employed for sensitive areas, and special areas are physically isolated by metal cages.

- **Account management and identity authentication**

Alibaba Cloud manages employee account lifecycle using a central account management and identity authentication system:

- Each employee is assigned a unique account;
- Password policy is set in place which requires employees to set a password meeting the length and complexity requirements, and to change the password regularly;
- Multiple logon authentication modes are supported, such as account password logon, one-time password logon, and digital certificate logon, etc.

- **Authorization management**

Alibaba Cloud grants minimum resource access permissions to each employee based on his/her position and role while ensuring separation of duties. An employee can log on to the central permission management platform to apply for access permissions to VPNs,

bastion hosts, control platforms, and production systems as needed. The requested permissions are granted to the employee upon the approvals of the supervisor, data or system owner, security administrator, and relevant departments.

- **Separation of duties**

Alibaba Cloud separates duties between maintenance permissions by role to prevent permission violations and audit failures. For example, duties are separated between maintenance and audit staff, and duties are separated between the database and system administrators, etc.

5.1.1.3 Operation Audit

- **Surveillance**

Alibaba Cloud data centers and server rooms are equipped with security surveillance systems covering all the areas and passages, and staffed with security guards for 24/7 patrol. All the surveillance videos and documents are saved and reviewed by dedicated personnel periodically.

- **Audit**

All the maintenance operations on production system can only be performed with bastion hosts. The entire operation process is recorded in logs, which are then transferred to a central log platform in real time. Alibaba Cloud defines audit rules for violations in accordance with its Account Usage Specification and Data Security Specification. Any violations would be handled by security personnel accordingly.

Internally, all the sensitive operations are logged in the management system that has a browser/server (B/S) structure, as stated in Alibaba Cloud's log audit specification, and such logs are transferred to the central log platform. The central log platform provides only the APIs for collecting and reading, not for modifying and deleting logs.

5.1.1.4 Storage Device Recycling/Decommissioning

Alibaba Cloud has established a security management system for the full life cycle of devices, from reception, storage, placement, maintenance, transfer, reuse, and to decommission. Access control and operation monitoring of devices are managed strictly, and maintenance and stocktaking of devices are conducted on a regular basis. In particular, when any device is recycled or decommissioned, Alibaba Cloud takes data erasure measures for the storage media, such as overwriting, degaussing, and physically bending.

Alibaba Cloud uses data erasure techniques that meet industry standards. The decommissioning operations are logged to prevent unauthorized access to customer data.

5.1.2 Hardware Security

5.1.2.1 Firmware Security

Secure firmware is one of the foundations for the overall cloud computing security. The firmware used within the Alibaba Cloud infrastructure is security hardened. Such hardening techniques include firmware baseline scanning, high-performance GPU instance protection, BIOS secure update, and BMC firmware protection.

- Firmware baseline scanning: The version and other related information of firmware are scanned regularly for any potential exception.
- High-performance GPU instance protection: Provide protection to critical GPU registers to ensure that the GPU flash cannot be modified by the users' virtual machines, thus making sure sensitive assets such as the GPU's firmware are not modified.
- BIOS secure update: Ensure that only the BIOS images signed by Alibaba Cloud are flashed to the relevant servers to avoid malicious BIOS flashing.
- BMC firmware protection: Prevent unauthorized BMC firmware flashing in the host operating system.

5.1.2.2 Encrypted Computing

Alibaba Cloud platform provides chip-level trusted execution environment. Users can establish a trusted execution environment to protect their sensitive data and encryption/decryption keys. With commercially available virtual machines that support encrypted computing, users can write program to ensure that the critical data is only accessed and operated by user authorized code. With the encrypted computing technology, Alibaba Cloud offers a more powerful data security solution to allow users to migrate data to cloud with high security.

5.1.3 Virtualization Security

Virtualization technology lays the foundation for cloud computing, and ensures isolation between multiple tenants in a cloud computing environment by means of virtualized

computing, storage, and network. Alibaba Cloud virtualization security technology involves three basic security features – instance isolation, hot fix dynamic patching, and escape detection – to ensure the security of Alibaba Cloud virtualization layer.

5.1.3.1 Tenant Isolation

VMM allows virtual machines at multiple computing nodes to be isolated from each other at the system level, preventing unauthorized access to system resources between tenants and thus guaranteeing the basic computing isolation between computing nodes. Virtualization management layer also provides storage isolation and network isolation. For details about tenant isolation, see "Cloud Product Security – Elastic Computing – ECS – Tenant Isolation" section.

- **Computing isolation**

Alibaba Cloud provides a variety of cloud-based computing instances and services that allow automatic scaling to meet application or business needs. These computing instances and services provide computing isolation at multiple levels to protect data, while ensuring configuration flexibility of users. Such isolation is provided by the hypervisor. Alibaba Cloud platform uses a virtualized environment where user instances run as standalone virtual machines and the isolation is enforced by using different processor ring levels to avoid unauthorized access of a user's virtual machine to the host and to another virtual machine.

- **Storage isolation**

In the basic design of cloud computing virtualization, Alibaba Cloud separates virtual machine-based computing from storage. This separation allows computing and storage to be scaled independently, and makes it easier to provide multi-tenant services. All the I/O operations of a virtual machine are intercepted by the hypervisor to ensure that the virtual machine can only access the physical disk space allocated to it, thus realizing the security isolation of hard disk space between different virtual machines. After an ECS instance is released, the original disk space and memory space are reliably scrubbed to ensure user data security.

- **Network isolation**

To provide network connections for ECS instances, Alibaba Cloud connects virtual machines to the Alibaba cloud virtual network. Alibaba Cloud virtual network is a logical structure built on top of the physical network structure. All the logical virtual networks are

isolated from each other. This isolation prevents the network traffic data from being snooped and/or intercepted by other malicious instances.

5.1.3.2 VM Escape Detection

VM Escape generally involves two steps: first, placing the virtual machine controlled by the attacker on the same physical host as the targeted victim virtual machine. Next, escape the isolation boundary to intercept any sensitive information from the victim or performing operations that compromise the functionality of the victim instance.

Alibaba Cloud's VMM uses advanced virtual machine distribution algorithm to prevent a malicious virtual machine from running on a specific physical machine. At the hypervisor level, Alibaba Cloud also provides three core technologies – hypervisor security hardening, attack detection, and hotfix – to mitigate the attacks from malicious virtual machines.

5.1.3.3 Hotfix Dynamic Patching

Alibaba Cloud virtualization platform supports hotfix dynamic patching technology, which can fix system defects or vulnerabilities without user intervention.

5.1.3.4 Change Management

Virtualization system is the foundation for cloud computing. Any changes to the virtualization system can directly affect the cloud operation. Alibaba Cloud has, according to ISO/IEC 20000, established a complete change management process, where changes are classified based on the degree of emergency, and are managed by categories based on their sources and targets. The criteria for judging possible outcomes from various changes are also clearly defined. The whole change process is standardized and is supported by automatic systems and tools. Any changes need to go through a series of stages from application, evaluation, approval, test, implementation, and to verify. The responsibilities of various personnel involved in the process are clearly defined. The changes are also documented in accordance with industry standards.

- Application phase of change: including application submission, documentation, reception, and approval are clearly defined.
- Implementation phase of change: including the change's scheme, plan, assessment, and implementation. All the changes are tested before being implemented. The change time window and change scheme are subject to strict review. In addition, Alibaba Cloud

will send a change notice to customers who may be affected by such change. Important change operations must be reviewed/confirmed by two persons.

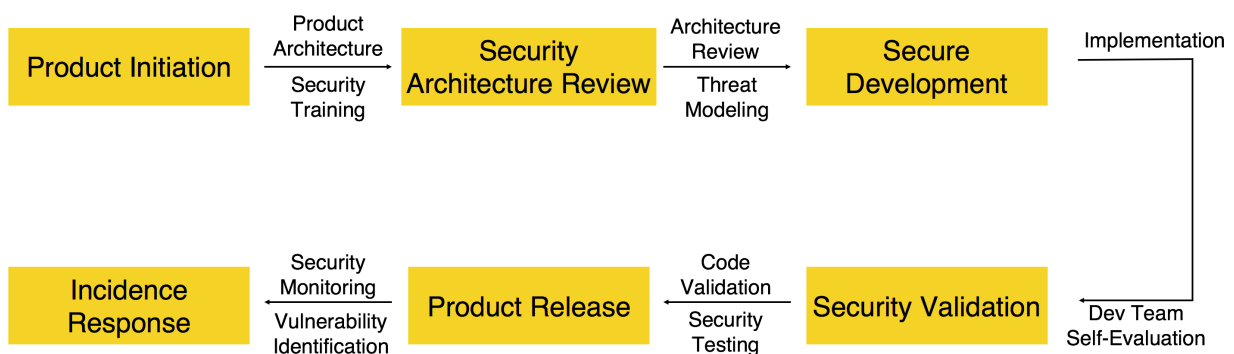
- Verification phase of change: including change verification, review of configuration changes, and change result notification. Alibaba Cloud records all the information throughout the change process and deploys an automatic configuration check tool to verify the configurations of systems after a change.

5.1.4 Cloud Product Security

Alibaba Cloud provides users with a variety of cloud products, including Elastic Compute Service (ECS), Relational Database Service (RDS), Object Storage Service (OSS), MaxCompute, and much more. For details about security features and capabilities of the cloud products, see the "Cloud Product Security" section.

5.1.4.1 Secure Product Lifecycle (SPLC)

SPLC (Secure Product Lifecycle) is a solution tailored for cloud products, designed to integrate the security into the entire product development life cycle. With SPLC, a complete security development mechanism is put into place at each stage, from product architecture review, development, validation, and to incidence response, to ensure the products meet the rigorous security requirements for cloud computing. As a result, SPLC helps cloud products to greatly improve their security capabilities and reduce their security risks.



As shown in the preceding figure, the entire security lifecycle of a cloud product can be divided into six stages: product initiation, security architecture review, secure development, security validation, product release, and incidence response.

In the product initiation stage, the security architect works together with the product team to establish an FRD (functional requirements document) and a detailed architecture

description based on the business requirements and the technical frameworks, and also extract the security baseline, privacy, and compliance requirements applicable to a product. Meanwhile, applicable security training courses and exams are arranged for the product team in this phase, to avoid significant security risks in the subsequent development stages of the product.

In the security architecture review stage, the security architect evaluates the security architecture of the product based on the FRD and the architecture description established in the previous stage, and create threat models for the product. In the process of threat modeling, the security architect creates detailed models for every assets that require protection, security requirements of assets, attack scenarios, and then proposes corresponding security solutions. The security architect then work with the product team to determine all the security requirements for the product, based on the preceding security baseline requirements and the security solutions proposed in threat modeling.

In the secure development stage, the product team must develop the product in accordance with the security requirements, and achieve the relevant security features and goals of the product. To ensure the rapid and continuous development, release and deployment cycle of cloud products, the product team carries out self-evaluation in this stage to confirm that the security requirements have been implemented, and provides corresponding test information to prepare for the next stage of security validation.

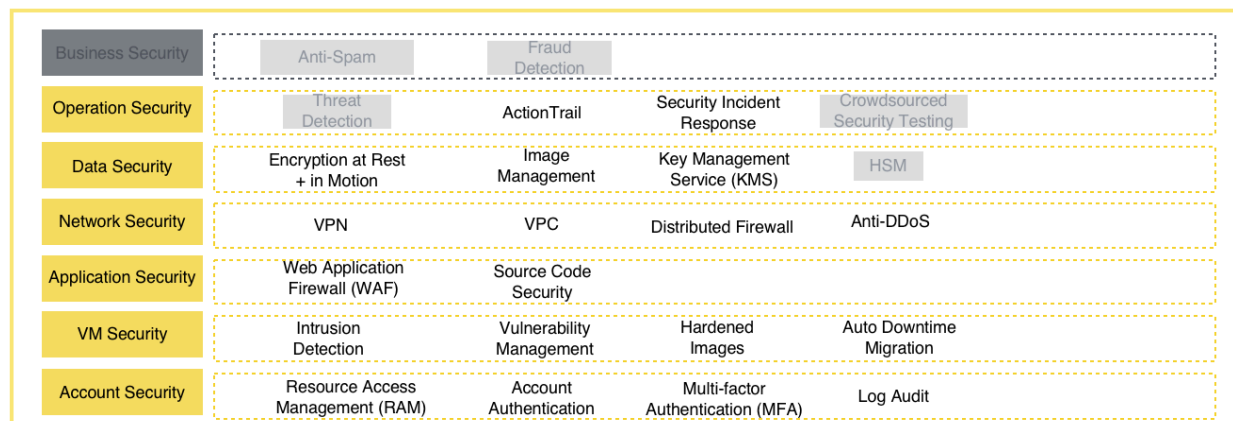
In the security validation stage, the security team implements comprehensive security reviews on the architecture, design, and server environment of the product according to the security requirements, and also performs code review and penetration testing on the product. Any product with security problems found in this stage must be amended correctly.

In the product release stage, only after the product pass the security validation and get the security approval can it be deployed to the production environment through a standard deployment system.

In the incidence response stage, the security incident response team constantly monitors possible security threats to the cloud platform, and identifies security vulnerabilities through external channels (such as ASRC) or internal channels (such as internal security scan, testing, etc.). Once a security vulnerability is detected, the incidence response team quickly rates it and determines its priority and schedule for a fixing. The incidence response team ensures a

rational allocation of resources in order to efficiently and effectively fix vulnerabilities to guarantee the security of Alibaba Cloud and its users.

5.2 User-Side Security Architecture



As shown in the preceding figure, Alibaba Cloud's user-side security architecture focuses on seven aspects: account security, VM security, application security, network security, data security, operation security, and business security.

5.2.1 Account Security

Alibaba Cloud provides a variety of security mechanisms to help you protect your account and resources from unauthorized operations. These security mechanisms include Alibaba Cloud account credentials and MFA tool, creation of separate user accounts, centralized management of individual user permissions, data transmission encryption via HTTPS endpoints, and user activity logging for security auditing. You can use these mechanisms to protect your Alibaba Cloud account. For details, see the "Cloud Product Security - Management and Monitoring - Identity and Access Management" section.

5.2.1.1 Account Authentication

Alibaba Cloud account authentication uses account credentials to verify the real identity of a user. Account credentials usually refers to a user's logon password or AccessKey (AK). Account credentials are confidential, and should be kept in secret by Alibaba Cloud users. With Alibaba Cloud's Resource Access Management (RAM) service, each Alibaba Cloud account can activate a RAM service independently and create individual subusers (i.e. RAM users). An Alibaba Cloud account owner can assign a different password or AccessKey to each RAM user, thus eliminates the security risk with sharing the Alibaba Cloud account's credentials. The account owner can also assign different permissions to different individual accounts by following the principle of least privilege. Password policies for individual RAM

users can be created on the RAM console to ensure that each subuser uses a strong password that is also rotated on a regular basis, thus improving the overall account security. Please note that it is also possible for account owner to create a RAM user who has the permission to create/manage other RAM users.

RAM also supports Security Token Service (STS) which issues temporary security credentials to users who need to access resources temporarily. The permission and auto expiration time (one hour by default) for a token can be changed as necessary to avoid the disclosure of permanent AKs.

- **Logon password**

The password specifications of Alibaba Cloud accounts and the associated risk control policies for logon security are managed by Alibaba Cloud. The password requirements of subusers (i.e. RAM users) can be defined by the Alibaba Cloud account owner, which include the required character combinations of a password, the number of logon retries and password rotation cycle, etc.

- **API AccessKey**

AccessKey(AK) is the credential used for accessing the Alibaba Cloud service APIs. An Alibaba Cloud account owner can log on to the Alibaba Cloud User Center to manage AccessKeys for itself and the RAM users under the account. The account owner can create, freeze, enable, and delete AccessKeys. An AccessKey consists of AccessKey ID and AccessKey secret. The AccessKey ID is public and used to identify a user, and the AccessKey secret is private and used to authenticate a user via signing the signature of an API request. It is recommended that the AccessKeys should be periodically rotated.

Alibaba Cloud strongly recommends that the AccessKey of a RAM user should be used instead of the AccessKey of an Alibaba Cloud account. An Alibaba Cloud account can be viewed as a “root” account, which has full control permissions to all cloud products and resources under such account. Hence, to avoid the risk of exposing the AccessKey of the root account, it is recommended that all users should operate resources at the RAM user level and follow the principle of least privilege.

- **Key pair management**

Alibaba Cloud RAM service provides the ability of key pair management in some regions. RAM users can create their own RSA key pairs. For each key pair, the public key is uploaded to RAM, while the private key is kept secret at the user's end. A user can use the key pair to access STS service in order to obtain a Session AccessKey. With this key, the user can access Alibaba Cloud service APIs where applicable.

5.2.1.2 Multi-Factor Authentication (MFA)

MFA is a simple and effective security best practice that provides an extra level of protection on top of user's logon password or AccessKey. With MFA enabled, a user is asked to enter its username and password (first security factor), and then a variable verification code (second security factor) from an MFA device when logging on to Alibaba Cloud. Alibaba Cloud supports software-based virtual MFA devices. The virtual MFA device is an application that generates a 6-digit verification code, and complies with time-based one-time password (TOTP) standard (RFC 6238). The virtual MFA application can run on mobile hardware devices (including smartphones).

5.2.1.3 Resource Access Management

Resource Access Management (RAM) is a centralized user identify management and resource access control service provided by Alibaba Cloud. By using RAM, an Alibaba Cloud account owner can create subusers (i.e. RAM user) for its employees, systems, or applications, and control their permissions to operate on cloud resources. With an individual logon password or AccessKey, each RAM user can log on to the Alibaba Cloud console or access the cloud service APIs. A newly created RAM user does not have any permissions to operate on resources by default. Only a RAM user that is explicitly authorized by an Alibaba Cloud account can operate on resources belong to that Alibaba Cloud account.

With RAM, Alibaba Cloud account owner can avoid sharing the Alibaba Cloud account credentials with others, and assign minimum operation permissions to different RAM users by following the principle of least privilege. It supports features such as MFA, strong password policies, separation of console users from API users, custom fine-grained authorization policies, grouped authorization management, temporary authorization token and account temporary suspension. The RAM service can be used to define fine-grained authorizations at an API action or resource ID level. The RAM service also supports various restrictive conditions on permission granting (such as constraints on source IP, required SSL/TLS channel, access time period, and MFA, etc.).

5.2.1.4 Log Auditing

Security logs can help Alibaba Cloud user to better understand and diagnose various security situations. Alibaba Cloud provides the ActionTrail service, which enables a unified log management for cloud resources. The ActionTrail service records user logon and resource access operations under each Alibaba Cloud account. Such record includes the operator, operation time, source IP address, resource object, operation name, and operation status. With all operation records saved by ActionTrail, customers can perform security analysis, intrusion detection, resource tracking, and compliance audit. Specifically, in a compliance audit, customers must acquire detailed operation records of the primary account and its subusers. The operation records from ActionTrail can meet these compliance audit requirements.

5.2.2 VM Security

5.2.2.1 Intrusion Detection

Alibaba Cloud users can install a lightweight software called Server Guard on their virtual machine instance, which can work together with the cloud security center for intrusion detection. The intrusion detection for the virtual machine includes remote logon alarm, identification of brute force attack behaviors, webshell detection and removal, and virtual machine anomaly detection.

5.2.2.2 Vulnerability Management

Alibaba Cloud users can install a lightweight software called Server Guard on the virtual machine, which can work together with the cloud security center for vulnerability management. The vulnerability management for the virtual machine incorporates multiple scanning engines (network and local scanning, and vulnerability verification) to thoroughly detect all vulnerabilities in the system at a given time. Features like remote logon alert and one-click webshell removal are provided for a complete vulnerability management experience.

5.2.2.3 Hardened Image

An image is an execution environment template for ECS virtual machine instances. It generally includes an operating system and preinstalled software. Alibaba Cloud ECS tenants can use an image to create an ECS instance or make changes to the system disk of an ECS instance. The security hardening of Alibaba Cloud public image (supports various Linux/Windows release versions) contains three parts: image security configuration, image

vulnerability repair, and default security software in an image. Alibaba Cloud monitors the vulnerabilities in Alibaba Cloud public image operating systems and third-party software in real time to ensure that all high-risk vulnerabilities in Alibaba Cloud public images are repaired in a timely manner. Alibaba Cloud public images are configured with security best practices for the virtual machine by default. Besides, all Alibaba Cloud public images includes Alibaba Cloud security software, such as Server Guard, by default to guarantee the security of instances upon boot up.

5.2.2.4 Auto Downtime Migration

ECS is deployed on the host machine (i.e. the physical server where ECS resides), which may fail due to performance anomaly or hardware failures. After detecting a fault on the host machine, the system would trigger a protective migration to migrate the ECS instance to a normal host machine automatically to ensure the high availability of applications.

5.2.3 Application Security

5.2.3.1 Web Application Firewall

Based on the big data analyzing capabilities of the cloud security, Web Application Firewall (WAF) filters out massive numbers of malicious accesses by defending against SQL injection, XSS, common web server plug-in vulnerabilities, trojan uploads, unauthorized access to resources, and other common OWASP attacks to prevent the leakage of website assets and data and safeguard website security and availability.

5.2.3.2 Source Code Security

In the secure product lifecycle (SPLC) of cloud products, Alibaba Cloud's security experts strictly review and validate the source code security to ensure a high level of code security for Alibaba Cloud products. Alibaba Cloud also constantly performs code security scanning for software in Alibaba Cloud Marketplace to effectively reduce security risks. Meanwhile, Alibaba Cloud strongly recommends that enterprise users should perform blackbox and whitebox code security validation and testing on their applications to prevent security vulnerabilities and improve the security robustness of their businesses.

5.2.4 Network Security

5.2.4.1 Network Isolation

Alibaba Cloud isolates production networks from non-production networks. Direct access is forbidden from a non-production network to any servers and network devices in a

production network. Alibaba Cloud isolates cloud service networks that provide services externally from the physical networks that supports the underlying cloud services functionalities. Network ACLs are configured to forbid access from cloud service networks to physical networks. Alibaba Cloud also takes network control measures to prevent unauthorized devices from connecting to the internal network of the Alibaba Cloud platform and prevent the servers of the internal platform from connecting to external devices.

Alibaba Cloud deploys bastion hosts on the production network boundaries. The maintenance personnel in the office network can access the production network for maintenance management through the bastion hosts only. When logging on to a bastion host, the maintenance staff must perform two-factor authentication by entering the domain account name and password plus a dynamic password. Bastion hosts use advanced encryption algorithms to ensure the confidentiality and integrity of data transmitted through maintenance channels.

5.2.4.2 Virtual Private Network (VPN)

VPN Gateway (Virtual Private Network Gateway) is an Internet-based service that establishes a secure and reliable connection among on-premise data centers, office networks, and Alibaba Cloud Virtual Private Clouds (VPC) over encrypted channels.

5.2.4.3 Virtual Private Cloud (VPC)

With Virtual Private Cloud (VPC), user can build an isolated network environment and customize IP address ranges, network segments, routing tables, and gateway. In addition, user can use connection methods like physical connection and VPN to connect VPCs with traditional IDCs, and thus build a hybrid cloud service.

5.2.4.4 Distributed Firewall

Security group offers the capability of distributed virtual firewall. A security group is a logical group that consists of instances with the same security requirements and mutual trust in the same region. Security groups are used to set network access control for one or more ECS instances. It is an important network security isolation tool and is used to divide network security domains on the cloud.

Each instance belongs to at least one security group. Instances in the same security group can communicate through the network. Instances in different security groups cannot

communicate with each other by default, but a source security group or a source network IP segment can be authorized to access a destination security group if needed.

5.2.4.5 Anti-DDoS

Alibaba Cloud keeps all data centers secure with a self-developed Anti-DDoS service that can automatically detect and route DDoS attacks away from user's network infrastructure. Attack discovery, traffic routing, and cleansing can be completed within five seconds, thus keeping the cloud platform network stable. In order to precisely identify DDoS attacks, the Anti-DDoS service not only triggers the protection mechanism by a traffic threshold value, but also based on network behavior analysis, thus further ensuring the customers' service availability in case of a DDoS attack.

5.2.5 Data Security

Data security and user privacy are the top most priorities of Alibaba Cloud. Alibaba cloud helps its customers to manage and control data security throughout the data lifecycle (production, storage, usage, transmission, propagation, and destruction).

5.2.5.1 Multi-copy Redundancy Storage

Alibaba Cloud uses a distributed storage system, in which a data file is divided into many data fragments that are redundantly stored on multiple devices. Distributed storage improves both data reliability and security.

5.2.5.2 Encryption at Rest and in Motion

Alibaba Cloud uses data encryption to ensure data security, including sensitive data encryption in applications, transparent data encryption in the RDS database, block storage data encryption, object storage system encryption, hardware security modules, and encryption for network data transmission.

- **Data Encryption in Motion**

Alibaba Cloud enables HTTPS encryption to ensure data transmission security. The Alibaba Cloud console uses HTTPS encryption for data transmission. Alibaba Cloud services provide customers with API access points with HTTPS encryption enabled, allowing customers to use AccessKeys to call Alibaba Cloud Service API securely. Alibaba Cloud uses industry standard SSL/TLS protocol with 256-bit key length to address the need for encrypted transmission of sensitive data.

- **Data Encryption at Rest**

Alibaba Cloud provides Key Management Service (KMS) for key management and data encryption capabilities for the encrypted storage of sensitive data on the cloud platform. Such sensitive data include authorization credentials, passwords, and encryption keys. In addition, data encryption is also enabled in different Alibaba Cloud products (please refer to related chapters for each product in the rest of the whitepaper):

- The RDS database product provides a transparent data encryption (TDE) layer.
- Block storage provides data encryption capability which supports automatic encryption of block storage devices used inside the virtual machines.
- Object Storage Service (OSS) supports storing the object's data encrypted, and supports both server-side and client-side encryption.

To achieve sensitive data encryption in customer applications, Alibaba Cloud also has an encryption solution in a hardware-trusted execution environment provided by the processor.

5.2.5.3 Image Management

The ECS provides snapshots and custom images. Snapshots can save the state of a system at a certain point in time for data backup, so that user can achieve quick disaster recovery. The customers can also create custom images using snapshots to include both the operating system and the data environment information of the snapshots in the images. Snapshots are incremental and only the data changed between two snapshots is copied.

5.2.5.4 Key Management

The Key Management Service (KMS) is a secure and easy-to-use management service provided by Alibaba Cloud. Users can ensure the confidentiality, integrity, and availability of the keys with minimal cost. The KMS allows the users to manage and use keys safely and easily, therefore the users can focus on developing the actual features that require encryption and decryption function.

5.2.5.5 Cleanup of Residual Data

The residual data in memory or disk is automatically overwritten with zero value, once such memory or disk is released and recycled. All replaced and scrapped storage devices must be degaussed and physically bent before they can be moved out of the data center.

5.2.5.6 Operation Data Security

Without proper consent of customers, operation personnel of Alibaba Cloud may not access customer data in any way. Additionally, following the principle that production data stays

within the production cluster, Alibaba Cloud restricts and prohibits any channel that would allow the production data to flow out of the production cluster, hence prevents any operation personnel from copying data from the production system.

5.2.6 Operation Security

5.2.6.1 ActionTrail

Alibaba Cloud provides the ActionTrail service, which enables a unified log management for cloud resources. The ActionTrail service records user logon and resource access operations under each Alibaba Cloud account. Such record includes the user name (i.e. operator), operation time, source IP address, resource object, operation name, and operation status. With all operation records saved by ActionTrail, customers can perform security analysis, intrusion detection, resource tracking, and compliance audit.

5.2.6.2 Security Incident Response

Alibaba Cloud has organized a specialized security emergency team to cope with possible security threats and actively respond to and deal with security incidences on the cloud platform.

The security incident response team constantly monitors possible security threats, and identifies security vulnerabilities through external channels (such as ASRC) or internal channels (such as internal vulnerability scanning, testing, etc.). Once a security vulnerability is detected, the incidence response team quickly rates it and determines its priority and schedule for a fixing. The team would also make an announcement to inform customers of the security issue and follow up actions in a timely manner.

Alibaba Cloud has established a specialized team to conduct security tests against Alibaba Cloud to ensure that the security incidence response process is technically sound and effective. Alibaba Cloud also regularly invites third-party teams to perform penetration tests to evaluate the effectiveness of the security protection in place and the smoothness of the incident response process.

6 Cloud Product Security

Alibaba Cloud offers a vast variety of products to its customers. In this section, several key products and their security related functionalities are described in detail to further explain Alibaba Cloud's security capabilities. The products covered in this section are not an exhaustive list of all products available internationally. Additionally, materials covered in this section are what's available from Alibaba Cloud's international offering today. Please refer to our website alibabacloud.com for a complete list of products and more details.

6.1 Elastic Computing

Alibaba Cloud uses the Elastic Compute Service (ECS) to provide various cloud-based elastic computing services for external customers.

6.1.1 Elastic Compute Service (ECS)

An ECS instance of Alibaba Cloud is a virtual computing machine which includes vCPU, memory, OS, disks, bandwidth, and other basic computing components. An instance is the actual operating entity offered by ECS to our customers. The user has the administrator permission for created instances and can log on to, use, and manage the instances at any time. The user can also perform basic operations on an instance, such as attaching a disk, creating a snapshot, creating an image, or deploying an environment, etc.

6.1.1.1 Tenant Isolation

As ECS instances are assigned to different users, isolation among instances provides the needed security barriers among tenants. VMM allows virtual machines at multiple computing nodes to be isolated from each other at the system level, preventing unauthorized access to system resources between tenants and thus guaranteeing the basic computing isolation between computing nodes. Virtualization management layer also provides storage isolation and network isolation.

Tenant isolation is achieved by providing isolation between the virtual machine management system and the customer's virtual machine and isolation between customer's virtual machines. The implementation mechanism of resource isolation between ECS instances is introduced as follows in terms of CPU isolation, memory isolation, storage isolation, and network isolation.

- **CPU isolation**

Based on the hardware virtualization technology Intel® VT-x, the hypervisor runs in the VMX root mode, while the virtual machines run in the VMX non-root mode. Different processor privilege levels (i.e. rings) can effectively avoid unauthorized access of user's virtual machine to physical hosts and the system resources on another user's virtual machines, while also isolating the virtual machines from each other. Hypervisor provides mutually isolated interaction channels between virtual machines and host resources, thus avoiding one tenant to gain read/write/execute access to the host system or other tenants and reducing the risks of system resource sharing.

- **Memory isolation**

On the virtualization layer, the hypervisor isolates the memory. The hardware-supported Extended Page Tables (EPT) technology ensures that the virtual machines cannot access each other's memory space. After an instance is released, all of its memory is reset to zero by the hypervisor to avoid potential information leakage.

- **Storage isolation**

In the basic design of cloud computing virtualization, Alibaba Cloud separates virtual machine-based computing from storage. This separation allows computing and storage to be scaled independently, and makes it easier to provide multi-tenant services. All the I/O operations of a virtual machine are intercepted by the hypervisor to ensure that the virtual machine can only access the disk space allocated to it, thus realizing the security isolation of hard disk space between different virtual machines. After an ECS instance is released, the original disk space and memory space are reliably scrubbed to ensure user data security.

Data disk encryption for ECS virtual machines (i.e. ECS cloud disk encryption) is an automatic storage encryption feature designed for the block storage devices used in virtual machines. ECS disk encryption provides volume encryption to ensure data security. ECS disk encryption provides tenants with a simple and secure encryption method in order to meet their business and certification needs. Since ECS disk encryption is transparent, Alibaba Cloud customers do not have to build, maintain, and protect their infrastructure of key management, to change any existing applications and operations and maintenance processes, or to perform any extra encryption and decryption operations.

ECS disk encryption runs on the network control nodes where ECS instances are located, and it encrypts the data transmitted from ECS instances to the cloud disks. The read and write operations of ECS instances are mapped to the corresponding files on the cloud disks. The user data is fully isolated in the backend storage service with a high reliability and security assurance.

- **Network isolation**

To provide network connections for ECS instances, Alibaba Cloud connects virtual machines to the Alibaba cloud virtual network. Alibaba Cloud virtual network is a logical structure built on top of the physical network structure. All the logical virtual networks are isolated from each other. This isolation prevents the network traffic data from being snooped and/or intercepted by other malicious instances. In addition, all ECS virtual machines can use the VPC and security group firewalls to segregate network access permissions in various scenarios.

To send a message from an ECS instance to another, such message would only be sent to the VSwitch port corresponding to the vNIC of the destination and is not visible to others. It is also not possible for a virtual instance running in promiscuous mode to receive or sniff the traffic flowing to other virtual instances. This is because the hypervisor would not transmit traffic flowing to other destination addresses to such instance. Even two instances on the same physical server and are owned by the same customer cannot sniff the traffic of each other.

6.1.1.2 Security Group Firewall

A security group is a virtual firewall that provides stateful packet inspection (SPI).

A security group is a logical group that groups instances in the same region with the same security requirements and mutual trust. Security groups are used to set network access control for one or more instances. As an important means of security isolation, security groups are used to divide security domains on the cloud.

Each ECS instance belongs to at least one security group. Instances in the same security group can communicate through the network by default. Instances in different security groups cannot communicate over intranet by default. However, mutual access can be authorized between two security groups.

6.1.1.3 SSH Key Pair

SSH key pair, or key pair for short, is a secure authentication method offered by Alibaba Cloud to remotely log on to a Linux instance. Compared with traditional methods using user names and passwords, the SSH key pair is more secure and reliable for logon authentication, making it convenient to remotely log on to a large number of Linux instances.

An SSH key pair is a pair of keys generated through a crypto algorithm: one key is made public, known as the "public key", and the other is kept secret by its user, known as the "private key". The Linux ECS instance stores the public key, and a user can use the private key to connect to the instance where the public key is already configured without the need to enter a password. This is done by using SSH commands or by other related tools. As a security measure, the username and password authentication method is disabled once SSH key pair is enabled on an ECS instance.

6.1.1.4 Anti-IP/MAC/ARP Spoofing

The IP/MAC/ARP spoofing is a known security challenge to a traditional network. Attackers may use IP/MAC/ARP spoofing to disrupt network environment and intercept network transmissions. Alibaba Cloud platform solves the IP/MAC/ARP spoofing problems by isolating any anomalous protocol requests initiated on the data link layer of the host and by avoiding IP spoofing on the network layer of the host.

6.1.1.5 High Availability

- **Server Load Balancer**

Multiple ECS servers can use the SLB service to create a cluster and eliminate single point of failure (SPOF), while improving the availability of the application system. See the "Networking - SLB" section for details.

- **High data reliability**

Through elastic block storage, data is stored via a triplicate distributed system for ECS instances. It also supports switchover, data snapshot backup and rollback, and system performance alerts. Additionally, each data segment has multiple copies, which guarantees rapid restoration when one data segment is physically damaged. Please refer to the product page on alibabacloud.com for detailed service availability and data reliability number.

- **Automatic fault recovery**

ECS instances are deployed on the host (i.e. physical server), which may fail due to abnormal performance or hardware failures. In case of a fault detected on the host, the instances on the faulty host would be automatically migrated to a normal host, thus restoring the normal operation of instances and ensuring the high availability of applications.

6.1.1.6 Snapshots and Images

The ECS provides snapshots and custom images. Snapshots can save the state of data at a certain point in time for data backup or image creating purposes. A user can easily create an automatic snapshot policy for cloud disks and define parameters such as the time of creation, repetition, and retention period, etc.

A user can create custom images using snapshots to include the operating system and data information of the snapshots in the images. With custom images, one can easily create multiple ECS instances with the same operating system and data information.

Snapshots on Alibaba Cloud are taken using an incremental method. In this method, two snapshots are compared and only the data that has changed is copied. A user should take snapshots in the following business scenarios:

- **Routine backup of system and data disks**

A user can back up business-critical data at regular intervals using snapshots to prevent data loss from misoperations, attacks, and viruses, etc.

- **Backup before risky operations**

Before important operations such as OS replacement, upgrading application software or migrating business data, a user should create one or more snapshots. In case of any issues occurring during the upgrade or migration process, one can timely restore data to a normal state of the system using the snapshots taken.

- **Use of multiple copies of production data**

A user can take snapshots of production data to provide close-to-real-time production data for data mining, report queries, and developing and testing applications. It is also possible to take snapshots to reuse data on a disk as the basic data for another disk.

Users can also create their own images and import them into the ECS.

6.1.1.7 Security Image

Alibaba Cloud images integrate patches for all known high-risk vulnerabilities, so that the host is not exposed to high risk attacks after being launched. Once a new high-risk vulnerability is detected/discovered, Alibaba Cloud would promptly update the images and deliver the updated images to customers. Additionally, Alibaba Cloud checks for data integrity of the images to detect any malicious tempering.

6.1.1.8 Hotfix Patch

Alibaba Cloud virtualization platform supports hotfix dynamic patching technology, which can fix system defects or vulnerabilities without user intervention.

6.1.1.9 RAM and STS Support

RAM is a resource access management service provided by Alibaba Cloud. Through RAM, users can create subusers and different groups to manage and control the permissions to access their owned resources.

RAM helps users manage the resource access permissions. For example, to enhance the control of network security, users can assign an authorization policy to specific groups. Such policy stipulates that, if the origin IP address is not from a specified enterprise intranet, such access requests must be denied.

Users can assign different permissions to different groups to manage ECS resources, for example:

- *SysAdmins*: This group requires permissions to create and manage ECS images, instances, snapshots, and security groups. Users can assign to this group an authorization policy which permits the group members to perform all ECS operations.
- *Developers*: This group only needs the permission to use ECS instances. Users can assign to this group an authorization policy that permits the group members to call *DescribeInstances*, *StartInstance*, *StopInstance*, *CreateInstance*, *DeleteInstance*, and other relevant APIs.

If a developer is transferred to the position of a system administrator, such subuser can be easily moved from the Developers group to the SysAdmins group.

ECS also supports the instance RAM role of ECS by using STS. The instance RAM role aims to enable ECS instances to play a role with certain permissions granted to instances directly.

The instance RAM role associates a RAM role to an ECS instance. It allows applications hosted on that instance to access other cloud services by using the STS temporary credential. This feature guarantees the security of the owner's AccessKey and allows fine-grained access control with the help of RAM.

6.1.1.10 Elastic Block Storage

Elastic block storage is a low-latency, persistent, and high-reliability random block level data storage service provided by Alibaba Cloud to ECS users. Elastic block storage supports the automatic replicating of users' data within the zone, thus preventing unexpected hardware faults from causing data unavailability and protecting services against the threat of component faults. Like a hard disk, users can partition the elastic block storage attached to an ECS instance, create a file system, and store data on it.

Elastic block storage supports automatic encryption of block storage devices (i.e. ECS cloud disks) used for ECS instances. The first time a user tries to encrypt a cloud disk in a given region, Alibaba Cloud would automatically create a Customer Master Key (CMK) for the user in the KMS for the region, and each cloud disk is encrypted using a unique 256-bit (AES-256) key protected by the CMK. The user cannot delete this CMK but can query it in the KMS console.

After an encrypted cloud disk is created and attached to an ECS instance, the data in the following list can be encrypted:

- Data on the cloud disk
- Data transmitted between the cloud disk and the instance. However, data in the instance operating system is not encrypted.
- All snapshots created from the encrypted cloud disk. These snapshots are called encrypted snapshots.

Elastic block storage uses a triple distributed system to provide a high data reliability for ECS instances. Please refer to the product page on alibabacloud.com for detailed service availability and data reliability number.

6.1.1.11 Best Practices

- **Security configuration for creating an ECS instance**

- Network type: VPC (Virtual Private Cloud). A VPC refers to a logically isolated private network, in which the network topologies and IP addresses can be customized. VPC supports express connect via physical connection and provides high network scalability.
- Network security group: security group (custom ports) default setting should only allow access from port 22 and port 3389. Note, port 22 is for Linux SSH logon and port 3389 is for Windows remote desktop logon.
- Image: select the needed image from the Alibaba Cloud officially provided images, and select the security hardening feature to load the ECS security components for free and obtain security features like webshell detection, remote logon alert, and brute-force cracking detection.
- SSH key pair logon settings: After an ECS instance is created, go to the ECS console and create an SSH key pair, select the created key pair, and bind it with the created instance.

- **Firewall - network security group settings**

Improper configuration of security group may expose the port or IP of a user's service to the Internet, and thus causing security risks. It is recommended to create a security group rule to prohibit access by default, and keep only the ports needed for the service open.

- **Create a snapshot**

Snapshots can save a copy of disk data at a certain point in time for data backup or custom image production. It is recommended to set up snapshot backup for important ECS instances. It is also a good practice to set up an automatic snapshot policy to backup key business data at regular intervals.

- **System security upgrades**

Security vulnerabilities are inevitable in any system, which requires the users to check for vulnerabilities and perform upgrades regularly. For a Windows system, make sure that the

Window security update is enabled. For a Linux system, one needs to use tools like yum to check for updates.

- **Use Server Guard to secure ECS instances**

Server Guard is a security agent for ECS instances. Through the lightweight agent installed on ECS, Server Guard works together with the cloud security center to provide remote login alert, brute-force cracking detection and webshell detection and removal capabilities.

6.1.2 Auto Scaling

Auto Scaling is a service to automatically adjust computing resources based on the volume of user requests. When demand for computing resources increase, Auto Scaling automatically adds ECS instances to serve additional user requests, or alternatively removes instances in the case of decreased user requests.

Auto Scaling can monitor user clusters and automatically replace unhealthy instances to save maintenance costs. It can also be used to manage user clusters, and automatically add and remove ECS instances as the service volume goes up and down, thus saving infrastructure costs. Auto Scaling is closely integrated with SLB/RDS and automatically attaches or detaches ECS instances to the SLB, and manages the whitelist of RDS accordingly. Auto Scaling also integrates with CloudMonitor to ensure uninterrupted service and streamlined maintenance.

6.1.2.1 Authentication

Auto Scaling performs authentication on each access request. Therefore, each request, whether being sent over HTTP or HTTPS protocol, must contain a proper signature. Auto Scaling uses AccessKey (AK) for identity authentication. For details of the authentication process, see "Cloud Product Security – Management and Monitor – Identity and Access Management –Authentication via AK" section.

6.1.2.2 RAM Support

Auto Scaling supports the RAM service. By enabling the RAM feature, users can grant access permissions to RAM users (i.e. subusers).

6.1.3 Resource Orchestration Service

Resource Orchestration Service (ROS) provides developers and system managers with a simple method to create and manage their Alibaba Cloud resources. Through ROS users can use text files in JSON format to define any required Alibaba Cloud resources, dependencies between resources, and configuration details.

ROS offers a template for resource aggregation and blueprint architecture that can be used as code for development, testing, and version control. Templates can be used to deliver Alibaba Cloud resources and system architectures. Based on the template, API, and SDK, users can then conveniently manage their Alibaba Cloud resource by code (i.e. Infrastructure as Code)

6.1.3.1 RAM Support

ROS supports the RAM service. By enabling the RAM feature, users can grant access permissions to RAM users (i.e. subusers).

6.2 Networking

6.2.1 Server Load Balancer (SLB)

Alibaba Cloud Server Load Balancer is a ready-to-use service that seamlessly integrates with ECS. It is a load balancing service that distributes varying traffic levels among multiple backend ECS instances without manual intervention. SLB ensures high availability by eliminating single point of failure, and protects against SYN flood and DDoS attacks.

6.2.1.1 High Availability

SLB supports cluster deployment, multi-zone deployment and guarantees multilayer disaster tolerance. SLB facilitates local disaster recovery by using multi-zone deployment model for some regions. SLB supports global load balancing and cross-region disaster recovery when used with DNS. Overall, SLB achieves auto scaling based on the application load and ensures that the service is not interrupted during traffic fluctuation. Please refer to the product page on alibabacloud.com for detailed service availability number.

6.2.1.2 Health Check

SLB checks the health status of ECS instances in the ECS pool, and automatically stop forwarding traffic towards instances that did not pass the health check. SLB also restores the isolated ECS instances after they returned to normal. SLB improves the overall service

capabilities of the application by running regular health checks at the backend of ECS instances.

6.2.1.3 Anti DDoS Attack

Alibaba Cloud provides the layer-4 (TCP and UDP protocol) and layer-7 (HTTP and HTTPS protocol) load balancing services. Layer 4 service uses an optimized and customized version of the open source software Linux Virtual Server (LVS) with Keepalived to achieve load balancing. Layer 7 service uses Tengine, a web server project based on Nginx, to achieve load balancing.

SLB when used with Anti-DDoS Basic service provides up to 5 Gbps Anti-DDoS capability free of charge. Additionally, the Layer 7 load balancing service provides the ability to defend against HTTP/S Flood attacks.

6.2.1.4 Access Control

SLB can shield the IP addresses of the backend servers, and only expose virtual IP addresses instead.

SLB provides a source IP address whitelist feature, which allows only whitelisted source IP addresses to access services through SLB.

6.2.1.5 HTTPS

SLB supports HTTPS/SSL/TLS load balancing services:

- SLB provides a certificate management feature for the HTTPS protocol listening. With certificate management, a user does not need to upload certificates to backend ECS instances.
- The deciphering of the ciphertext is offloaded to SLB service to reduce the CPU overheads of backend ECS instances.

SLB provides a certificate management system where user certificates and keys are managed and stored. Private keys uploaded to the certificate management system will be stored encrypted.

6.2.1.6 Log Feature

SLB supports log management, which allows users to view the operation and health check logs of instances.

6.2.1.7 RAM and STS Support

SLB supports the RAM service. By enabling the RAM feature, users can grant access and management permissions to RAM users (i.e. subusers). SLB also supports STS to provides RAM users with authorization credentials for short-term resource access.

6.2.1.8 Best Practices

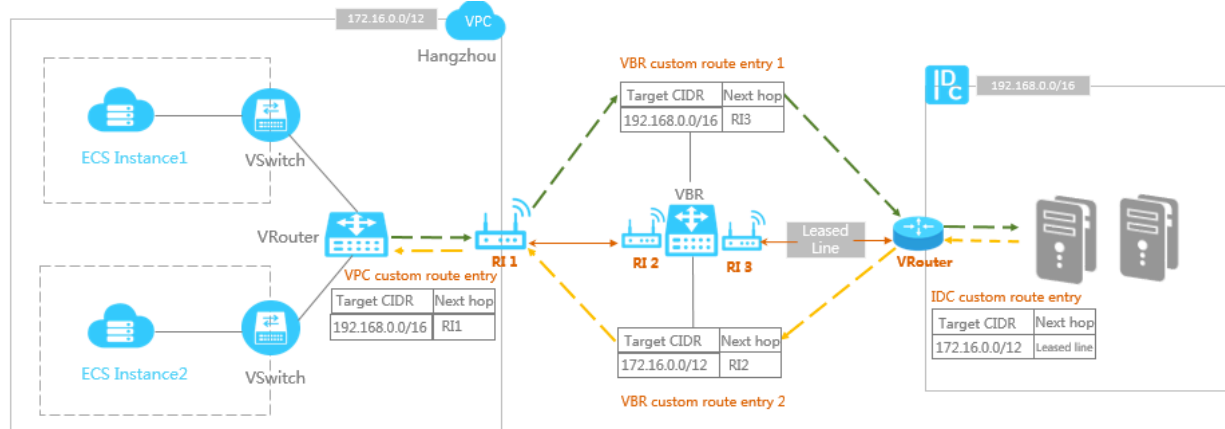
- When using SLB, users must make sure the management port of the back-end ECS instances are not forwarded via SLB to the Internet, such high-risk ECS ports includes 22, 3389, 3306, and 6379.
- If load balancing is used for intranet communications only, users can choose to use intranet SLB instances. Differences between internet instances and intranet instances include:
 - Internet instances: SLB instances only provide public IP addresses. The load balancing service is accessible over the Internet.
 - Intranet instances: SLB instances only provide private IP addresses of Alibaba Cloud. The load balancing service is accessible only over the Alibaba Cloud intranet, and not to Internet users.
- It is recommended that the users should use SLB for HTTPS protocol listening, and it supports both one-way and two-way authentication.

6.2.2 Virtual Private Cloud (VPC)

Virtual Private Cloud (VPC) is a private network established on Alibaba Cloud. VPCs are logically isolated from other virtual networks in Alibaba Cloud.

With Alibaba Cloud' VPC, users can build an isolated network environment with customized IP address ranges, network segments, routing tables, and gateways. Additionally, users can connect a VPC to another VPC or to an on-premises IDC network to form an on-demand network environment, which allows users to smoothly migrate applications to Alibaba Cloud and expand the on-premises IDC.

A typical VPC network architecture is shown in the following figure:



6.2.2.1 Custom Network

Users can customize the network addresses, subnets, and routers in a VPC.

6.2.2.2 Firewall

Security group firewall can be used to partition network security regions among ECS instances within a VPC.

6.2.2.3 Network Access Control

VPC relies on the network access control capabilities of each cloud product within it. For example, ECS instances can be partitioned using the security group firewall, and SLB and RDS services can implement whitelists for access control.

6.2.2.4 Tenant Isolation

ECS instances of different tenants are deployed in different VPCs.

Different VPCs are isolated based on VxLAN tunnel IDs. A VPC can be segmented to one or more subnets similarly to the classic network environment. Different instances in each subnet are interconnected using the same VSwitch. Different subnets are interconnected using VRouters.

6.2.2.5 VPC Communication

VPC itself cannot communicate with other VPCs, classic networks, or the Internet by default. Users can use Internet-facing products such as the EIP, Express Connect, NAT, VPN Gateway, and Internet-facing SLB to implement communications channels.

6.2.2.6 Elastic IP Address (EIP)

An Elastic IP (EIP) address is a public IP address resource that users can purchase and possess independently. It can be dynamically bind to a VPC ECS instance.

An EIP is a NAT IP address. It is located in a public network gateway of Alibaba Cloud, and is mapped to the private network interface card (NIC) of the bound ECS instance via NAT. Therefore, the ECS instance bound with the EIP address can communicate with the Internet without disclosing the EIP address to the NIC. For instances without EIPs, users can configure the routing within a VPC to direct the traffic to go through another instance, which has an EIP and configured with SNAT, to access the Internet.

6.2.2.7 VPN Gateway

VPN Gateway is an Internet-based service that establishes a secure and reliable connection between a customer on-premise data centers and an Alibaba Cloud VPC over an encrypted channel. Alibaba Cloud supports Internet Protocol security (IPsec) VPN connections.

- **Building a hybrid cloud quickly**

VPN Gateway enables communication between a VPC and on-premises data centers by creating an encrypted tunnel, which is easier and faster than applying for a physical connection.

- **Physical connection verification**

For users who want to build a hybrid cloud with the physical connection function provided by Alibaba Cloud Express Connect, they can first use VPN Gateway to establish a VPN tunnel between a VPC and on-premises IDC in order to verify the functionality of such hybrid cloud.

- **Remote disaster recovery and data backup**

For remote disaster recovery scenarios with relatively less demanding requirements and remote data backup scenarios with low data traffic, VPN Gateway can be used to build an encrypted connection between a VPC and on-premises data centers.

6.2.2.8 Express Connect

Express Connect creates private network channels between two VPCs and between a VPC and an on-premises IDC (i.e. physical connection), enhancing the flexibility of network topology as well as the quality and security of cross-network communication. Express Connect also avoids unstable network quality common in public networks and prevents data leakage.

- **Intranet communications between VPCs**

Express Connect supports intranet communications between VPCs in the same region or different regions, of the same account or different accounts. By creating virtual router interfaces on both sides of the VPCs to build an Express Connect on Alibaba's backbone transmission network, users can implement fast, secure, reliable, and convenient communications between VPCs.

- **Intranet communication between an on-premises IDC and a VPC**

Users can connect an on-premises IDC to a VPC through physical connection on the physical layer, and create Virtual Border Router (VBR) and Router Interface (RI) to connect the on-premises IDC to a VPC.

6.2.2.9 NAT Gateway

NAT Gateway is an enterprise-class public network gateway that provides NAT proxy services (SNAT and DNAT), up to 10 Gbps forwarding capacity, and cross-zone disaster recovery capabilities.

As a public network gateway, NAT Gateway must be configured with public IP addresses and bandwidth to work properly. The public IP addresses and bandwidth for NAT Gateway are abstractly grouped into a shared bandwidth package. Such bandwidth package is shared among all IP addresses in the group.

Compare to EIP, the NAT Gateway supports sharing bandwidth among multiple IPs and binding one public IP to multiple ECS instances. NAT Gateway and shared bandwidth packages together provide a high-performance, flexible, and enterprise-class public gateway service.

6.2.2.10 RAM and STS Support

VPC supports the RAM service. By enabling the RAM feature, users can grant access and management permissions of VPC resources to RAM users (i.e. subusers). VPC also supports STS to provides RAM users with authorization credentials for short-term resource access.

6.3 Database

6.3.1 ApsaraDB for RDS

Alibaba Cloud Relational Database Service (RDS) is a stable and reliable online database service with the elastic scaling capability. Based on the Apsara distributed system and high-performance storage of ephemeral SSD, RDS supports MySQL, SQL Server, and PostgreSQL engines. It offers a complete set of solutions for backup, recovery, monitoring, migration, disaster recovery, and troubleshooting database operation and maintenance.

ApsaraDB for RDS provides a variety of security reinforcement features to secure user data, including but not limited to:

- **Network:** IP whitelist, VPC network, SSL/TLS
- **Storage:** TDE (Transparent Data Encryption), automatic backup
- **Disaster recovery:** intra-city disaster recovery (multi-zone zones), remote disaster recovery (disaster recovery instances)

6.3.1.1 Tenant Isolation

RDS enforces that each tenant can only view and operate its own database. Additionally, Alibaba Cloud provides an extra layer of security for the servers on which the RDS services are running. For example, users are prohibited from reading and writing OS files via any database, which ensures that a user's data is not accessible to other users.

6.3.1.2 High Availability

A high availability RDS instance has two database nodes for a master-slave hot standby. If the master node fails, the system can quickly switch to the slave node in seconds. Please refer to the product page on alibabacloud.com for detailed service availability number.

RDS offers an automatic backup mechanism. Users can set a backup schedule or initiate a temporary backup at any time. Data can be recovered from a backup to a temporary RDS instance. After the recovered data is verified, the data can then be migrated back to the master RDS instance.

6.3.1.3 Access Control

- **Database account**

After an RDS instance is created, RDS does not create any initial database account for the user. A user can create a classic mode database account in the RDS console or by using OpenAPI, and set up database-level read/write permissions. If fine-grained access control (i.e. at the table/view/field level) permissions is needed, the user can create a master mode database account.

The initial master mode account must be created in the RDS console or by using OpenAPI, after which the user can logon with the initial account to create and manage additional accounts using SQL commands or Alibaba Cloud's Data Management System (DMS). Please note the initial account may not change the password for the additional accounts that were created. Instead one must remove those accounts and create new ones if their passwords need to be changed.

- **IP whitelist**

By default, RDS instances are set to be inaccessible to any IP addresses. Therefore, the list contains only 127.0.0.1. To add IP whitelist rules, users can use the security controls module in the console or the OpenAPI. The IP whitelist can be updated without restarting RDS instances and does not affect their usage. The IP whitelist function can be set with multiple groups, each of which can contain up to 1,000 IP addresses or IP address segments (800 for SQL Server instances).

6.3.1.4 Network Isolation

- **VPC**

In addition to the IP whitelist function, RDS also supports a VPC implementation. VPC is an isolated network environment on Alibaba Cloud. By using the VPN or physical connection, users can connect their on-premises IDC network to Alibaba Cloud to form an on-demand network environment, thus allowing RDS instances access from both on-premises IDC servers and Alibaba Cloud ECS instances. Please note the users can use a custom RDS IP address segment of the VPC to resolve any IP address conflicts.

By using the VPC and the IP whitelist, the security of RDS instances can be dramatically improved.

- **Internet Access**

By default, RDS instances deployed in a VPC are only accessible to the ECS instances in the same VPC. While not recommended, if necessary, an RDS instance can also be configured to accept requests from Internet by requesting a public IP. Such requests including but not limited to:

- Access requests from ECS EIPs.
- Access requests from on-premises IDC's public IPs.

The IP whitelist applies to all connection methods of RDS instances. Hence, it is strongly recommended to set the whitelist rules on the RDS instance before requesting the public IP.

6.3.1.5 Data Encryption

- **SSL/TLS**

RDS supports the SSL/TLS protocols for MySQL and SQL Server. Users can use the server root certificate provided by RDS to verify that the database service is provided by RDS to prevent man-in-the-middle attacks. Also, RDS supports updating SSL/TLS certificate on the server side so that the users can replace the certificate as needed.

- **TDE**

RDS supports transparent data encryption (TDE) for MySQL and SQL Server. Currently MySQL 5.6 and SQL Server 2008 R2 is supported. TDE can be used to perform real-time I/O encryption and decryption on instance data files. When a user first turn on the TDE option in a region, a Customer Master Key (CMK) for the user is created automatically in the KMS for the region. Each database instance is encrypted with a unique 256-bit (AES-256) database key that is protected by the CMK. The user cannot delete this CMK but can query it in the KMS console.

6.3.1.6 SQL Audit

RDS supports querying SQL audit logs so that users can audit SQL regularly to detect potential issues in a timely manner. The RDS Proxy records all the SQL statements sent to RDS, including the connected IP addresses, names of the accessed databases, accounts used,

SQL statements, running durations, number of returned records, and the time of executions, etc.

6.3.1.7 Backup Recovery

To ensure data integrity and reliability, regular automatic backup is required for databases to ensure data restorability. RDS provides two backup functions, namely data backup and log backup.

6.3.1.8 Instance Disaster Recovery

Alibaba Cloud provides cloud computing services for multiple regions around the world. Each region contains multiple availability zones.

To provide higher availability than single-zone instances, RDS supports multi-zone instances (also known as intra-city dual data centers or intra-city disaster recovery instances). A multi-zone instance deploys physical servers in different zones. When a zone (zone A for example) fails, traffic can be quickly switched to another zone (zone B for example). The entire switchover is transparent to users and requires no application code changes.

RDS also supports cross-region data disaster recovery. A user can copy the RDS instance A' in region A to the RDS instance B' in region B using Data Transmission (where instance B' is a complete and independent RDS instance with a separate connection address, account, and permissions).

6.3.1.9 Software Upgrades

RDS provides users with new versions of database software when applicable. Generally, version upgrades are not mandatory. The database of an RDS instance is upgraded only when an RDS instance is restarted. In rare cases (such as with critical bugs and security vulnerabilities), RDS enforces database upgrades during the maintenance of the instance. Such mandatory upgrades only result in transient database disconnections but no obvious negative effects for the application provided that the database connection pool is correctly configured. Users can change the maintenance time either in the RDS console or by using the OpenAPI to prevent mandatory upgrades during their services peak hours.

6.3.1.10 Support for RAM and STS

RDS supports the RAM service. By enabling the RAM feature, users can grant access and management permissions to RAM users (i.e. subusers). RDS also supports STS to provides RAM users with authorization credentials for short-term resource access.

6.3.1.11 Best Practices

- **Network configuration**

Users can access RDS over either an intranet or the Internet. If the RDS instances are used inside of the Alibaba Cloud, it is recommended to configure the RDS instances to allow accesses over intranet only. If you access RDS over a VPC, you are advised to select a VPC instance.

- **IP whitelist configuration**

In the RDS console, users should always set the RDS IP whitelist to only include the corresponding ECS IP or the VPC IP over the intranet, and not allowing public Internet IPs unless it is absolutely necessary.

6.3.2 ApsaraDB for Redis

ApsaraDB for Redis is compatible with open-source Redis protocol standards and provides persistent memory database services. Based on its high-reliability dual-machine hot standby architecture and seamlessly scalable cluster architecture, this service can meet the needs of businesses that require high read/write performance and flexible capacity adjustment.

6.3.2.1 High Availability

In addition to the single-node architecture, ApsaraDB for Redis also supports dual-node hot standby and cluster architectures to ensure the high availability of system services.

- **Single-node architecture**

This is suitable for cache-only scenarios. It supports flexible configuration changes for single-node clusters and provides cost-effective performance that suits high QPS scenarios.

- **Hot standby architecture**

During system operation, data is synchronized between the master and slave nodes. If the master node fails, the system automatically switches over to the slave node in a matter of seconds. The entire process is automatic without affecting users' services and businesses. The master/slave architecture guarantees the high availability of system services.

- **Cluster architecture**

Cluster instances adopt a distributed architecture, with each node working in master/slave mode. This supports automatic disaster recovery switchover and failover. Users can choose from multiple cluster specifications based on their business needs and the service allows scalable database performance.

6.3.2.2 Authentication

ApsaraDB for Redis provides an identity authentication mechanism based on instance IDs and passwords.

6.3.2.3 Access Control

- **Intranet access**

ApsaraDB for Redis only supports accesses through the Alibaba Cloud intranet. In other words, only applications on Alibaba Cloud ECS instances can be connected to ApsaraDB for Redis for data operations.

- **IP whitelist**

Before using Redis instances, the IP addresses or IP address segments of the requestor instances must be added to the whitelist of the target Redis instance for accessing the database.

6.3.2.4 Backup Recovery

ApsaraDB for Redis provides a backup archiving feature, which can retain automatic and manual backup files for seven days without charge. After seven days, the backup files are automatically deleted.

6.3.2.5 RAM Support

ApsaraDB for Redis supports the RAM service. By enabling the RAM feature, users can grant access and management permissions to RAM users (i.e. subusers).

6.3.3 ApsaraDB for Memcache

ApsaraDB for Memcache is a scalable memory-based cache service that supports high-speed access to large amounts of small data. ApsaraDB for Memcache can greatly cut down the backend storage load and speed up the response of websites and applications. Based on the Apsara distributed system and high performance storage, ApsaraDB for Memcache provides a complete set of solutions for master/slave hot standby, disaster recovery, business monitoring, data migration, and other scenarios.

6.3.3.1 High Availability

ApsaraDB for Memcache consists of three components, namely, the proxy server (service proxy), the partitioning server, and the configuration server.

- **Proxy server**

It is single-noded. A cluster structure may contain multiple proxies and the system automatically implements load balancing and failover for the proxies.

- **Partitioning server**

Each partitioning server is in a high availability dual-copy architecture. The system automatically implements the master/slave switchover in case of a fault in the master node to ensure the high availability of services.

- **Configuration server**

It is used to store cluster configuration information and partitioning policies. It currently adopts the dual-copy architecture to ensure high availability.

6.3.3.2 Authentication

ApsaraDB for Memcache supports SASL (Simple Authentication and Security Layer) authentication, which requires users to enter the correct user name and password before operating on the data.

ApsaraDB for Memcache is accessed using SASL authentication by default. It also supports password-free login provided that the IP whitelist is enabled.

6.3.3.3 Access Control

- **IP whitelist**

To guarantee database security and stability, users must add IP addresses or IP address segments used for database access to the whitelist of the target instance before using ApsaraDB for Memcache. Correct use of the whitelist improves access security protection for ApsaraDB for Memcache. It is recommended that the whitelist be regularly maintained.

- **VPC**

ApsaraDB for Memcache is fully connected to VPC, and users can build an isolated network environment based on Alibaba Cloud.

6.3.3.4 Backup Recovery

ApsaraDB for Memcache provides a backup archiving feature, which can retain automatic and manual backup files for seven days without charge. After seven days, the backup files are automatically deleted.

6.3.3.5 RAM Support

ApsaraDB supports the RAM service. By enabling the RAM feature, users can grant access and management permissions to RAM users (i.e. subusers).

6.4 Storage and CDN

6.4.1 Object Storage Service (OSS)

The Object Storage Service (OSS) is a massive, secure, cost-effective, and highly reliable cloud storage service provided by Alibaba Cloud. OSS enables users to focus on their core business needs with platform-independent RESTful APIs, elastically scalable capacity and processing capabilities, and its Pay-As-You-Go billing method.

6.4.1.1 Authentication

OSS uses AccessKey (AK) for identity authentication. For details of the authentication process, see "Cloud Product Security – Management and Monitor – Identity and Access Management –Authentication via AK" section.

6.4.1.2 Access Control

The accesses to OSS resources are divided into access by the owner and access by third-party users. The owner here refers to the bucket owner, also known as the developer. Third-party users are users other than owner who access resources in a bucket. Also, accesses include anonymous accesses and signature-based accesses. In OSS, a request that contains no authentication code (i.e. signature information) is considered anonymous access. Signature-based access refers to a request that contains signature information in the request header or URL as stipulated in the OSS API documentation.

OSS provides access control for buckets and objects.

Three access permissions are available for a bucket:

- **public-read-write:** Anyone (including anonymous users) can perform read, write, and delete operations on the objects in the bucket. The fees incurred by such operations are borne by the owner of the bucket. **Use this permission with caution.**
- **public-read:** Only the owner of the bucket or any authorized users can perform write or delete operations on the objects in the bucket. Anyone (including anonymous users) can read the objects in the bucket.
- **private:** Only the owner of the bucket or any authorized users can perform read, write, and delete operations on the objects in the bucket. Others cannot access the objects in the bucket without authorization.

When a new bucket is created without a permission specified, OSS automatically sets the bucket permission as private.

Four access permissions are available for an object:

- **public-read-write:** All users can perform read and write operations on the object.
- **public-read:** Only the owner of the object can perform read and write operations on the object. Others can only read the object.
- **private:** Only the owner of the object can perform read and write operations on the object. Others have no permissions to operate on the object.
- **default:** Object follows the access permission of bucket.

If an object is uploaded without the assigned permission, its permission is set as default by the OSS.

6.4.1.3 Support for RAM and STS

OSS supports the RAM service. By enabling the RAM feature, users can grant access and management permissions to RAM users (i.e. subusers). OSS also supports STS to provides RAM users with authorization credentials for short-term resource access.

6.4.1.4 High Availability

OSS provides automatical scaling and delivers a high service availability. Additionally, the OSS data is stored in three copies, and is designed for a high data persistence rate. Please refer to the product page on alibabacloud.com for detailed service availability and data reliability number.

6.4.1.5 Tenant Isolation

The OSS segments user data and adds user tags to them. The user data is stored discretely in a distributed file system and separated from the data index. The OSS users are authenticated using symmetric AccessKeys. The signature information is verified upon each request. After a user passes the authentication, OSS reassembles the discrete data based on the user tags. This achieves the separation of data storage among multiple tenants.

6.4.1.6 Access Log

The OSS automatically saves the access logs. A bucket owner can enable the access logging feature through the OSS console. After access logging is enabled for a source bucket, the OSS generates a data object that contains access logs for that bucket (by hour) and writes the object into the user-specified target bucket according to predefined naming rules.

6.4.1.7 Anti-leech

The OSS is a Pay-As-You-Go service. To prevent user data on the OSS from being leeches, the OSS supports anti-leech based on the field referer in the HTTP header. Users can log on to the OSS console or use OpenAPI to set a referer whitelist for a bucket or whether to allow empty referrer requests. For example, for a bucket named oss-example, set its referer whitelist as <http://www.alibabacloud.com/>. Then, only requests with the referer of <http://www.alibabacloud.com/> can access the objects in the oss-example bucket.

6.4.1.8 Cross-Origin Resource Sharing Access

Cross-origin access, or cross-origin of JavaScript, is a type of browser restriction for security consideration, namely, the same-origin policy. When Website A tries to use the JavaScript code on its webpage to access Website B, the attempt is rejected by the browser because A and B are two websites of different origins.

However, cross-origin access is a commonly used on a day-to-day basis. For example, OSS is used at the backend for the website www.a.com and a JavaScript-based

upload function is provided on the webpage. However, requests on the webpage are only sent to `www.a.com`, whereas all requests sent to other websites are rejected by the browser. As a result, user-uploaded data must be relayed to other sites through `www.a.com`. If cross-origin access is configured, data can be uploaded directly to OSS instead of relaying it through `www.a.com`.

6.4.1.9 Server-side Encryption

OSS supports server-side encryption for data uploaded by the users. When a user uploads data, OSS encrypts the user data and permanently stores the data with encryption; when the user downloads the data, OSS automatically decrypts the encrypted data, returns the original data to the user, and also declares in the header of the returned HTTP request that the data has been encrypted on the server.

OSS server-side encryption provides the following options for users to choose from (based on the key management policy):

- **Fully managed by OSS**

The generation and management of data encryption keys are conducted by OSS, which provides strong and multi-factor security measures to protect data. The data encryption algorithm adopted is AES-256.

- **CMK managed by KMS**

In addition to the AES-256 encryption algorithm, KMS is used to manage CMKs (keys that encrypt data keys) and generate data keys. Envelope encryption is used so that the encrypted data key blobs are stored with the encrypted data objects, and the data key blobs would be decrypted and used to decrypt the data ciphertext upon data retrieval. Please note that an extra small fee for using KMS keys is collected in this option. This option is only supported in Mainland China and Singapore regions currently, and will be available to other regions as soon as possible.

The OSS server-side encryption code is an attribute of objects. When creating an object, the user only needs to add the HTTP header, "x-oss-server-side-encryption", to the Put Object request and specify its value as "AES256"/ "KMS" to encrypt and store the object on the server, where "AES256" refers to OSS fully managed server-side encryption, and "KMS" refers to CMK managed by KMS server-side encryption option.

6.4.1.10 Client-side Encryption

Client encryption refers to that the encryption is completed before the user data is sent to the remote server while the encryption key used is only kept locally at the client-side. Therefore, others cannot obtain the original data without the secret encryption key even if the data is leaked.

6.4.1.11 Best Practices

- **Data access control**

The OSS provides permission access control for buckets. Normally, when storing static images, CSS, and JS resources of applications in the OSS, the users should set the permission of the bucket as public-read.

For sensitive data, the permission of the bucket must be set as private so it can only be accessed with AK authentication.

- **Integrity of data transfer**

An error may occur when data is transferred between the client and the server.

Currently, the OSS supports returning the CRC64 value of the object uploaded. The client can compare the CRC64 value with that calculated locally to verify the integrity of data transfer.

6.4.2 Table Store

Table Store is a NoSQL database service built upon the Alibaba Cloud's Apsara distributed system, enabling the users to store and access large volumes of structured data in real time. Table Store organizes data into instances and tables that can seamlessly scale using data partitioning and load balancing. Applications use the Table Store service through the Table Store API/SDK or the Table Store console.

6.4.2.1 Authentication

Table Store uses AccessKey (AK) for identity authentication. For details of the authentication process, see "Cloud Product Security – Management and Monitor – Identity and Access Management –Authentication via AK" section.

6.4.2.2 High Availability

Through automatic fault detection and data migration, Table Store protects applications from faults and errors that may occur on the underlying hardware platform, and delivers a

high service availability. Table Store manages data with multiple backup copies and enables quick recovery from a backup failure to provide a high service reliability. Please refer to the product page on alibabacloud.com for detailed service availability and reliability number.

6.4.2.3 Strong Consistency

Table Store ensures strong consistency on data writing. Once a write operation is returned success, applications using Table Store can read the latest data written.

6.4.2.4 Monitoring Integration

Users can log on to the Table Store console to obtain monitoring information in real time, including the requests per second and the average response latency, etc.

6.4.2.5 Support for RAM and STS

Table Store supports the RAM service. By enabling the RAM feature, users can grant access and management permissions to RAM users (i.e. subusers). Table Store also supports STS to provides RAM users with authorization credentials for short-term resource access.

6.4.2.6 Support for VPC

Table Store supports accessing within a VPC. Users can activate and bind the VPC on the Table Store console.

6.4.3 Network Attached Storage

Alibaba Cloud Network Attached Storage (NAS) is a file storage service oriented towards Alibaba Cloud ECS, HPC, and Docker computing nodes. It features a distributed file system with unlimited capacity and performance scaling, with namespace and multiple client access support. NAS supports standard file access protocols (NFSv3, NFSv4), so existing applications do not need to be modified.

6.4.3.1 Access Control

NAS supports standard directory/file permission operations on a file system, and supports read/write/execute permission settings for specific users/groups. NAS supports mount points in both a VPC and a classic network, and allows ECS instances within the same VPC or under the same Alibaba Cloud account to access the file system. NAS also provides IP-level permission groups for fine-grained access control.

6.4.3.2 RAM Support

NAS supports the RAM service. By enabling the RAM feature, users can grant access permissions to RAM users (i.e. subusers).

6.4.3.3 High Availability

With a no single point of failure design, NAS provides a high availability and data reliability. Alibaba Cloud NAS substantially saves maintenance cost and reduces data security risk in comparison to a self-built NAS. Please refer to the product page on alibabacloud.com for detailed service availability and data reliability number.

6.4.4 Alibaba Cloud CDN

Alibaba Cloud Content Delivery Network (CDN) is a distributed network built on, and overlaying, the bearer network, and is composed of edge node server clusters distributed across different regions. A CDN replaces the traditional data transmission mode centered on web servers. CDN delivers the source content to edge nodes and works with a precise scheduling system. It distributes user requests to the appropriate nodes, allowing users to retrieve the desired content as quick as possible, effectively reducing the Internet congestion problem and increasing the response time of user accesses.

6.4.4.1 Authentication

Alibaba Cloud CDN uses AccessKey (AK) for identity authentication. For details of the authentication process, see "Cloud Product Security – Management and Monitor – Identity and Access Management –Authentication via AK" section.

6.4.4.2 Tenant Isolation

The user data cached on the CDN is tagged. The data is stored discretely in a distributed file system and separated from the data index. Users are authenticated using symmetric AccessKeys. User requests are differentiated by domain granularity. After a user passes authentication, CDN reassembles the discrete data based on the user's domain. This achieves the separation of data storage among multiple tenants.

6.4.4.3 URL Authentication

The URL authentication feature protects user's site resources from illegal download and misuse. Leeching issues are only partially solved by adding the referer blacklist or whitelist, this is because the referer content may be forged. Applying URL authentication is recommended to protect the origin site resources in a more secure and effective manner.

The URL authentication function uses Alibaba Cloud CDN nodes in combination with client resource sites to provide a more secure anti-theft protection for origin site resources. The CDN client site provides a user with an encrypted URL (including permission verification information) and the user uses it to initiate a request to the CDN node. The CDN node verifies the permission information in the encrypted URL to determine the legality of the request. Legal requests will receive a normal response and illegal requests will be rejected. This protects CDN client site resources.

Alibaba Cloud CDN is compatible with various authentication modes. Users can select an appropriate mode based on their business needs to effectively protect their origin server resources.

6.4.4.4 HTTPS Acceleration

Alibaba Cloud CDN provides HTTPS acceleration. Users only need to enable the secure acceleration mode and then upload the certificate and private key for the CDN domains. The service also supports viewing, disabling, enabling, and editing certificates. Advantages of HTTPS acceleration includes:

- Key user information is encrypted during transmission, thus preventing leakage of sensitive information, such as session IDs or cookies, etc.
- Integrity verification is performed on all data during transmission, protecting the DNS or content from being hijacked, tampered with, or suffering from other “man in the middle” attacks.

6.4.4.5 Anti-leech

Alibaba Cloud CDN provides an anti-leech feature. The anti-leech feature is based on the HTTP referer mechanism where the referer, namely an HTTP header field, is used for source tracking, source recognition and processing. Users can configure a referer blacklist or whitelist to identify and filter visitors in order to limit access to their CDN resources.

6.4.4.6 IP Blacklist

Alibaba Cloud CDN provides an IP blacklist feature. If an IP address is added to the blacklist, it cannot access the corresponding CDN domain.

6.4.4.7 httpDNS

A traditional DNS resolution is implemented by accessing the local DNS of an ISP in order to obtain the resolution result, which is prone to DNS hijacking, DNS errors, and cross-network traffic and leads to failed or slow website access.

The httpDNS is a DNS service that uses HTTP protocol to directly access the Alibaba Cloud CDN server to perform domain resolution. Because it bypasses the local DNS of an ISP, it can avoid DNS hijacking and obtain real-time accurate DNS resolution results.

6.4.4.8 RAM Support

Alibaba Cloud CDN supports the RAM service. By enabling the RAM feature, users can grant access and management permissions to RAM users (i.e. subusers).

6.5 Analytic and Big Data

6.5.1 MaxCompute

MaxCompute is a big data processing platform that is mainly used for batch structural data storage and processing, which can provide massive data warehouse solution and big data modeling service. MaxCompute provides a convenient way to analyze and process big data. Users are able to analyze big data without concerning details of distributed computing.

6.5.1.1 Authentication

MaxCompute supports two account systems: Alibaba Cloud account system and RAM account system. Please note, MaxCompute project can only identify the Alibaba Cloud account system by default.

6.5.1.2 Authorization Management

MaxCompute is a data processing platform that supports multiple tenants. Different tenants have different data security requirements. MaxCompute offers project-level security configuration to satisfy the flexible data security requirements of different tenants. After a project is created by a user, the user is the owner of the project. That is, all objects in the project (such as tables, instances, resources, and UDF) belong to the user. Unless authorized by the owner, no one except the owner has the permission to access the objects in the project.

When the owner of the project decides to authorize another user, the owner needs to add the user to the project. Only users belong to the same project can be authorized by the project owner.

A role is a set of access permissions. When an owner wants to assign a group of users with the same permission, the owner can authorize them through role. Role-based authorization can greatly simplify the authorization process and reduce the authorization management costs.

MaxCompute can assign different permissions to users or roles in a project based on four types of objects (project, table, function, and resource instance).

MaxCompute supports ACL authorization, which is an object-based authorization method. The permission data authorized by the Access Control List (ACL) is considered as a type of sub-resource of the object. Authorization can be performed only when the object exists. When the object is deleted, the authorized permission data is automatically deleted. MaxCompute authorization supports the syntax similar to the GRANT and REVOKE commands defined by SQL92. It grants or revokes permissions to/from the existing project object through simple authorization statements.

6.5.1.3 Cross Project Resource Sharing

Package is a mechanism for sharing data and resources across projects. It is used for cross project user authorization.

The administrator/owner of project A can perform authorization on the objects to be used in project B by creating a package that includes all resources needed by B, and then permits project B to install the package. After project B's administrator/owner installs the package, the administrator/owner can then determine whether to grant permissions of accessing the package to other users of project B as needed.

6.5.1.4 Data Protection

If a project contains highly sensitive data that cannot be shared with other projects, the ProjectProtection option can be set to true which prohibits any access to data in the project.

6.6 Application Service

6.6.1 Log Service

The Log Service (or Log for short) is an all-in-one service for log-type data that allows users to quickly complete log data collection, consumption, shipping, query, and analysis. Log Service can help users increase O&M efficiency and build processing capabilities to handle massive log volumes.

6.6.1.1 High Availability

Log data is stored in a distributed file system, and multiple copies are stored to ensure storage reliability.

6.6.1.2 Read-only Log System

Tampering prevention is an important feature of Log. Log provides an append-only log system, which only allows users to append logs but not modify previous logs to prevent log tampering.

6.6.1.3 Offline Archiving

In addition to the real-time query and analysis functions, Log also provides the capability to store log archives to MaxCompute and OSS, so that users can analyze data with MaxCompute and opensource big data analytical software.

6.6.1.4 Authentication

Log authentication uses AccessKey (AK) for identity authentication. For details of the authentication process, see "Cloud Product Security – Management and Monitor – Identity and Access Management –Authentication via AK" section.

6.6.1.5 RAM Support

Log supports the RAM service. By enabling the RAM feature, users can grant access and management permissions to RAM users (i.e. subusers).

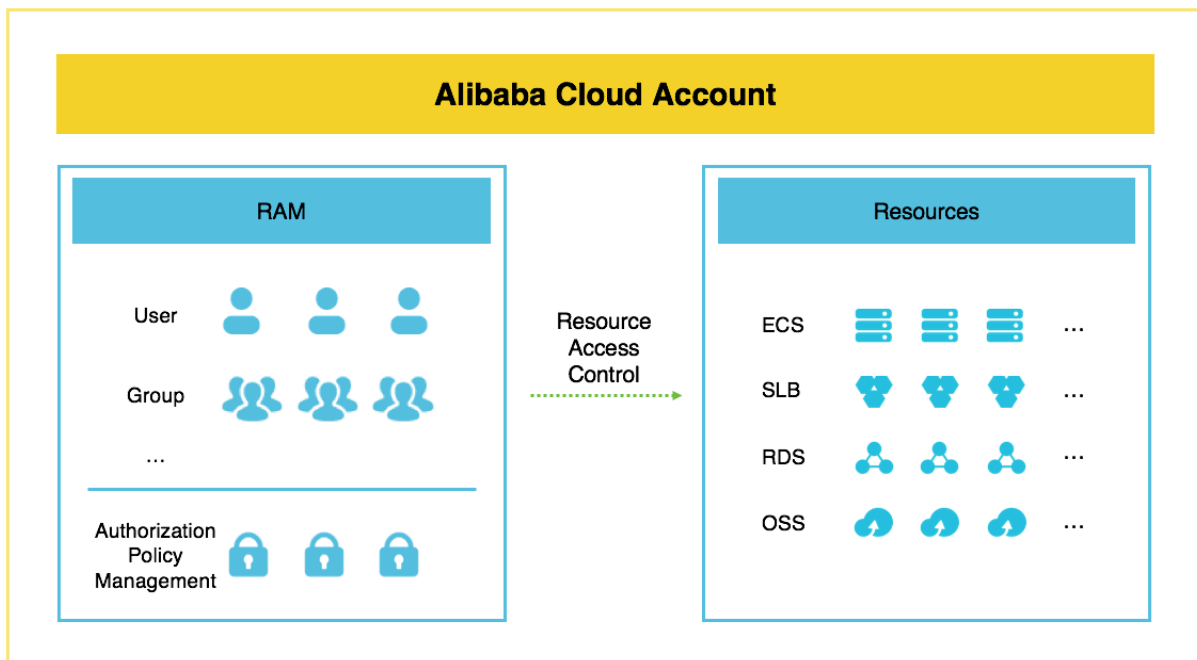
6.7 Management and Monitoring

6.7.1 Identity and Access Management (RAM&STS)

Alibaba Cloud provides multiple tools and features to help users securely authorize access to resources for different scenarios. Among those, the Resource Access Management (RAM)

service is provided for user identity management and resource access control. RAM enables an Alibaba Cloud account (i.e. primary account) to have multiple independent subusers (i.e. RAM users). It also supports such features as multi-factor authentication, strong password policies, separation of console users from API users, custom fine-grained authorization policies, grouped authorization, temporary authorization token and account temporary suspension. The RAM service can be used to define fine-grained authorizations at an API action or resource ID level. The RAM service also supports various restrictive conditions on permission granting (such as constraints on source IP, required SSL/TLS channel, access time period, and MFA, etc.).

RAM provides centralized services on user identity and access control management. The following figure shows the relations between RAM and other cloud services:



RAM is the basis for the security management and security operations of Alibaba Cloud accounts. With RAM, the account owner (or a RAM user with granted permission) can assign a different password or API AccessKey to each subuser, which eliminates the security risk with sharing the Alibaba Cloud account's credentials. In the meantime, assigning different work permissions to different subusers could also reduce the risks associated with excessive permissions that comes with using a single cloud account only.

6.7.1.1 User Management

An Alibaba Cloud account can create one or more independent subusers through the RAM service. The relationship between an Alibaba Cloud account and its RAM subusers is as

follows: (1) From the perspective of ownership, an Alibaba Cloud account is the owner of Alibaba Cloud resources and the basis on which resource usage is billed, while RAM users only exist in RAM instances of an Alibaba Cloud account. RAM users do not possess resources, and the resources they create under authorization belong to the Alibaba Cloud account. RAM users do not possess bills, and all expenses incurred by their authorized operations are also debited to the Alibaba Cloud account. (2) From the perspective of permission management, the relationship between an Alibaba Cloud account and its RAM users is analogous to that of Root and User in Linux. The Alibaba Cloud account has all the permissions to operate on and control resources, and a RAM user only possesses certain permissions authorized by the Alibaba Cloud account. In addition, the Alibaba Cloud account can revoke the permissions authorized to a RAM user at any time. Meanwhile, the Alibaba Cloud account can grant RAM users permission to operate on RAM resources such as RAM user creation and management.

Each RAM user represents the identity of a security principal, such as some operational personnel or an application. If a new user or application wants to access the cloud resources under an Alibaba Cloud account, a RAM user must be created and be granted access to relevant resources. If multiple RAM users are created under one Alibaba Cloud account, as a more convenient approach to manage RAM users and permissions, it is recommended to create groups to categorize RAM users based on job function and assign permissions to the groups.

The administrator can also create a kind of users called "RAM role" through RAM. Both the RAM role and RAM users are identity objects for RAM management. The difference between the RAM role and RAM users is that the RAM role is, in a sense, a virtual user who does not have a long-term authentication AccessKey and cannot be used without being assumed by a trusted RAM user identity.

6.7.1.2 Account Credential

Account credentials is used to verify the real identity of a user. Account credentials usually refers to a user's logon password or AccessKey (AK). Account credentials are confidential, and should be kept in secret by Alibaba Cloud users.

RAM supports the following credentials:

- Logon password

The password specifications of Alibaba Cloud accounts and the associated risk control policies for logon security are managed by Alibaba Cloud. The password requirements of subusers (i.e. RAM users) can be defined by the Alibaba Cloud account owner, which includes the required character combinations of a password, the number of logon retries and password rotation cycle, etc.

- **Multi-factor authentication (MFA)**

MFA is a simple and effective security best practice that provides an extra level of protection on top of user's logon password or AccessKey. With MFA enabled, A user is asked to enter its username and password (first security factor), and then a variable verification code (second security factor) from an MFA device when logging on to Alibaba Cloud. Alibaba Cloud supports software-based virtual MFA devices. The virtual MFA device is an application that generates a 6-digit verification code, and complies with time-based one-time password (TOTP) standard (RFC 6238). The virtual MFA application can run on mobile hardware devices (including smartphones).

- **API AccessKey**

AccessKey(AK) is the credential used for accessing the Alibaba Cloud service APIs. An Alibaba Cloud account owner can log on to the Alibaba Cloud User Center to manage AccessKeys for itself and the RAM users under the account. The account owner can create, freeze, enable, and delete AccessKeys. An AccessKey consists of AccessKey ID and AccessKey secret. The AccessKey ID is public and used to identify a user, and the AccessKey secret is private and used to authenticate a user via the signature of an API request. It is recommended that the AccessKeys should be periodically rotated.

Alibaba Cloud strongly recommends that the AccessKey of a RAM user should be used instead of the AccessKey of an Alibaba Cloud account. An Alibaba Cloud account can be understood as a “root” account, which has full control permissions to all cloud products and resources under such account. Hence, to avoid the risk of exposing the AccessKey of the root account, it is recommended that all users should operate resources at the RAM user level and follow the principle of least privilege.

- **Key pair management**

Alibaba Cloud RAM service provides the ability of key pair management in some regions. RAM users can create their own RSA key pairs. For each key pair, the public key is

uploaded to RAM, while the private key is kept secret at the user's end. A user can use the key pair to access STS service in order to obtain a Session AccessKey. With this key, the user can access Alibaba Cloud service APIs where applicable.

6.7.1.3 Authentication via AK

There are many Alibaba Cloud products that use AccessKey (AK) for identity authentication. Before sending a request to a cloud service as an individual identity, a user must generate a signature string for the request according to the format specified by the product and then sign the signature string using the AccessKeySecret (often using HMAC-SHA1) to generate an authentication code.

After receiving the request, the service finds the corresponding AccessKeySecret based on the AccessKeyID, and obtains the signature string and authentication code in the same way. If the obtained authentication code is the same to the one provided, the request is assumed valid. If not, the service rejects the request and returns an HTTP 403 error.

Users can also directly use the SDKs provided by the cloud services with different AccessKeys for different types of identity authentication.

6.7.1.4 Group Management

If an Alibaba Cloud account owner has created multiple RAM users under the account, it is recommended that the account owner uses groups to better manage the users and their permissions. The account owner can create groups for RAM users who share the same responsibilities (such as Admins, Developers, and Accounting), and categorize and grant permissions to each group. When the responsibilities of a user change, one only needs to move the user to another corresponding group accordingly without affecting other users. When the permissions of a group change, one only needs to modify the authorization policies for the group, so that the policies can be directly applied to all RAM users associated with the group.

6.7.1.5 Permission and Authorization Policy Management

1) Permission

Alibaba Cloud uses permission to describe the ability of an operating security principal (such as user, user group, and RAM role) to access a specific resource. A permission is used to allow or deny the execution of certain operations on certain resources under certain conditions.

- **The Alibaba Cloud account (i.e. primary account/root account/resource owner) controls all permissions**
 - Each resource has only one owner (resource owner). The owner must be an Alibaba Cloud account. This account pays for and has full control permissions to the resource.
 - The resource owner is not necessarily the resource creator. For example, if a RAM user is granted the permission to create resources, the resources created by this user belong to the Alibaba Cloud account that the RAM user is under. In this case, the RAM user is the resource creator, but not the resource owner.
- **By default, a new RAM user has no permission**
 - A RAM user represents an operator and must be explicitly authorized to perform any operation.
 - A new RAM user has no operation permission by default, and cannot operate on resources in the console or by using APIs until authorization is granted.
- **RAM users are not automatically granted the permissions to the resources they created**
 - The RAM user is not automatically granted any permissions to the created resources unless the resource owner explicitly grants permissions to the user.

2) Authorization policy

Authorization policies are a set of permissions described in a policy language. By attaching authorization policies for a user or a group, the user or all users in the group can acquire the access permissions specified in the authorization policies (access denial takes priority by default).

RAM supports two types of authorization policies: system authorization policies and custom authorization policies.

- **System authorization policies**

System authorization policies is a set of generic authorization policies provided by Alibaba Cloud. This set of policies specifies read-only permission or full permission to different products, such as read-only permission to ECSs, and full permission to ECSs. The system authorization policies provided by Alibaba Cloud can only be assigned, but cannot be edited or modified by users. System authorization policies are automatically updated or modified by Alibaba Cloud.

- **Custom authorization policies**

RAM allows custom authorization policies. Each custom authorization policy mainly includes “effect”, “resource”, “action”, and “condition”. For example, the following fine-grained authorization can be implemented: read-only operations (i.e. action) is allowed (i.e. effect) on OSS sampleBucket (i.e. resource) on conditions that the requester's IP address is 42.160.1.0 and the access time is within the period from 9:00 am to 9:00 pm. Otherwise, access is denied.

6.7.1.6 RAM Role Management

A RAM role can be viewed as a virtual RAM user. It does not have any long-term authentication key (such as logon password or AccessKey). A RAM role must be assumed by an authorized RAM user before it can work properly. RAM roles can be used in scenarios like cross-cloud-account resource authorization, resource access authorization among various cloud services, and issuance of temporary authorization tokens to mobile apps.

Two types of RAM roles are available:

- **User role**

A user role is a type of role that can be assumed by a subuser. The role can be assumed by a subuser under the customer's own Alibaba Cloud account or other accounts. User roles are typically used to solve problems such as cross-account access and temporary authorization.

- **Service role**

A service role is a type of role that can be assumed by a cloud service. A cloud service with a service role can access resources of other cloud services.

Roles differs from RAM users in the way they are used. A RAM role must be assumed by an authorized RAM user. After assuming the role successfully, the RAM user receives an STS security token for this RAM role. Then, the user can use this security token to access the resources authorized to the role.

Please note, after switching to a RAM role identity, the RAM user can only perform operations authorized to this role identity, but the access permissions of the user's real identity upon logon will not be available. When switching back to the logon identity, the RAM user has the access permissions of its own identity, but not those of the role identity.

6.7.1.7 STS (Security Token Service)

Alibaba Cloud Security Token Service (STS) provides RAM users with authorization credentials for short-term resource access. Some scenarios involve users (persons or applications) that do not regularly access the resources of a cloud account but only need occasional access, such users are called "temporary users". For some other users, such as applications running on untrusted mobile devices, it is undesirable to issue them long-term AccessKeys due to the insecure nature of the execution environment. In these cases, STS can be used to issue temporary authorization credentials to these users. When issuing a token, the administrator can define the permissions and expiration time (one hour by default) for the token as needed.

An STS access token is a triplet that includes a security token, an AccessKey ID, and an AccessKey secret. The user passes in the security token and the AccessKey ID to call resource APIs, and uses the AccessKey secret to sign the request. Security tokens issued by STS are not used together with other AccessKeys.

STS eliminates the need to create and manage a long-term RAM user account and an AccessKey for a temporary user or a user with low security level. Moreover, the authorization credentials are automatically issued by STS, therefore they are not embedded in insecure locations such as the client side code. By default, tokens rotate automatically on an hourly basis to improve security.

6.7.1.8 Best Practices for Account Security Management

- **Enable MFA for Alibaba Cloud account and RAM users with high-risk permissions**
It is recommended that the account owner should bind MFA to the account so that multi-factor authentication is performed each time the account is used. If a RAM user is granted with high-risk permissions (such as stopping virtual machines and deleting buckets), MFA should also be bind to the RAM user.
- **Define authorization policies with IP and MFA restrictive conditions for high-risk privileges**
A good practice is to have two persons control the account password and the MFA device separately to ensure that any operation can only be done in the presence of both persons.

- **Separate RAM users by user management, permission management, and resource management (i.e. separation of duties)**

A good system with separation of duties supports checks and balances to minimize security risks. When using RAM, consider creating different RAM users that are separately responsible for RAM user management, RAM permission management, and the management of resource operations under various products.

- **Manage permissions for RAM users using groups**

Normally, it is not necessary to bind an authorization policy to a RAM user individually. It is more convenient to create a group (such as admins, developers, and accounting groups) related to the role and responsibilities of the users, bind an appropriate authorization policy to the group, and then add the related users to the group. All users in a group share the same permissions. Therefore, one can modify the permissions of all users in the group in one operation. When a user changes role within an organization, the admin can simply move the user from one group to another.

- **Use STS to grant short-term access permission credentials to temporary users**

STS makes resource authorization eliminates the need to create and manage a long-term RAM user account and a AccessKey for a temporary user or a user with low security level. Moreover, the authorization credentials are automatically issued by STS, therefore they are not embedded in insecure locations such as client side code. By default, tokens rotate automatically on an hourly basis to improve security.

- **Mandate strong password policy for Alibaba Cloud account and RAM users**

It is recommended to create a strong password policy that requires a sufficiently long minimum length, non-letter characters, and an adequate rotation cycle, for RAM users on the RAM console.

- **Do not create an AccessKey for the Alibaba Cloud account (i.e. primary account)**

An Alibaba Cloud account has full permission to all resources under it. Hence, it is not recommended for the account owner to create/use an AccessKey for the Alibaba Cloud account in order to avoid the disastrous consequences of losing it. Today, the account owner must log on to the Alibaba Cloud console to create an AccessKey for the primary account. This operation requires multi-factor authentication.

- **Rotate logon passwords and AccessKeys regularly**

It is recommended that the account owner and the RAM users rotate logon passwords and AccessKeys regularly. The user can set a password policy to force RAM users to rotate their logon passwords and AccessKeys in a regular cycle.

- **Separate console users and API users**

It is not recommended to create both a password for console operations and an AccessKey for API operations under the same RAM user. Generally, logon passwords should be created for employees, and AccessKey should be created for systems and applications.

- **Enhance security with policy conditions, such as access source IP and time restrictions**

It is recommended to use more restrictive policy conditions when possible for high privilege operations. For example, one can authorize the user Alice the permission to shut down ECS instances with the condition that Alice does it at a specified time and only on the company network.

- **Adjust and revoke RAM user permissions that are no longer needed in a timely fashion**

When a user's role changes and a permission is no longer necessary, the admin should revoke the permission at once. This can minimize the security risk caused by potential leakage of the access credentials of the user.

- **Follow the principle of least privilege and avoid granting excessive permissions**

The principle of least privilege is a primary rule for security design. For example, in an organization, if the responsibilities of the developers group only require reading data stored in an OSS buckets, authorize the group with the read-only permission rather than read-write permissions to the OSS resources. As a rule of thumb, a permission to access all resources of every product should not be granted.

6.7.1.9 Typical Authorization Management Scenarios

- **Enterprise employee sub-account management and permission assignment**

Cloud account A represents an enterprise's super administrator who has full permission to all the cloud products the enterprise has purchased. To minimize uncontrollable risks resulting from leakage of Alibaba Cloud account passwords or AccessKeys, we do not recommend that employees directly use an Alibaba Cloud account (root permission) for daily resource operations. It is better to activate the RAM service and assign different subuser accounts to employees and allocate different permissions to the subusers. Follow

the principle of least privilege to ensure the best possible alignment of power and responsibility and reduce the information security risk associated with the enterprise's migration to cloud.

- **Cross-enterprise (tenant) resource authorization management**

Suppose Alibaba Cloud account A and Alibaba Cloud account B are two enterprises. A has purchased various resources (such as ECS, RDS, SLB instances, and OSS buckets) to support its businesses. Enterprise A wants to focus on the research and development of its business systems, so it authorizes Enterprise B to perform the tasks of maintaining, monitoring, and managing cloud resources. Enterprise B can further allocate permissions for the resources of A to one or more employees for management, and B should precisely control the permissions assigned to its employees who operate on the resources of A. If the maintenance entrustment relationship between A and B is terminated, Enterprise A can revoke the permissions assigned to Enterprise B at any time. With RAM cross-account role authorization, Enterprise A can manage its authorization to Enterprise B easily.

- **Temporary permission token management for mobile applications**

Enterprise A has developed a mobile application. Mobile applications generally run on untrusted user devices which are not under control of A. If the application needs to operate on the cloud resources of A, how to ensure the security of the application is a great challenge. It requires that the application must not store long-term AccessKeys because the application runs on untrusted devices which can be fully controlled by potential attackers. Therefore, it requires that the authorization protocol supports granting the application minimum temporary permission. Even if the temporary permission token is controlled by malicious users, the potential loss/impact can be minimized. To solve the problem, A can use RAM role token management to grant different minimum temporary permissions to apps running on different devices.

- **Manage access permissions between different cloud services**

Alibaba Cloud platform supports selling multiple cloud products, such as ECS (Elastic Computing Service) and OSS (Object Storage Service). On the cloud platform, cloud resources between different cloud products are completely isolated. Without authorization from the resource owner (the customer who has purchased the cloud resources), no party (including the cloud product itself) has the permission to operate on

the customer's cloud resources. For example, Enterprise A has purchased cloud products ECS and OSS. If the applications deployed by A on ECS need to access the data on OSS, by default no operation permission is given unless A grants explicit authorization to allow its ECS instances to access OSS data. With RAM, the customers can authorize a service to operate on the resources of another service to ensure that the operation permissions for all cloud resources are fully controlled by the resource owners.

6.7.2 Key Management Service

Key Management Service (KMS) is a secure and easy-to-use management service provided by Alibaba Cloud. With KMS, users will no longer have to spend excessively to protect the confidentiality, integrity, and availability of keys. Instead, KMS securely and conveniently manages keys for the users.

KMS has implemented multiple strict security measures to ensure data security. The key management infrastructure of Alibaba Cloud conforms to the recommendations in (NIST) 800-57 and uses cryptographic algorithms that comply with the (FIPS) 140-2 standard.

6.7.2.1 Authentication and Resource Access Management

- **Client authentication**

Like other Alibaba Cloud services, KMS requires users to use AccessKey ID and AccessKey Secret to authenticate requests (based on HMAC algorithm). The server ensures request integrity by verifying its verification code. KMS uses HTTPS protocol, so the client can verify the KMS service identity by verifying the server certificates. The data confidentiality of the communication between the client and the server is protected by the HTTPS protocol.

- **Resource Access Management**

KMS supports RAM for resource access management.

6.7.2.2 Secure Channels

In KMS, all the internal communications of Key Management Service use two-way authentication TLSv1.2 protocol to securely protect the communication of internal nodes.

6.7.2.3 Data Security

- **Domain Master Key (DMK)**

The Domain Master Key in KMS are the core master key of the service. It is held and used by a dedicated distributed system (i.e. Virtual HSM, or VHSM) and its plain texts only exist

in the memory of VHSM nodes and the encrypted communication channels between them.

VHSM nodes use hardware devices to encrypt and protect the data stored locally. VHSM regularly rotates DMK. After rotation, the old DMK are only used to decrypt the customer master keys generated previously.

- **Customer Master Key (CMK)**

Users can create and manage Customer Master Keys using APIs or the console of KMS service. CMKs cannot be exported from the KMS service. When CMKs are generated on VHSM, they are encrypted by the DMK and the relevant context data (for example, the owner of a CMK and the CMK ID) is also protected. The encrypted CMKs are stored in a highly reliable storage service.

When a user uses a CMK, the KMS service retrieves the ciphertext CMK and enters it into the VHSM. The VHSM decrypts the CMK and performs the operations requested using the plaintext CMK. In the whole process, the plaintext CMK is only stored in the VHSM memory, and is released immediately after use.

6.7.2.4 BYOK

With Alibaba Cloud Key Management Service, users can also import keys from their own Key Management Infrastructure (KMI) and achieve a greater control over the keys used to protect critical business information. This scenario is commonly referred to as Bring Your Own Key (BYOK).

With BYOK, a user can create a CMK without KMS generating the key material and then import the user's own key material into this CMK. When an imported key material is used, the user controls the generation, lifecycle and durability of the associated CMK while allowing KMS to use a copy of it. The user may choose to do this for one or more of the following reasons:

- To generate the key material with a known entropy that meets the user's business or compliance requirements.
- To lease the key material managed by user's KMI to KMS, only temporarily or periodically. This can be achieved by setting an expiration time for the imported key material or

manually deleting it without a waiting period. The user may optionally re-import the key material again in the future if needed.

- To own the original copy of the key material in the user's KMI and to take better control over the durability of it for additional disaster recovery.

Please note that the key material imported must be a 256-bit symmetric encryption key.

6.7.2.5 Management and Maintenance Security

In addition to complying with the maintenance security protocols of Alibaba Cloud, KMS implements a multi-user control mechanism. For sensitive operations (including but not limited to certificate issuance for internal nodes, security sensitive API calling, etc.), passwords from multiple users must be entered. After the passwords are verified, the operation can be performed.

6.7.2.6 Best Practices

- **Encryption and decryption using KMS**

A user can directly call the KMS API, and use a specified CMK to encrypt and decrypt data directly. This scenario applies to encryption and decryption of a small amount of data (less than 4 KB). Data is transmitted to the KMS server by using secure channels, encrypted or decrypted at the server, and returned by using secure channels.

- **Envelope encryption**

Envelope encryption is a mechanism that allows users to store, transfer, and use encrypted data by encapsulating its data keys in an envelope. Instead of using CMKs to directly encrypting and decrypting data. This is recommended when a large volume of data needs to be encrypted or decrypted. To use envelope encryption, a user can directly call the KMS API, generate a data key, use a specified CMK to protect the data key, and use the data key for data encryption and decryption operations.

6.7.3 ActionTrail

Alibaba Cloud provides the ActionTrail service, which enables a unified log management for cloud resources. The ActionTrail service records user login and resource access operations under each Alibaba Cloud account. Such record includes the operator, operation time, source IP address, resource object, operation name, and operation status. ActionTrail

provides operation record query, and saves record files to the specified OSS bucket. With all operation records saved by ActionTrail, users can perform security analysis, intrusion detection, resource tracking, and compliance audit.

ActionTrail collects API calling records of cloud services (including API calling records triggered by using the console). It standardizes the operation records and saves them to specified OSS buckets as files. Users can manage the records files using all the management functions provided by OSS, such as authorization, enabling lifecycle management, and archiving management. Meanwhile, users can also use OSS data encryption and permission management functions to ensure the data security of the event records. ActionTrail supports operation event query from such dimensions as username (i.e. operator), operation time, source IP address, resource object, operation name, and operation status, and it can help to diagnose problems quickly or track security incidents.

Generally, when a user initiates an operation calling using the console or SDK, ActionTrail transfers the operation records to the specified OSS Bucket within 5 minutes. Users can view the operation records for the last 7 days using the ActionTrail console. For older records, users need to access the OSS bucket set for ActionTrail directly.

ActionTrail can be mainly used for the following scenarios:

- **Security analysis**

Logs recorded by ActionTrail can be used for security analysis for any potential security issues. For example, ActionTrail records all account logon operations, including detailed records such as the logon time, which IP was used, whether multi-factor authentication logon was used, etc. With these records, Users can determine whether their accounts has any security issues.

- **Resource change tracking**

When any cloud resources were changed unexpectedly, the operation logs recorded by ActionTrail can help users identify how the changes took place. For example, when a user noticed that an ECS instance had stopped, with the help of ActionTrail, the user can find out who initiated the operation, from which IP, and at what time, etc.

- **Compliance audit**

If a user's organization has multiple members and the Alibaba Cloud RAM service is used to manage the identities of members, it is generally required for the user to obtain the detailed operation records of each member to meet the compliance auditing requirements. The operation events recorded by ActionTrail can meet these compliance auditing requirements.

6.7.4 CloudMonitor

CloudMonitor monitors the Alibaba Cloud resources and Internet applications. The functions of CloudMonitor are to collect monitoring metrics for Alibaba Cloud resources, to monitor network connectivity, and to set alarms for the monitoring metrics.

CloudMonitor can monitor ECS, ApsaraDB, Server Load Balancer and other types of Alibaba Cloud service resources. Moreover, it monitors Internet application availability using common network protocols such as HTTP and ICMP. CloudMonitor gives users a comprehensive understanding of the usage, performance, and operational states of Alibaba Cloud resources.

Users can set different alarm rules for the alarm service. An alarm message is sent when certain metric data reaches an alarm threshold based on the preset rules, enabling users to make quick responses.

6.7.4.1 Resource Access Management

CloudMonitor supports RAM and allows users to control permissions for Cloud Service Monitoring metric data, alarm rule management, alarm contact and alarm contact group management for RAM users. Additionally, it supports time, MFA, and IP authentication types.

7. Alibaba Cloud Security

Built on Alibaba Group's security technologies and experiences over the years, Alibaba Cloud Security provides customers with one-stop security services. All Alibaba Cloud Security products support the RAM service. Users can create and authorize RAM subusers with the RAM service.

Please note that the products covered in this section are what's available from Alibaba Cloud's international offering today. Please refer to our website alibabacloud.com for a complete list of security products and more details.

7.1 Basic Protection

7.1.1 Anti-DDoS Basic

By default, Alibaba Cloud provides up to 5 Gbps Anti-DDoS capability free of charge.

Alibaba Cloud provides free Anti-DDoS service to a certain extent for all users. Please find the Anti-DDoS basic feature explained in product specification as it may vary in different regions.

7.1.2 Best Practices

The Anti-DDoS Basic's traffic threshold can be set automatically or manually. If you select automatic configurations, the system dynamically adjusts the cleaning threshold value based on ECS's traffic size. If you select manual configurations, you can set the thresholds of the traffic size and PPS (packets per second). Alibaba Cloud Security enables traffic cleaning immediately when the threshold value is exceeded. The threshold value needs to be higher than the normal traffic size. However, if the threshold value is set too high, it cannot provide effective protection as the Anti-DDoS feature would be not triggered. Likewise, if the threshold value is too low, the traffic cleaning triggered by Anti-DDoS may affect normal traffic access. Additionally, the maximum traffic threshold value is directly tied to the network traffic handling capabilities of all other Alibaba Cloud products that a customer had purchased. This is done to prevent the threshold value to be set too high thus render the Anti-DDoS protection non-affective.

7.2 Advanced Protection

7.2.1 Anti-DDoS Pro

Alibaba Cloud Security's Anti-DDoS Pro is a paid service to address the problem of service interruptions caused by high-volume DDoS attack on Internet servers (including non-Alibaba Cloud hosts). The attack traffic will be mitigated by the service and ensure the stability and reliability of the servers under protection.

Alibaba Cloud Security's Anti-DDoS Pro provides the following features and advantages:

- **Full coverage of common DDoS attacks**

Alibaba Cloud Security's DDoS cleaning system can protect Alibaba Cloud users against various types of DDoS attacks targeting the network layer, transport layer, or application layer (including HTTP/S Flood, SYN Flood, UDP Flood, UDP DNS Query Flood, (M)Stream Flood, ICMP Flood, HTTP Get Flood, and all other types of DDoS attacks).

- **Five second auto protection**

Alibaba Cloud Security's DDoS cleaning system applies world-class detection and protection technologies, which enable the Anti-DDoS Pro to complete attack discovery, traffic redirection, and traffic cleaning in five seconds, thus greatly reducing the network jitter. To identify DDoS attacks, the system triggers the protection when traffic meets threshold, thus ensuring the service availability in case of a DDoS attack.

- **High scalability and redundancy**

Each basic unit in Alibaba Cloud Security's DDoS cleaning system can filter upto 10 Gbps attack traffic. Powered by the high scalability and high redundancy of the cloud computing architecture, the anti-DDoS system supports seamless scale-up in the cloud environment to achieve the highly scalable anti-DDoS capability.

7.2.1.1 Best Practices

To activate Anti-DDoS Pro, a user must go to the DNS service provider and update the DNS record in order to resolve domain names to the Anti-DDoS CNAME or A Record assigned. Although both CNAME and A record are supported for rerouting traffic, it is recommended to use CNAME record when possible. Additionally, a user should also configure the Anti-DDoS Pro to use the CNAME of the origin site domain.

After the domain name resolution is modified, all public network traffic is routed to the Anti-DDoS Pro Cleaning Center. The incoming traffic is forwarded to the origin server IP by Anti-DDoS Pro using the port/protocol forwarding. The malicious attack traffic is cleaned and filtered in the Anti-DDoS Cleaning Center before returning to the origin server IP, which helps ensure the continuity and stability of the origin server services.

Anti-DDoS Pro is fully compatible with Alibaba Cloud Security WAF (Web Application Firewall) and Alibaba Cloud CDN (Content Delivery Network). Therefore, the best architecture is to combine Anti-DDoS Pro, CDN, and WAF to provide protection and acceleration for the origin server deployment:

- Anti-DDoS Pro (entry-level anti-DDoS) -> CDN (cached resource acceleration) -> WAF (application layer protection) -> origin server (ECS/SLB/VPC/IDC...)

7.2.2 Mobile Security

Mobile Security provides security services for the full life cycle of mobile apps. Mobile Security can precisely detect potential risks such as security vulnerabilities in apps, malicious codes, and fake apps. By using features such as app hardening and security components, it would significantly improve apps' anti-cracking and anti-reverse engineering capabilities.

Mobile Security offers the following features:

1) Vulnerability scan

This feature scans Android apps to quickly locate vulnerabilities, and provides remediation solutions for any vulnerabilities detected. Static and dynamic scanning methods are combined to discover as many security vulnerabilities as possible:

- Static scanning uses taint analysis to precisely backtrack variable values, thus analyzes and tracks vulnerabilities at the register granularity.
- Dynamic scanning uses fuzz testing to detect potential vulnerabilities in the Android environment with accurate results.

2) Malicious code scan

The feature is designed for enterprise users who use third-party plug-ins, subcontract app development to a third party, or who serve as app distribution channels. The feature scans Android apps to precisely detect any malicious codes embedded in the apps.

- The Alibaba Cloud self-developed malicious code scanning engine has received the perfect score numerous times in the renown AV-Test assessments.

- The feature also adopts advanced machine learning algorithms and big data analysis to extract code attributes automatically in order to detect any malicious codes.

3) Fake app detection

The feature is used for brand risk identification. It can detect the distribution channels of fake Android apps to help you prevent the distribution of such apps and minimize the damages to your brands. The detection feature applies to:

- More than 300 app distribution channels worldwide.
- Non-typical channels such as cloud storage, forums, enterprise websites, and phishing scams.

4) Application hardening

This feature enhances the anti-cracking capabilities of Android apps by using recompilation, shelling, and modifying the command calling sequence. Common hardening features tend to emphasize on the reinforcement intensity, making apps unusable after hardening. In contrast, Alibaba Cloud Security's app hardening feature focuses equally on the reinforcement intensity and compatibility. The application hardening feature includes:

- Defense against common static analysis tools: effectively prevents hackers from using static analysis tools such as APKTool, dex2jar, and JEB to analyze Java-layer codes of apps.
- SO shelling: shells SO files to effectively prevent malicious users from using tools such as IDA and readelf to analyze logics in SO files.
- DEX shelling: DEX files are shelled to prevent against different types of analysis tools. Such as preventing hackers from dumping the memory of Java-layer codes.
- Constant encryption: encrypts plaintext constant strings in DEX files and uses the decryption functions to dynamically decrypt the strings at run time, greatly increasing the difficulty of reverse analysis.
- Java command translation: modifies the call chain of the Java-layer logic. In this way, even if hackers get the Java-layer code, they still cannot completely analyze the business logic.
- Java execution simulation: detaches the instructions in DEX files and simulates their execution in a custom execution environment, effectively preventing malicious users from dumping Java-layer codes at the instruction level.

7.2.3 Web Application Firewall (WAF)

Alibaba Cloud Web Application Firewall (WAF) is a SaaS-based web application security service which detects illegal web requests through its built-in security strategy. Web Application Firewall (WAF) filters out massive numbers of malicious accesses by defending against SQL injection, XSS, common web server plug-in vulnerabilities, trojan uploads, unauthorized access, and other common OWASP attacks to prevent the leakage of website assets and data and ensure website security and availability.

WAF Feature	Sub-Features	Description
Service configuration	Protocols supported	<ul style="list-style-type: none">• Supports web security protection for HTTP and HTTPS (available in the Premium version) traffic.
Web application protection	Defense against common web application attacks	<ul style="list-style-type: none">• Defense against common OWASP attacks, including SQL injection, XSS, Webshell uploads, command injection, illegal HTTP protocol requests, common web server vulnerability attacks, unauthorized access, path traversal, etc. Also includes backdoor isolation and scan protection.• Website stealth: Your website address is not exposed to attackers, so attack packets cannot bypass the WAF to attack your website directly.• Regular 0-day patch updates: Latest vulnerability patches are provided to global users simultaneously to secure websites.• Friendly observation mode: With observation mode enabled for new website services, possible attacks matching the protection rules trigger warnings but are not blocked.• Prevention Mode: Actively blocks intrusions and attacks detected by its set rules. Attackers' requests are denied and their connections are terminated. This mode continues to log all attacks in the WAF logs file.

WAF Feature	Sub-Features	Description
Web application protection	Protection against malicious HTTP/S attacks	<ul style="list-style-type: none"> Controls the frequency of access from a single source IP addresses, provides redirect jump verification, and determines if access requests come from a human or machine. Identifies massive and slow request attacks based on the statistics of response codes, URL request distribution, abnormal referers, and user-agent features, and works with the precise website protection rule to provide comprehensive protection. Fully utilizes Alibaba Cloud's advantages in big data security to establish threat intelligence and trusted access analysis models. This allows you to quickly identify malicious traffic.
Web application protection	Precision access control	<ul style="list-style-type: none"> Provides a friendly configuration console interface and supports conditional combinations of IP, URL, Referer, User-Agent, and other common HTTP fields. This allows you to create powerful yet precise access control policies. Establishes comprehensive multi-layer protection with the security modules for protection against common Web attacks and HTTP/S Flood attacks by precisely distinguishing between trusted and malicious traffic based on your needs.
Web application protection	Virtual patches	<ul style="list-style-type: none"> Adjusts web protection policies to provide quick protection before actual patches are released.
Management	Attack event management	<ul style="list-style-type: none"> Supports centralized management and statistics logging of attack events, traffic sizes, and scales.
Reliability	Load balance support	<ul style="list-style-type: none"> Provides the service in clusters, balancing loads on multiple servers (multiple load balancing policies supported).
Reliability	Scalability	<ul style="list-style-type: none"> Reduces or increases the number of machines in a cluster based on actual traffic, thus enabling elastic scaling.

WAF Feature	Sub-Features	Description
Reliability	No single point of failure	<ul style="list-style-type: none">• Downtime or maintenance of a single machine does not affect the service availability.

As a cloud firewall service, WAF modifies your website's DNS records, so that all requests through WAF are detected, filtered, and cleaned in order to direct safe traffic to the site server and prohibit attacks from reaching the server.

WAF also filters out large numbers of malicious access attempts and alleviates the performance impact of HTTP/HTTPS flood attacks on servers.

7.2.3.1 Best Practices

- **Configure an ECS security group or an SLB whitelist for the origin server to prevent the origin server being exposed**

Note: Origin server protection is not mandatory. Normal service forwarding is not affected if no configuration is made. However, without such protection, attackers can bypass WAF protection to attack the origin server directly when the origin server IP is known.

- **Prevent leakage of sensitive information with WAF**

The anti-leakage feature mainly covers the leakage of sensitive information on websites, especially it provides filtering of sensitive information such as mobile phone numbers, ID numbers, and credit card numbers, etc. This feature provides defense against scenarios such as unauthorized access of websites, and malicious crawling of sensitive information on websites.

- **Prevent WordPress reflection attacks with WAF**

The users of WAF Premium (and above) Version can prevent WordPress reflection attacks effectively with precision access control rules set in place.

7.2.4 Server Guard

Server Guard is a security agent for ECS instances. Through the lightweight agent installed on ECS, Server Guard works together with the cloud security center to provide remote login alert, brute-force cracking detection and webshell detection and removal capabilities.

- **Remote logon alert**

Server Guard records all logon operations and provides real-time reminder for logon attempts from uncommon locations. Users can configure common logon locations as needed.

- **Brute-force cracking detection**

Server Guard detects brute-force type of attempts to crack a user's password and reports such events to Alibaba Cloud for further interception. It prevents attackers from cracking one's password through repeated guesses.

- **Webshell detection and removal**

Powered by webshell scanning and removal engine developed by Alibaba Cloud, Server Guard supports both local and on-the-cloud scanning and one-click removal of webshell files. Additionally, it supports scheduled scanning and removal policies. Common webshell file types including php and jsp are supported.

7.2.5 Alibaba Cloud SSL Certificates Service

Alibaba Cloud SSL Certificates Service allows customers to directly apply, purchase and manage SSL certificates on Alibaba Cloud. This service is offered in cooperation with qualified certificate authorities. Customers can select a third party certificate authority and its certificate products to implement full-site HTTPS security solutions, thus making customers' websites credible and being protected from hijacking, snooping, and tampering, etc. The service also provides integrated lifecycle management on the cloud to simplify the certificate deployment, so that you can distribute the certificates to the cloud products with one click.

SSL Certificates Service provides the following features:

- HTTPS-secured websites are built to encrypt the communications between users and the websites to ensure that the information displayed to the user is reliable and the websites are protected from hijacking, snooping, and tampering.
- Digital certificates issued by trusted CAs are provided. Different levels of digital certificates are issued after the review and authentication by CAs.
- Certificate lifecycle management are provided for the users to manage digital certificates of multiple channels in a unified manner. The users can view their cloud

service certificates and manage their certificate orders on a unified platform.

- The users can deploy digital certificates to other activated Alibaba Cloud products (such as CDN, SLB, Anti-DDoS, and WAF) with a single click on the Alibaba Cloud platform, enabling a simplified certificate deployment.
- Digital certificates can be securely revoked according to standard revocation procedures reviewed by CAs.

7.2.5.1 Best Practices

By deploying SSL digital certificates to web servers or other resources (such as SLB, CDN, etc.), customers can enable HTTPS-secured access to their website, thus significantly improving the website security, even in public environments such as airports, coffee shops, and Internet cafes to minimize the risks of sensitive information leakage.

A simplified certificate purchase process is provided to Alibaba Cloud customers. Secure certificate and key storage solutions are also provided.

The certificates in the Alibaba Cloud SSL Certificates Service are connected with Alibaba Cloud products, which allows a simplified one-click deployment of digital certificates to Alibaba Cloud products.

7.2.6 Managed Security Service

Alibaba Cloud Managed Security Service delivers WAF and Anti-DDoS management and security assessment services. With the managed service, potential losses caused by security breaches are reduced and customer's online business is constantly monitored and protected from incoming attacks.

Managed Security Service is suitable for customers who have purchased Alibaba Cloud's Anti-DDoS and WAF products but lack continuous monitoring capabilities and security engineering resources to defend against vulnerabilities. The service is therefore ideal for customers who would otherwise need to outsource professionals to assist in ensuring effective security service operations.

Managed Security Service provides customers a completely managed service for Alibaba Cloud WAF and Anti-DDoS, including product deployment, security

monitoring, policy optimization, security incident response service and security report service.

The management security service provides the following features:

- **Setup service**
 - Create WAF/Anti-DDoS implementation plan
 - Set up WAF/Anti-DDoS for your online web applications
 - Configure and enable HTTPS certificate to your website
 - Implement WAF/Anti-DDoS with other products, including Alibaba Cloud CDN
- **Monitoring service**
 - Continuous monitoring and detection of attacks targeting your system
 - Monitor the impact of security attacks on your system
 - Analyze web attack log to identify potential threats
- **Policy optimization**
 - Introduce tighter WAF policies by reviewing monitoring results and security test results
 - Optimize WAF policies to protect your website from HTTP floods, SQL injections, cross-site scripting (XSS), and bot attacks
 - Adjust Anti-DDoS policies to reduce the impact of attacks
 - Modify WAF/Anti-DDoS policies based on business requirements
- **Incident response**
 - Quick response to security incidents such as high volumn DDoS attacks
 - Experienced security experts help to respond to attacks
- **Reports**
 - Supports daily, weekly, and monthly reports
 - Supports customized reports according to customer needs

Additionally, a security assessment service is also provided in order to help customers to perform the following scans and tests:

- **Services scan:** Detects open port and port banner, and identify unnecessary services.
- **System vulnerability scan:** Detects CVE vulnerabilities, weak passwords, and system misconfigurations.
- **Web application scan:** Detects OWASP Top 10 Web Vulnerabilities.
- **Manual test:** Manually analyze target systems and find vulnerabilities courtesy of our security experts.

8. Alibaba Cloud Security Ecosystem

In the spirit of open resource and cooperation, Alibaba Cloud collaborates with security partners to establish Alibaba Cloud Security Industry Ecosystem and provide customers with industry-leading security solutions that are consistent with their existing deployed security control measures.

Alibaba Cloud marketplace has provided customers with such security solutions as VPN, next-generation firewall, WAF, etc.

9. Version History

April 2018: International Edition - Version 1.0