Nigerian FinTech's Cloud Adoption White Paper

Issue 1.0

Date 2023-05-26



HUAWEI TECHNOLOGIES CO., LTD.

Copyright © Huawei Technologies Co., Ltd. 2023. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademarks and Permissions

 and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: https://www.huaweicloud.com/

DEFINITIONS

The words hereunder shall have the meaning stated against them.

ACL	Network Access Control List
API (s)	Application Programming Interface
AZ	Application Zone
BIOS	Basic input/output System
BMC	Baseboard Management Controller
CapEx	Capital Expenditure
CBN	Central Bank of Nigeria
CI/CD	Continuous Integration and Continuous Delivery/ Continuous Deployment
CVE	Common Vulnerabilities and Exposures
DDOS	Distributed Denial of Service
DJCP MLPS	The "Classified Protection of Cybersecurity", also known as the "Multi-Level Protection Scheme"
DR	Disaster Recovery
EIP	Enterprise Integration Pattern
Fintech	Financial Technology
GPU	Graphics Processing Unit
HA	High Availability
KPI(s)	Key Performance Indicator(s)
I/O	Input/Output
ISO27001	International Standard for Information Security
IT	Information Technology
NDPR	Nigerian Data Protection Regulation
NDRP	National Disaster Recovery Plan
OpEx	Operating Expenditure
SDK	Software Development Kit
SLA	Service Level Agreement

SME(s)	Small and Medium sized Enterprise(s)
O&M	Operations and Maintenance
PCI/DSS	Payment Card Industry Data Security Standard
PoLP	Principle of Least Privilege
R & D	Research and Development
RDS	Remote Desktop Services
TPS	Transaction Processing System
VM	Virtual Machine
VPC	Virtual Private Cloud
VPN	Virtual Private Network

Contents

1 Introduction 1
1.1 What Is FinTech? 1
1.2 Nigerian FinTech Landscape 1
1.3 Purpose of This White Paper 2
2 Nigerian FinTech's Cloud Adoption 3
2.1 Cloud Computing Overview 3
2.1.1 What Is Cloud Computing? 3
2.1.2 Cloud Computing Deployment Models 3
2.1.3 Cloud Computing Service Models 4
2.2 Cloud Computing Trends in Nigeria 5
2.3 Cloud Adoption in Nigerian FinTech Industry 6
3 Significance of Nigerian FinTech's Cloud Adoption 7
3.1 Challenges Facing Nigeria's FinTech 7
3.2 Public Cloud Facilitates Cloud Transformation of FinTech 8
3.2.1 Regulatory Compliance 8
3.2.2 Cloud Security 12
3.2.3 Cloud Native 18
3.2.4 Operational Costs 25
4 Cloud Adoption Framework 29
4.1 Plan 29
4.2 Landing Zone Construction 30
4.2.1 Why Is the Landing Zone Required? 30
4.2.2 Landing Zone Reference Architecture 31
4.2.2.1 Account and Organizational Structure 32
4.2.2.2 Accounting Management 34
4.2.2.3 Network Planning 35
4.2.2.4 Identity and Permissions Design 37
4.2.2.5 Compliance Audit 39
4.3 Cloud Architecture Design 40
4.3.1 Architecture Design Objectives and Principles 40
4.3.2 HA 40
4.3.3 Scalability 44
4.3.4 Performance 46
4.3.5 Security 46
4.3.6 Cost 49
4.4 Adopt 49

4.4.1 Application Migration to the Cloud 49
4.4.1.1 Rehost 51
4.4.1.2 Replatform 53
4.4.1.3 Rearchitect 55
4.4.1.3.1 Traditional Application Architecture Issues 55
4.4.1.3.2 Cloud Native Application Architecture 55
4.4.1.3.3 Application Reconstruction 56
4.4.1.4 Migration Implementation 57
4.4.1.4.1 Migration Team Establishment 57
4.4.1.4.2 Migration Assurance 58
4.4.1.4.3 Testing and Verification 59
4.4.1.4.4 Service Cutover and Rollout 59
4.4.2 Data Migration to the Cloud 63
4.4.2.1 Typical Data Lake Scenarios 63
4.4.2.1.1 Data Warehouse and Report Analysis 63
4.4.2.1.2 Integrated Solution for Streaming and Batch Processing and Query 64
4.4.2.1.3 Real-Time Import of Data from Databases to a Data Lake 64
4.4.2.1.4 Real-Time Import of Messages and Logs to a Data Lake 65
4.4.2.2 Big Data Migration 66
4.4.2.2.1 Process 66
4.4.2.2.2 Migration Solution Design 67
4.4.2.2.3 Example Migration Plan and Period 68
4.4.2.2.4 Big Data Migration Tools 68
4.5 Govern and Manage 70
4.5.1 Cost Management 70
4.5.2 Security Compliance and Governance 71
4.5.2.1 Security Management Organization 71
4.5.2.2 Personnel Security Management 74
4.5.2.3 Security PMI 76
4.5.2.4 Account Security Management 77
4.5.3 Cloud-based O&M 79
4.5.3.1 Trends and Challenges 79
4.5.3.2 Multi-Dimensional O&M 80
4.5.3.3 Backup and Restoration 86
4.5.3.4 Change Management 88
4.5.3.4.1 Roles and Responsibilities 89
4.5.3.4.2 Change Classification 89
4.5.3.4.3 Change Process 90
4.5.3.4.4 KPIs 91
4.5.3.5 Emergency Handling 92
4.5.3.6 O&M Service 94

Best Practices 96
5.1 Customer Background 96
5.2 Solutions 96
5.2.1 One Global Network 96
5.2.2 Cloud Native 96
5.2.3 One-Stop Big Data 97
5.2.4 Migration Team and Solutions 97
5.3 Customer Benefits 97
6 About Huawei Cloud 98
7 FINTECH ASSOCIATION OF NIGERIA 99

1 Introduction

1.1 What Is FinTech?

FinTech is the combination of innovative business models and technology to enable, enhance, and disrupt financial services. It refers to the use of new forms of technologies like the blockchain, mobile payment, artificial intelligence and machine learning to disrupt traditional financial services models and improve efficiency. Software and computer algorithms are used to help financiers and consumers better manage their finances, enjoy an improved user experience, and improved lives.

1.2 Nigerian FinTech Landscape

The Nigerian FinTech industry was in its infancy a few years ago, but has since matured, becoming a major sector in Africa. Nigeria FinTech is gradually gaining international recognition and is increasingly embraced by consumers with stronger collaboration between FinTech and existing enterprises.

The COVID-19 pandemic and the impact it had on the Nigerian economy has created an opportunity for FinTech to contribute to economic growth. For this to happen, entrepreneurs need to be able to scale-up and improve their value propositions and deliver on Nigeria's FinTech potential.

There are six broad FinTech segments.

FinTech Segment	Category	Value Proposition
Payments, Mobile Money & digital Banking	• Digital payments • Mobile money & wallets • Payments processing & switching • Digital banks & accounts • Global transfers & remittances • Terminal & infrastructure providers	Provide payments and accounting solutions for consumers and financial institutions. This includes mobile money and wallets, digital banking, payment processors, and payments back-end infrastructure providers.
Lending	• Retail lending	Provide loans to consumers and SMEs

FinTech Segment	Category	Value Proposition
	• SME lending • Credit data analytics & loan assessments	online. FinTech lenders assess borrowers' credit worthiness and automate the underwriting process. Some FinTechs in this segment also provide lending platforms.
Savings, investment & crowdfunding	• Digital wealth & asset management • Automatic savings platforms • Alternative investments • International investment platforms • Crowdfunding	Provide solutions to help consumers save money, and help private and institutional investors to buy, sell, and manage assets and securities. FinTechs in this segment may also provide crowdfunding platforms.
Enterprise services & infrastructure	• Financial solutions • Credit infrastructure • APIs and connectors	Provide enterprise software and solutions for the financial industry. It includes tools for managing accounts, and finance processes. Providers here also provide core infrastructure and platforms to enable and enhance banks' services.
Cryptocurrency	Cryptocurrency exchanges & wallets	Provide access to digital cryptocurrency markets and exchanges to enable users to pay or accept cryptocurrency.
InsurTech	• Digital insurers • Insurance comparison services • Digital agents and distribution platforms • Insurance system providers	Provide access to insurance for consumers either directly or through marketplaces. It also provides new technologies to enable traditional insurers.

1.3 Purpose of This White Paper

FinTech pursues innovation, efficiency, security, and reliability. Cloud-based infrastructure is well suited to supporting these four goals. This white paper aims to share Huawei Cloud experience in cloud deployment and operations with the FinTech industry and provide guidance for FinTech enterprises on cloud transformation.

2 Nigerian FinTech's Cloud Adoption

2.1 Cloud Computing Overview

2.1.1 What Is Cloud Computing?

Cloud computing is a revolutionary technology that has transformed the way businesses and organizations access and use computing resources. It allows users to access and use shared resources, software, and information over the Internet, rather than having to manage and maintain these resources themselves.

One of the main benefits of cloud computing is how it can reduce costs and increase efficiency. By using cloud services, businesses and organizations can access advanced technology and resources without having to invest in expensive hardware and infrastructure. This allows them to stay focused on their core operations and drive growth, while also reducing their operating expenses.

Cloud computing also offers more flexibility and scalability than traditional computing models. Users can scale their use of resources up or down as needed, which helps to optimize costs and quickly respond to changes in demand or market conditions. This can be particularly useful for businesses that experience seasonal fluctuations or rapid growth.

In addition to these benefits, cloud computing offers enhanced security and reliability. Cloud vendors typically have advanced security measures in place to protect their systems and data, which can provide better security than what an organization could achieve on its own. Similarly, cloud services are often more reliable than on-premises systems, as they are typically backed by redundant infrastructure and built to withstand outages and other disruptions.

Overall, cloud computing and cloud technology are transforming the way businesses and organizations operate and compete in the digital age. By providing access to advanced technology and resources, as well as improved efficiency, flexibility, and security, cloud computing is helping businesses and organizations to innovate and succeed in today's fast-paced and constantly evolving world.

2.1.2 Cloud Computing Deployment Models

Cloud computing deployment models refer to the way in which cloud services are delivered to customers. There are three main deployment models: public, private, and hybrid.

- Public cloud

Public cloud deployment involves using cloud services that are provided by third-party service providers and made available to the general public over the Internet. These services are typically accessed on a pay-per-use basis, and customers do not need to purchase any hardware or software. Public cloud deployment is the most popular option because it is cost-effective, easy to use, and requires minimal maintenance.

- Private cloud

  Private cloud deployment involves using cloud services that are hosted on an enterprise's own infrastructure. This option is suitable for businesses that have strict security and compliance requirements and want to have more control over their data and resources. Private cloud deployment can be more expensive than public cloud deployment because it requires the purchase and maintenance of hardware and software.

- Hybrid cloud

  Hybrid cloud deployment involves using a combination of public and private cloud services. This option allows businesses to take advantage of the benefits of both deployment models. For example, a company may use a public cloud for storing and processing non-critical data, while using a private cloud for storing and processing sensitive data. Hybrid cloud deployment can be more complex to set up and manage than the other deployment models, but it offers greater flexibility and scalability.

- Community Cloud

  A community cloud is a cloud computing model where the cloud services are shared by a group of organizations with similar interests or requirements. For example, a community cloud could be used by healthcare organizations to store and share patient data securely. In a community cloud, the infrastructure and resources are shared among the members of the community.

In conclusion, cloud computing deployment models refer to the way in which cloud services are delivered to customers. There are three main deployment models: public, private, and hybrid. Each deployment model has its own set of benefits and drawbacks, and the best option for a business will depend on its specific needs and requirements.

2.1.3 Cloud Computing Service Models

Cloud computing service models refer to the different types of cloud services that are available to customers. There are three main service models, each defined by the different types of resources provided: infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS).

- Infrastructure as a service (IaaS)

  Infrastructure as a service (IaaS) is the most basic service model. It involves providing customers with access to infrastructure resources such as servers, storage, and networking. Customers can use these resources to build and run their own applications and services. IaaS is suitable for businesses that want to have more control over their IT infrastructure and want to take responsibility for managing and maintaining it.

- Platform as a service (PaaS)

  Platform as a service (PaaS) is a higher-level service model that provides customers with access to a cloud-based platform for developing, testing, and deploying applications. PaaS includes everything that is needed to build and run applications, including operating systems, development tools, and middleware.

PaaS is suitable for businesses that want to focus on developing applications without worrying about the underlying infrastructure.

- Software as a service (SaaS)

  Software as a service (SaaS) is the highest-level service model. It involves providing customers with access to software applications that are hosted and managed by the service provider. SaaS is suitable for businesses that want to use software applications without having to install and maintain them on their own systems.

2.2 Cloud Computing Trends in Nigeria

Cloud computing has seen significant adoption in Nigeria in recent years, with more and more businesses and organizations turning to cloud services to improve their operations and competitiveness. Here are a few of the key trends in cloud computing in Nigeria:

Increased adoption: The adoption of cloud computing in Nigeria has been on the rise as more businesses and organizations in the country recognize the benefits it offers. According to a survey conducted by the Nigeria Internet Group (NIG), the adoption of cloud computing in Nigeria increased by 23% between 2019 and 2020.

A growing focus on security: As businesses become more reliant on cloud computing, there is a growing focus on security. Cloud vendors in Nigeria are investing in robust security measures to protect their clients' data and meet compliance requirements.

Increasing use of hybrid cloud: Many businesses in Nigeria are opting for hybrid cloud solutions, which combine on-premises and cloud-based infrastructure. Hybrid cloud enables businesses to take advantage of the scalability and flexibility of the cloud while still retaining control over sensitive data.

Rising demand for managed services: As businesses adopt cloud computing, there is a growing demand for managed services, such as cloud migration, data backup and recovery, and cloud security. Managed services enable businesses to outsource the management of their cloud infrastructure to specialized providers, allowing them to focus on their core competencies.

Growing use of artificial intelligence and machine learning: As businesses in Nigeria adopt cloud computing, there is a growing trend towards the use of artificial intelligence (AI) and machine learning. These technologies enable businesses to process and analyze large amounts of data quickly and accurately, allowing them to make better informed decisions and improve their operations.

Growing Demand for Cloud Skills: As cloud computing becomes more prevalent in Nigeria, there is a growing demand for skilled professionals who can implement, manage, and maintain cloud solutions. This has led to an increase in cloud training programs and certifications, as well as the development of specialized cloud job roles.

Overall, cloud computing is set to continue its growth trajectory in Nigeria, driven by the need for cost-effective, scalable, and secure IT solutions. As more businesses and organizations adopt cloud services, the industry is expected to become even more innovative and competitive in the coming years.

2.3 Cloud Adoption in Nigerian FinTech Industry

Adoption of cloud computing in the Nigerian FinTech industry has been on the rise in recent years, driven by increasing demand for digital financial services and the need for improved efficiency and cost-effectiveness. However, the regulatory landscape for cloud adoption in the Nigerian FinTech industry is still in its early stages, with a number of challenges and uncertainties that need to be addressed.

One major challenge facing the adoption of cloud computing in the Nigerian FinTech industry is the lack of clear and consistent regulations. While there are some existing regulations related to data protection and cybersecurity, these are often vague and open to interpretation, making it difficult for FinTech firms to understand their legal obligations and liabilities when using cloud services.

Another challenge is the lack of infrastructure and connectivity in many parts of the country, which can make it difficult for FinTech firms to access and use cloud services effectively. This is particularly true in rural areas, where Internet connectivity is often limited or non-existent.

Despite these challenges, there are also a number of opportunities for the adoption of cloud computing in the Nigerian FinTech industry. For example, the use of cloud services can help FinTech firms to improve their agility and scalability, allowing them to respond more quickly to changing market conditions and customer needs. Additionally, the use of cloud services can help FinTech firms to reduce their operational costs and improve their operational efficiency.

Overall, the adoption of cloud computing in the Nigerian FinTech industry is still in the early stages, and there are a number of regulatory challenges and uncertainties that need to be addressed. However, with the right regulatory framework and infrastructure in place, the use of cloud services has the potential to bring significant benefits to the industry, including improved efficiency and cost-effectiveness.

3 Significance of Nigerian FinTech's Cloud Adoption

This chapter will describe key FinTech technologies, promising FinTech applications, and public cloud's significance in the Nigerian FinTech industry.

3.1 Challenges Facing Nigeria's FinTech

Cloud computing is critical to FinTech. Nigeria's FinTech enterprises hope to migrate more services to the cloud as soon as possible, but some problems hinder the journey.

- **Challenge 1: Complex and outdated approval processes**

  Regulatory authorities accept cloud adoption but may avoid making difficult decisions and suspend the approval process, leaving decision making to their successors. This is commonly seen in positions with high employee turnover.

- **Challenge 2: Complex data retention and security issues**

  One of cloud computing's highlights is that it has no geographic restrictions, but storing FinTech data on the cloud in another country or region may raise judiciary issues. For example, to comply with data protection laws and regulations in their own countries, a cloud vendor may reject investigation from the customer's local government or other even more powerful institutions.

  Data retention measures are subtle when the regulatory environment for data and public cloud computing has been growing increasingly complex. Additionally, cross-border economic activities may run up against conflicting data retention policies in different countries or regions and some required data may be unobtainable.

- **Challenge 3: Fierce competition**

  When product managers and operational teams compete for FinTech products to be rapidly evolved and quickly rolled out, the underlying infrastructure and applications grow more and more complex. This puts a lot of pressure on traditional IT systems.

  In a traditional system, labor is expensive and engineers struggle to keep up with technological changes. Some financial organizations still use a traditional IT architecture which is incapable of supporting cloud native FinTech services.

- **Challenge 4: Security risks due to fast development**

FinTech focuses on cooperation, innovation, convergence, and quick responses. The quick launch of new services benefits customers but may introduce security risks. Global cooperation increases the risk of external attacks. A variety of new services, such as different kinds of apps, blur network borders and make border security more difficult. The convergence of online and offline services tends to make risks propagate. In the financial industry, many business lines, complex functions, and requirements for quick response reduce the security of new products and services. When the underground financial service chain is complete, criminals have a deeper understanding of new technologies and services of banks, which introduces more risks to financial systems.

- **Challenge 5: External attacks due to big data channels**

  Big data analytics is widely used to maximize data value. Massive volumes of correlative financial data may disclose customer privacy and will be exploited by hackers for illegal activities. They steal and analyze big data for precision attacks on customers. They steal funds from bank customers by illegally obtaining databases, sorting data, and launching credential stuffing attacks.

- **Challenge 6: Difficult data isolation**

  Cloud computing provides shared resource pools that can be accessed on demand at any time. Resource reuse, blurred physical boundaries, uncertain locations of virtual storage, and hypervisor virtualization make it difficult to isolate data at each layer, which may result in data leaks.

3.2 Public Cloud Facilitates Cloud Transformation of FinTech

When thinking about cloud adoption in the FinTech industry, there are four aspects to consider: regulatory compliance, cloud security, cloud native, and operational costs.

- Regulatory compliance: Cloud adoption must comply with regulatory requirements in the financial industry. This white paper describes how Nigerian FinTech on Huawei Cloud is compliant with the international and local laws and regulations.
- Cloud security: Cloud vendors must ensure cloud security. It is not simply the security of single points. Cloud vendors need to consider all aspects of security for financial services to move to the cloud. Huawei Cloud provides a security framework and success cases to help enterprises formulate comprehensive cloud migration policies.
- Cloud native: The digital and cloud transformation of FinTech is based on a cloud native infrastructure and a distributed FinTech architecture.
- Operational costs: Management framework is provided. To provide support for decision makers of FinTech enterprises, every aspect of cloud operational costs are analyzed based on a cloud architecture.

3.2.1 Regulatory Compliance

The Nigerian's FinTech industry is facing more regulatory challenges than other countries. The changing FinTech environment and complex emerging technologies (such as blockchain, AI, and digital banking) exacerbate conflicts between market competition and regulatory compliance. The good news is that Nigeria's regulatory authorities have adopted loose market policies and open communication channels to support the development of FinTech.

Currently, the main regulatory authorities and their regulations in Nigeria are as follows.

Table 3-1 FinTech regulatory authorities of Nigeria

Regulatory Authority	Function	Role in FinTech	FinTech Segment
CBN	The Central Bank of Nigeria (CBN) is the principal regulatory authority mandated to issue licenses to banks and other financial institutions by virtue of the Banks and Other Financial Institutions Act (BOFIA), 2020. The CBN has the primary responsibility for the regulation of the payment system and forms of banking business and other financial activities.	Administering operating licenses and prudential supervision of FinTechs	Banks, Payment Companies and Other Financial Institutions.
SEC(Securities and Exchange Commission)	Develop and regulate a capital market that is dynamic, fair, transparent and efficient, to contribute to Nigeria's economic development.	Deploying innovative products or processes in the Nigerian capital market. Issuing operating licences and rules of conducts for licensed players.	FinTechs playing in the capital market
NDIC(Nigerian Deposit Insurance Company)	Protect depositors and guarantee the settlement of insured funds when a deposit-taking financial institution can no longer repay its deposits, thereby helping to maintain financial system stability.	Identifying, developing, and promoting technology solutions that will protect depositors with insured financial institutions	Deposit-receiving FinTechs
NOTAP(National Office for Technology Acqusition and Promotion)	Evaluation/Registration of technology transfer agreements; promotion of intellectual property; oversight over R&D commercialization, activities, and industry	Issuing approvals for the transfer of technology	FinTechs technology arrangements with foreign companies.

Regulatory Authority	Function	Role in FinTech	FinTech Segment
	linkage		
NCC(Nigerian Communications Commission)	Responsible for economic and technical regulation of the communications industry	Regulating FinTech services leveraging mobile phones pursuant to the license framework for Value Added Service (VAS) issued by the NCC	Mobile money operators and other Fintechs who utilize Value Added Services in their offerings.
NITDA(National Information Technology Development Agency)	Operate and implement the Information Technology policies, frameworks and standards for the development of IT in Nigeria and to give effect to provisions of the National Information Technology Development Agency Act (NITDA Act) of 2007.	Regulating the use of policies such as Cloud Computing Policy, data privacy and protection across Nigeria	All FinTechs
NAICOM(National Insurance Commission)	Oversee the ministration, supervision, regulation, and control of insurance business in Nigeria and protect insurance policyholders, beneficiaries, and third parties to insurance contracts.	Overseeing the activities of InsurTechs that overlap with the mandate of NAICOM	InsurTech
NFIU(Nigerian Financial Intelligence Unit)	Responsible for the receipt of disclosures from reporting organizations, the analysis of these disclosures, and the production of intelligence for dissemination to competent authorities	Monitoring money laundering and terrorist financing risks within the ecosystem	All FinTechs
NIMC(National Identity Management Commission)	Regulate matters of national identity in Nigeria with services covering National Identification Number	Identifying verification and authentication, KYC processes	Mobile money, digital banking, lending, wealth management, and

Regulatory Authority	Function	Role in FinTech	FinTech Segment
	(NIN) enrollment and issuance, National e-ID card issuance, identity verification as well as data harmonization and authentication		InvesTech
NSE(Nigerian Stock Exchange)	Regulated by the Securities and Exchange Commission (SEC) of Nigeria. The Exchange offers listing and trading services, licensing services, market data solutions, ancillary technology services and more.	Providing access to capital	FinTechs seeking listing on the capital market
NIBSS(Nigerian Interbank Settlement System)	Deliver payment/settlement services to financial institutions, and mitigate operational and credit risks in funds transfer across financial institutions	Providing infrastructure for the automated processing and settlement of transactions between financial institutions	FinTechs leveraging inter-bank transactions
SANEF(Shared Agent Network Expansion Facilities)	Deepen financial inclusion in Nigeria through a robust ecosystem with strong regulatory oversight, consumer protection, and interoperable systems.	Providing support to super agents, banks, and other stakeholders to acquire agents in the 6 geo-political zones.	Mobile money, digital banking, and lending

In conclusion, the relationship between regulatory authorities and stakeholders in the industry is critical to the success of financial services in Nigeria. The main challenges facing the financial industry are as follows.

- Slow update of rules and regulations

  The rules and regulations evolve too slowly, which limits growth potential and makes strategic planning difficult.
- Unfair competition

  Some cloud vendors or regulatory authorities, are allowed to compete in the ecosystem, which is unfair for the rest FinTech players.

Compliance with Laws and Regulations

Cloud vendors must comply with laws and regulations issued by international and national legislatures listed in Table 3-1.

Nigeria's Certifications

Nigeria's regulatory authorities such as the information and information security departments evaluate and certify cloud vendors' capabilities, service quality, credibility, and network and information security, aiming to promote the vendors to strengthen their management and improve service quality and their credibility.

International Certifications

Cloud services deployed in different countries or regions must comply with international and local laws and regulations.

An important certification for management systems is from the International Organization for Standardization (ISO). It reflects that a cloud vendor has a mature management mechanism and complies with industry standards.

The compliance assessment reports for management systems demonstrate to customers the integrity and effectiveness of cloud platform management, such as the management is continuous and effective and the responsibility division between the cloud vendors and customers is clear.

Huawei Cloud Security Compliance

Huawei Cloud is the only vendor certified by the CBN and that has issued an NDRP compliance white paper (see *Huawei Cloud Compliance with Nigeria NDPR* and *Huawei Cloud User Guide to Financial Services Regulations & Guidelines in Nigeria)*.

Huawei Cloud has earned or completed more than 100 international compliance certifications and third-party audits. In terms of privacy protection, Huawei Cloud is committed to protecting the personal data of each customer. Customers own and control all personal data provided to Huawei Cloud. Huawei Cloud eagerly cooperates with international and local regulatory authorities in Nigeria to take responsibility for protecting personal data and continuously improve the personal data protection system. Huawei Cloud has been certified by a large number of authoritative organizations for its ability to protect personal data. For example, Huawei Cloud has been certified for ISO 27017 and Payment Card Industry Data Security Standard (PCI DSS). Huawei Cloud will continue to improve the personal data protection system. In this process, Huawei Cloud will focus on how to better safeguard data security when managing and processing personal data.

3.2.2 Cloud Security

Security is the topic that most concerns FinTech enterprises during their cloud adoption.

Advantages of Cloud

Cloud is secure because of its:

- Advanced technologies

  Traditionally, security technologies are fixed to a level applied when a service system was established and cannot adapt to full-stack IT risk control. When enterprises try to adopt new technologies, old infrastructure impairs inspection and upgrade and drives up costs. As a result, the security technologies of most enterprises are out of date. In contrast, security services are continuously rolled

out and updated on the cloud. Network security teams of cloud vendors can safeguard enterprise data using the latest security technologies.

- Large scale

  Cloud vendors can leverage the large reserves of IT resources they possess to easily solve security problems. For example, cloud vendors can use their large data centers and servers to provide low-cost Disaster Recovery (DR) solutions. If enterprises build their own DR data centers, the costs are huge. Cloud vendors have the large bandwidth needed to help enterprises defend against DDoS attacks. It is impractical for enterprises to purchase bandwidth at that scale.

- Excellent experts and experience models

  Cloud vendors typically have a deep pool of IT talent available to ensure system stability on the cloud and develop technical means to defend against security attacks in a timely manner.

  Cloud vendors are faced with more severe security threats than FinTech enterprises. In the long-term process of fighting against viruses and network attacks, cloud vendors have accumulated extensive experience which enables them to fix vulnerabilities and address risks quickly. Enterprises cannot handle security problems as effectively as cloud vendors due to a lack of knowledge and experience.

Security Architecture

Ensuring the security of the cloud requires a security assessment, a secure design, and secure operational practices.

- Cloud security assessment
  - Best practices need to be defined for identity authentication and access control, including single sign-on, role- or policy-based permissions control, and active directory (AD) integration.
  - Appropriate data protection and encryption mechanisms need to be identified for applications on the cloud. These mechanisms should cover data lifecycle management, encryption key management, and system protection processes.
  - The network architecture needs to be verified to ensure that network segmentation and protection mechanisms are in place, including microsegmentation, VPC and subnets, security groups, VPN connections, and the principle of least privilege (PoLP).
  - Tools or methods used for improving PaaS security need to be evaluated. There needs to be application threat modeling, web scanning tools, and data masking methods.
  - Potential log sources and log analysis and storage methods need to be evaluated. This includes the log collection and storage system, tool choices, and log lifecycle management.
  - Security information and event management (SIEM) requirements need to be reviewed and methods for implementing SIEM on the cloud need to be defined. This includes the SIEM architecture, vulnerability assessment methods, and monitoring policies.
  - Security patches and infrastructure updates need to be installed. This means patching the operating systems and using various system hardening tools, persistent storage policies, and database patching methods.

  - There needs to be a model in pace for continuously improving the security architecture. There should be periodic environmental audits, security tests, and evaluation of new methods and security products.

- A secure design

  A secure design includes:

  - Multi-layer isolation (in-depth defense)

    Security isolation and access control are used to isolate networks, systems, applications, and data.

  - Authentication and authorization (PoLP)

    Accounts and subaccounts are granted with only the necessary permissions. Two-factor authentication and secondary authentication for key operations are enabled.

  - Encryption

    Data is encrypted for transmission and storage.

  - Monitoring and alarming

    Logs and monitoring measures are used to detect configuration changes, abnormal logins and operations, data leaks, and attacks in a timely manner.

- Huawei Cloud security practices

  Huawei Cloud is secured by:

  - Cloud security services, such as Web Application Firewall (WAF), Data Security Center (DSC), and Key Management Service (KMS)
  - Cloud service features, such as VPC network isolation, and data transmission and storage encryption
  - Security management and monitoring activities of users on the cloud, such as log analysis, operation auditing, and server and container monitoring

  Security measures related to the following factors are recommended on the cloud:

  - Basic security
  - Service security
  - Application security
  - System security
  - Isolation measures
  - Data lifecycle management
  - Security monitoring
  - Operational security

For more information, see *Huawei Cloud Security White Paper*, *White Paper for Huawei Cloud Data Security*, *White Paper for Huawei Cloud Privacy Protection*, and *White Paper for Huawei Cloud Trustworthiness*.

Cloud Platform Security

Cloud platform security includes the physical environment, virtualization, infrastructure, and cloud service security, identity and access control, and security monitoring and operations.

- Physical security

Data center disaster recovery, access control, account and identity authentication, O&M audits, data destruction policies and process, data erasure policies, customer data disposal policies, and network isolation policies

Huawei Cloud's data centers provide T3+ level security.

- Infrastructure security
  - Physical network security, hardware firmware security, hardware firmware baseline scanning, GPU-accelerated instance protection, BIOS signature verification, and BMC firmware protection
  - Encrypted computing (encrypted data can only be processed in a trusted environment, providing hardware-based data protection)
  - Trusted computing (the cloud platform is measured and verified to ensure the computing environment security, and whitelisted applications are monitored to ensure the application security)
- Virtualization security

  The security of different virtualized resources needs to be ensured. There needs to be a focus on the virtualization providers, managers, and users. The following security factors require attention:
  - Virtual networks: multi-tenant security isolation, network defense, security domain, and multi-instance GPU
  - Virtual machines (VM): security patches, security vulnerabilities, virus defense, secure policy migration, and access control
  - Virtual storage: storage isolation, storage backup, and access control
  - Virtualization management: virtual image security, authorization and authentication, data destruction security, audit logs, and monitoring
- Cloud service security
  - Security of computing, storage, and network services
  - Service registration, service application and deregistration, permission authentication, service monitoring, service flow control, and service routing
  - Network security isolation, communication protocols, port authentication, and data encryption
  - Database services, big data services, component/service security, permission authentication, audit logs, code specifications, code vulnerabilities, data encryption, and data verification
  - Component/service management, service metering and billing security, component/service transmission security, service access security, and component/service integration security
- Identity authentication and access control
  - Account lifecycle management
  - Unified login management, account and password management policies, and access control policies
  - User permissions management

    Permissions are allocated based on their rights, roles, user groups, and departments. Each user applies for, uses, and releases their permissions in the permission management system.
- Security monitoring and operations
  - Cloud service security management

The security of cloud services is ensured in each stage (design, development, test, and release) before they are launched.

- Cloud platform monitoring

  Attacks on applications, servers, and networks need to be detected in a timely manner. After a security event is detected, the cloud platform immediately handles the event and addresses any impacts. The vulnerabilities and security events reported from external systems also need to be properly handled in a timely manner.

Cloud platform security measures may include:

- Identity and access management (IAM): This includes authentication and authorization measures to ensure only authorized individuals have access to resources within the cloud environment.
- Encryption: This includes encrypting data in transit and at rest, to prevent unauthorized access or interception.
- Network security: This includes using firewalls and network segmentation to protect the cloud environment from unauthorized access.
- Security monitoring: This includes continuous monitoring of the cloud environment for security threats, using tools like intrusion detection systems (IDS) and security information and event management (SIEM) systems.
- Incident response: This includes having a plan in place to respond to security incidents, such as data breaches or system failures.
- Compliance and auditing: This include ensuring the cloud platform meets relevant regulatory requirements and standards, and conducting regular audits to ensure security controls are effective.

Huawei Cloud Security Architecture

Focusing on customer requirements, Huawei Cloud Service Cybersecurity & Compliance Standard covers the security architecture design on both the cloud and customer sides. Vertically, it includes account security, security monitoring, and operations management, and horizontally it includes cloud platform security, basic user security, user data security, application security, and service security. Huawei Cloud can use this standard to design a cloud security management system and customers can use it to improve their system security.

- Account security

  Huawei Cloud's Identity and Access Management (IAM) centrally manages users and their permissions. Each cloud resource belongs to a unique Huawei Cloud account that will pay for the resource and have absolute control over the resource. An enterprise can create multiple IAM users under a Huawei Cloud account for their employees, departments, or applications and manage permissions of the IAM users. Each IAM user has an independent login password or access key to perform operations on the Huawei Cloud console or using APIs. Huawei Cloud accounts support multi-factor authentication, strong passwords, separation of console users and API users, fine-grained permissions control, user groups, and temporary authorization tokens. Huawei Cloud's resource catalog service allows administrators to easily create a resource catalog structure that reflects service relationships. Then, they can distribute

accounts to different positions of the structure for hierarchical resource management.

- Basic user security

  Server data is collected to monitor server status. Big data and threat intelligence are used to detect threats. Logs are used to quickly locate causes. Take containers as an example. Each container has an exclusive kernel to isolate its memory, network, and I/Os from other containers. In this way, users can be isolated from each other when their containers are running on the same host. Containers are protected by Web-CMS vulnerability detection and repair, webshell detection and elimination, cloud antivirus, process behavior monitoring, network connection monitoring, process startup logging, and network connection logging. Container images are scanned based on the latest CVE list. Only container images signed by the container users can be used in Huawei Cloud container clusters. Huawei Cloud uses tunneling technologies to provide isolation at the data link layer so that user networks can be kept isolated from each other. Instance-level virtual firewalls provide status detection and data packet filtering. The firewalls can also divide security domains between computing instances (for example, container clusters) on the cloud. Both north-south and east-west network traffic are analyzed. Network-wide traffic (such as Internet access and inter-security group access) is visualized. External connections originated from inside the network can be identified and blocked. A full-traffic proxy is used to defend against heavy-traffic attacks and refined resource starvation attacks at the web application layer.

- User data security

  Massive pools of data are scanned for sensitive content based on AI and machine learning using built-in algorithms or custom models. The use of sensitive data is recorded and analyzed to identify security issues in a timely manner and generate alarms, helping prevent data leaks and meet compliance requirements such as the NDPR. Data can be securely shared in an anonymized format. A variety of anonymization algorithms and adaptive parameters ensure that anonymization results are matched to different development, testing, and analysis environments. To reduce data risks, data exceptions can be identified and alarms will be generated. Big data analysis identifies abnormal access to sensitive data and files and suggestions will be provided accordingly. Abnormal behaviors are audited and traced, and alarms will be reported. Operations on both structured and unstructured data are audited to improve enterprise responses to security risks and help meet compliance requirements.

- Application security

  Batch vulnerability scanning and other attacks are automatically blocked. Threat intelligence tools are used collaboratively to defend against attacks. Strict Web Application Firewall (WAF) rules and policies are configured to defend against 0-day vulnerability attacks. A whitelist is used to control access to applications. Whitelist changes are identified by comparison with a baseline. Sensitive information, for example database connection strings (including passwords), is stored to configuration files and distributed to the production environment. During Docker orchestration, sensitive information is stored in environment variables.

- Service security

  SDK and API invoking are used to identify scalpers and customers cheating in marketing promotions like flash sales. Anti-bot solutions reduce and eliminate the impact of malicious tools on user websites. Cutting-edge technologies are

used to precisely detect sensitive information, reducing any non-compliance risks. The content moderation costs are less.

- Security monitoring and operations

  An in-depth, three-dimensional threat detection architecture is built to detect threats based on clues hackers leave behind. Compliance results can be viewed and downloaded online. Notifications of resource changes and compliance events help learn about resource status in real time. Methods are built into rules to automatically correct non-compliance in a timely manner. All auditing logs of a FinTech enterprise are centrally stored, meeting NDPR requirements. Third-party audit systems, Security Operation Center (SOC), and backup systems are interconnected for automatic monitoring and alarming. Penetration testing is performed during each phase of security development. This helps enterprises discover system defects and vulnerabilities as early as possible. Rectification suggestions will be provided for them. After being authorized by customers, Huawei Cloud's security experts will take over the security of customer assets on the cloud. Customers only need to make sure their services are secure. They do not have to develop their own security professionals, security technologies and tools, and security systems.

3.2.3 Cloud Native

FinTech's cloud adoption places new requirements on IT infrastructure.

- Continuous delivery

  Developers work on software branches to build new features. When the new features are ready, they are merged back to the main branch. The main branch version can always be reliably released at any time. This improves development efficiency to quickly cater for customer needs. A traditional monolithic architecture is not capable of continuous delivery because the code maintenance, integration, and compilation are too difficult.

- Superlative scalability

  It can take several minutes to provision VMs, which is not fast enough for some scaling needs. Containers are more lightweight than VMs, which means they can scale faster. In addition, containers ensure that if a microservice runs properly in the development environment, it will run properly anywhere.

- Integrated development and operations

  Agile development enables faster iteration but also makes releasing versions riskier and complicates O&M. To address this, a platform that integrates and automates development, testing, and O&M is needed to complete high-frequency deployment, improve reliability, stability, scalability, and security of the production environment, and eliminate release risks.

FinTech IT Agility Accelerated by Cloud Native

- Definition and principles of a cloud native architecture

  A cloud native architecture separates non-service code from cloud applications as much as possible so that cloud infrastructure takes over non-functional features (such as scalability, resilience, security, and observability). In this way, services will be more lightweight, agile, and automated.

  A cloud native architecture complies with the following principles:

  - Service-oriented: Relationships between service modules are abstracted and standardized to decouple modules with different lifecycles from each other.

The overall development is faster and more stable when the models iterate in parallel.

- Scalable: The system deployment can be automatically scaled to keep up with service changes. This will reduce resource waste.
- Observable: Cloud computing is built on a distributed architecture, so to ensure service health, log analysis, link tracing, metrics measurement, and various other methods are introduced to record, analyze, and drill down cross-server service invocations from end to end.
- Resilient: A cloud native architecture is able to resist hardware or software damage so that services will not be interrupted.
- Automated: Descriptive configurations and a result-oriented delivery process tune up automation tools. The tools must adapt delivery objectives to actual environments so that delivery and O&M can be automated.
- Zero-trust: A cloud native architecture focuses on identity-centric access control. Any personnel, devices, and systems must be authenticated and authorized every time they try to access the network.

- Differences between the cloud native and traditional architectures

  A cloud native architecture has changed the way enterprises build and use computing resources. Hardware utilization has been improved and resources usage is more flexible. With cloud native, applications depend less on running environments, so they can be deployed more quickly and be migrated easier. According to a report from Gartner, 70% of global enterprises will deploy their applications on a cloud native architecture by 2021. The trend is turning away from traditional IT architectures and transitioning to cloud native architectures.

  Traditional architectures are suitable for stable service deployment. Cloud servers have been the most important infrastructure since the birth of cloud computing. Over the past decade, the performance of cloud servers has undergone significant changes. Low latency, high performance, and isolated environments make cloud servers the best option for deploying monolithic and performance-sensitive applications (such as high-performance computing, video transcoding, and machine learning). For large applications, especially stateful applications and applications requiring high data durability, cloud servers are also a good choice. Because of the excellent security and isolation performance, cloud servers are widely admired by IT operations, risk management, and development personnel.

  Cloud native architectures are more suitable for sensitive service deployment. Leveraging application-oriented scaling, powerful portability, and lightweight deployment, cloud native architectures are becoming increasingly popular. Cloud native applications are based on microservices. A monolithic application is split into multiple loosely coupled microservices based on functions or logical associations. These microservices can be developed, tested, and deployed in parallel, greatly improving the iteration efficiency. A container is a standardized software unit that packages applications and all their dependencies so that applications can run quickly and reliably across computing environments. A shared kernel makes containers more lightweight. Containers can be scaled in seconds to respond to service requirements. They can be scaled more quickly than VMs.

Benefits of Cloud Native to FinTech

Cloud adoption is more than just migration to the cloud. To adapt to the cloud, application design, development, deployment, and maintenance all need to be improved, so they can take full advantages of the cloud platform.

- IT architecture

  Excellence of the cloud: Cloud native greatly unleashes the technology dividends of cloud computing and fully inherits the design ideas of cloud computing. Cloud native applications can be well supported in the cloud environment in terms of resource orchestration, distributed deployment, and high-availability architecture, which will be helpful for building a robust application system.

  More elastic deployment: Cloud native removes deployment inconsistency across environments. This promotes resource standardization and provides a basis for service-oriented and automated processes. Cloud native uses containers as basic scheduling units. Compared with process-level resource division in VMs, a shared kernel that containers use makes resources more elastic.

  Compatible with diverse technology stacks: Unlike a traditional architecture where applications are bound to specific languages and technology stacks, a cloud native architecture separates applications in different service domains from each other. This allows applications to have different technical options. For example, using Python in a recommendation system is much more efficient than using Java. Cloud native architectures make it possible to use multiple technology stacks for application development. A service team can choose whichever technical stacks are best suited to their specific needs. There is no need to worry about compatibility issues.

  Improved service stability: Excellent automation and self-healing make cloud native applications more resilient so that they can quickly recover from exceptions. Service statuses, system health, interface invocations, and real-time alarms are displayed on the console. Based on monitoring results and health check results, container-level automatic recovery is triggered to address service risks in a timely manner.

- Enterprise operations

  Reduced enterprise IT costs: The incredible auto scaling of cloud native avoids wasting resources reserved for peak hours, improves resource utilization, and reduces resource costs. In a traditional architecture, non-service modules are embedded into applications, resulting in high R&D costs. In a cloud native architecture, applications are separated from infrastructure (such as the programming framework and middleware). Infrastructure becomes loosely coupled modules that are sharable in resource pools. Application infrastructure is integrated with the cloud platform to provide a technical platform that provides simple, consistent, and convenient interfaces for applications, reducing labor and resource costs in R&D.

  Faster service delivery: Digital transformation places more pressure on service channels, intensifies competition, and raises the bar for user experience. Rapid response to customer requirements has become a core competency for enterprises in today's digital world. In a cloud native architecture, applications are developed as microservices that use standard APIs for interaction. The loosely coupled architecture reduces iteration costs, provides a basis for co-development of multiple teams, and accelerates service delivery. Cloud native enables agile application development, speeds up service delivery, reduces trial-

and-error costs, enables faster response to customer requirements, improves user experience, and accelerates service innovation.

Relieved pressure: In a traditional architecture, middleware is usually bound to specific services and cannot be reused for other purposes. During application deployment, a large number of labors are needed to build middleware, increasing the risk of errors. In a cloud native architecture, cloud vendors take care of O&M. Enterprises can easily deploy and start cloud native clusters. Applications can be quickly rolled out with the help of middleware services provided by the cloud platform. Enterprises are free to focus on their own services. There is no need to worry about other issues.

Streamlined development and operations: DevOps improves the efficiency of enterprises throughout a development and operations lifecycle. A standard tool chain streamlines the cognition, process, and tools, breaks department silos, and standardizes and automates development processes, improving production efficiency.

Cloud Native Concepts and Technologies

Cloud native is a collection of cloud computing technologies and enterprise management methods. Cloud native technologies include containers, service meshes, microservices, immutable infrastructure, and declarative APIs. Cloud native features a loosely coupled system that is resilient and can be easily managed and observed. By automating the process, frequent, reliable, and predictable updates are allowed for services, and applications can be released at any time. Cloud native helps enterprises build and run scalable applications in public, private, and hybrid clouds. By automating cloud platforms, microservices can be built across clouds and continuously delivered to production environments.

- Containers

  Containers are a lightweight virtualization technology that provides multiple isolated operating system environments on a single host. Processes are isolated using namespaces. Each container has a unique writable file system and resource quota. Container runtimes are responsible for running individual containers, such as providing computing, storage, and network resources. Container orchestration is responsible for container cluster scheduling, service discovery, and resource management.

  With containers, a microservice and all required configurations, dependencies, and environment variables can be packaged into an image which can be used to easily deploy the microservice on new servers. This makes containers ideal for deploying a single microservice.

- Microservices

  In a microservice architecture, applications are structured as a collection of very small, loosely coupled services. Each "microservice" can be independently upgraded, deployed, expanded, and restarted. In this way, frequent updates will not affect end users. Unlike a traditional monolithic architecture, in a microservice architecture the system is simple. Individual microservices can be independently deployed and expanded. Cross-language programming is not an issue.

  To address the O&M challenges of a flexible architecture and agile development, a microservice framework is used for microservice registration, discovery, and governance. Currently, there are two microservice architectures: intrusive and non-intrusive. In an intrusive architecture (for example, Spring Cloud), a microservice framework is embedded into code for class inheritance. In a non-

intrusive architecture (for example, Istio), the framework is deployed together with applications as a transparent proxy covering the application network.

- DevOps

  DevOps is a set of processes, methods, and systems that emphasizes collaboration and communication between software development (applications or software engineering), technical operations, and quality assurance departments. DevOps emerged as a response to the software industry's awareness that close collaboration between development and operations is essential for rapid delivery of software products and services.

  DevOps automates and monitors the entire process from software building and integration, to testing, release, deployment, and O&M. It aims to accelerate the development cycle, increase the deployment frequency, and produce more reliable releases. A complete tool chain is used to deeply integrate code repositories, artifact repositories, project management, and automatic testing. This will help enterprises quickly achieve the collaboration of software development and operations and move to the cloud without additional costs.

  DevOps provides a new platform for developers and O&M personnel to optimize resources and scale applications through continuous delivery. DevOps works with cloud native to provide a product development process that adapts to rapidly changing market conditions and better serves the business objectives.

- Service mesh

  A service mesh is a dedicated infrastructure layer that provides capabilities like management, observability, and security to service-to-service communication over a network. A service mesh works like a set of lightweight network proxies deployed with application programs. The applications are unaware of the service mesh they run on. A service mesh consists of a control plane and a data plane. The data plane runs in proxy Sidecars where transparent infrastructure is provided for network communications with low or zero interference.

  Service meshes, often thought of as a second-generation microservice architecture, have greatly changed microservices.

  - Service governance decoupled from service logic

    Service meshes separate most SDK capabilities from applications, split them into independent processes, and deploy them in Sidecars. In this way, service communications and related management and control functions are separated from the service system and offloaded to the infrastructure. Developers can focus on services without worrying about other issues.

  - Unified governance of heterogeneous systems

    Service meshes offload service governance to infrastructure to enable unified traffic control across languages and protocols.

- Serverless architecture

  Serverless computing (or serverless for short) is a model for building and managing microservices-based architectures. It allows you to manage your application deployment on a per-service level, rather than a per-server level. A serverless architecture runs in stateless containers triggered by events. By using a serverless architecture, developers can offload the responsibility for managing or operating local or cloud servers to a third-party provider, enabling them to focus better on writing application code. A serverless architecture allows developers to build, deploy, and start services without having to manage the underlying infrastructure.

  The core idea of a serverless architecture is to abstract infrastructure into services and provide them as scalable APIs. Users can call them on demand

and pay for what they use. This architecture eliminates the demand for massive online server components, simplifies development and O&M, reduces operational costs, and shortens the delivery period.

Currently, there are three mature serverless-based technologies: function as a service (FaaS), backend as a service (BaaS), and serverless containers.

Huawei Cloud Native Architecture

Huawei Cloud provides solutions and best practices for a large number of enterprises in different industries to help them complete digital transformation, especially the adoption of cloud native technologies. Huawei Cloud has focused its attention on integrating enterprises' core concerns, organizational and IT culture, along with Huawei's engineering capabilities to create a unique cloud native architecture. When the architecture is designed, the following issues need to be considered:

- FinTech strategies

  In the financial industry, any architecture has to serve the strategies of financial enterprises, and a cloud native architecture is no exception. A cloud native architecture means not only a technical upgrade, but also the reconstruction of the production process of enterprise core services. In this process, digital services will be built through software development and operations. It is like replacing manual workshops with automated assembly lines in the industrial era.

  Generally, FinTech companies have more demanding requirements for computing technologies. For example, they use big data services provided by cloud vendors to provide a more intelligent user experience, and they use technologies like cloud native to establish wider and intuitive connections. More and more enterprises believe that IT strategies are needed to enable business innovation. The existence of C-suite roles such as Chief Technology Officer (CTO), Chief Information Officer (CIO), Chief Data Officer (CDO), and Chief Information Security Officer (CISO) all speak to the important role technologies play in enterprise strategies.

- Business development

  When providing cloud and consulting services for enterprises, Huawei Cloud finds that the main requirements of digital services for a cloud architecture are service continuity, quick service rollout, low costs, and technology-enabled service innovation. Service continuity refers to how digital services must be able to continuously provide services for users. It means there cannot be service interruptions due to software or hardware faults or bugs, and there have to be defenses in place against hackers. Data centers must be always available, and measures need to be in place in the event of a natural disaster. When services expand rapidly, there needs to be enough software and hardware resources so that new users can be continuously acquired.

  Digital services are more flexible than traditional services. To keep up with fast changing market conditions, digital services need to be innovated and updated quickly. A cloud native architecture can provide these capabilities at the product, tool, and process levels. To make these capabilities possible, enterprises' organizational structures and applications may need to be reconstructed (for example, applications need to be split into microservices).

  Cloud computing helps enterprises shift CapEx investments into OpEx. With cloud computing, enterprises do not need to purchase software and hardware resources ahead of time. They can pay as they go, buying cloud resources on-demand. A cloud native architecture reduces enterprise development and O&M costs. Especially, containers reduce O&M costs by more than 30%.

Traditionally, when new technologies are introduced, there is a lengthy selection process, proof-of-concept testing, pilot projects, and promotions. In contrast, cloud services involve less trial-and-error costs, and they can combine different technologies on the same platform. So, cloud services can adopt new technologies faster, which enables faster innovation.

- Organizational structure

  The adoption of a cloud native architecture has a great impact on the development, testing, and O&M personnel of enterprises. To ensure the continuous evolution of the architecture, a department like an Architecture Governance Committee, needs to be established to continuously evaluate whether the architecture is performing as planned.

  According to Conway's Law, organizations design systems that mirror their own communication structure. This means that enterprises need to adjust their organizational structures to adapt to a service-oriented architecture (microservices, small services).

- Cloud native architecture

  Service-oriented

  Application modules with different iteration periods are split into microservices or small services and invoked using standard APIs. Service integration is driven by events, reducing service dependencies. The SLA is continuously improved through a strengthened measurement system.

  Scalable

  Scalable resource pools ensure that systems can automatically scale in or out based on workloads.

  Serverless

  More cloud services should be used instead of third-party services, especially when an enterprise performs O&M on open source software. Applications should be built stateless as much as possible and run in serverless cloud services such as FaaS and containers. The stateful components, on the other hand, should run in other cloud services.

  Observable

  IT infrastructure needs to be quickly rectified if an error occurs. The system needs to be observable. There needs to be good logging, monitoring, application performance management (APM), link tracing, and QoS measurement.

  Resilient

  In addition to common patterns such as circuit breakers, rate limiters, fallback, retry, and cache, the resilience design also includes high availability, disaster recovery, and asynchronization.

  Automated

  The development, testing, and O&M need to be agile. Containers are recommended to automate software building. Open Application Model (OAM), Infrastructure as Code (IaC), and GitOps are used to automate the CI/CD pipeline and O&M process.

  Secure

  The key point is digital security. Cloud security services safeguard application running environments. Application security is hardened throughout the entire software development process to ensure that the applications comply with security standards such as ISO27001, PCI/DSS, and DJCP MLPS.

3.2.4 Operational Costs

Different from a traditional data center, the total cost of ownership (TCO) of a cloud data center includes cloud operational costs (cloud infrastructure costs and on-cloud costs) and one-time installation costs (cloud setup and migration costs).

Cloud Infrastructure Costs

- Servers or containers
- Storage
- Network
- Databases
- Middleware
- Licenses
- Operating systems

On-Cloud Costs

- Data transmission
- Extra connections
- Monitoring
- Application development
- Enterprise support

Cloud Setup Costs

- Networking connections
- Security tools
- Monitoring and management tools
- Cloud migration tools
- Cloud deployment

Migration Costs

- Establishment of cloud management organizations
- Application migration
- Operational model changes
- Trainings in skill improvement
- Centralized budgeting and internal accounting

Cloud Adoption Benefits in Operational Costs

Traditional data centers have the following shortcomings:

- Enterprises need to purchase more resources for peak hours than is needed at other times, which drives up CapEx and increases TCO.
- It takes a long time and a dedicated team to purchase and create infrastructure, resulting in extra time and labor costs.

- The time-consuming purchasing process slows down service innovation and makes innovation more expensive.

In contrast, cloud data centers have the following highlights:

- No fixed costs. Enterprises only need to pay for what they use.
- Quick resource provisioning. Infrastructure resources can be automatically provisioned in minutes or even seconds.
- Faster Innovation. Enterprise services can iterate faster when backed by the cloud.

The cloud lets FinTech enterprises convert CapEx to OPEX, reduces time and labor costs, and reduces the cost of innovation.

Operational Cost Optimization

For FinTech enterprises, operational costs on the cloud can be continuously optimized in the following ways:

- Instance selection

  An enterprise can select more cost-effective instances to meet performance and workload requirements.
- Optimized billing modes

  An enterprise can take advantage of different billing modes ato save money. For example, reserved instances and pay-per-use instances can be used by the same enterprise. Reserved instances are a great way to reduce costs.
- Auto Scaling

  Instances can be scaled up or down automatically based on service demands. For example, instances can be automatically scaled in during off-peak hours and scaled out during peak hours.
- Smart consumption

  Idle instances can be automatically identified and stopped. For example, development and testing instances can be stopped on weekends.
- Monitoring

  Metrics are configured to continuously monitor the resource utilization and find opportunities to reduce costs. Because it is easy to provision instances on the cloud, it is more important than ever to monitor the consumption of services and avoid "cloud sprawl" (the uncontrolled scaling out of cloud instances or cloud services).

To further reduce costs on the cloud, FinTech enterprises need to periodically check the costs and benefits of applications. When necessary, applications need to be reconstructed to make full use of new technologies, which is helpful for improving efficiency and reducing costs.

Construction of Cloud Operational System

- Professional services

  Professional services accelerate service development, ensure delivery quality, and reduce deployment risks. They provide highly integrated and automated tools for development teams and support for agile development, management, and deployment of components and applications.
  - Management and control

Professional services provision and supply infrastructure and provide technical support services. Architectures and development practices adaptable to changes are used. Customer requirements are collected and managed to ensure that the infrastructure can keep up with changing service requirements. Applications are periodically tested to provide a reference for better planning.

- Collaboration

  Project teams are provided with the tools and technical support required for an automated and visualized development process.

- Standard compliance

  The delivery process, including development, configuration, building, testing, deployment, and release, must comply with cloud management standards.

- Backend O&M

  O&M refers to the unified maintenance of infrastructure, cloud platforms, and technical support services as well as the security assurance of the cloud platforms and networks.

  - Monitoring and management

    Cloud platforms and infrastructure are monitored and managed in a unified manner. Infrastructure and technical support services are maintained. Applications are maintained together with project teams. Service logs are collected, stored, and analyzed to associate with events, detect problems, locate root causes, and automatically rectify faults. Higher automated, predictive performance analytics are used to develop potential remedies and automate solution deployment.

  - Collaboration

    Project teams are provided with logs and monitoring results of applications. The application O&M process is automated and visualized for project teams.

  - Standard compliance

    The cloud platform O&M must comply with O&M standards.

  - Support

    Unified service and delivery management enable technical support to be provided for cloud resources and services.

- Global service catalog

  The service catalog is helpful for cloud service management and the unified operations of cloud services and technical support services. It covers service planning, capacity planning, cloud access standards, technical standards, security and compliance management, order management, metering, billing, settlement, service release, SLA management, budgeting, purchase, and supplier management.

  - Management and control

    The service catalog enables the cloud operational team to manage the lifecycle of cloud services and monitor and maintain infrastructure in a unified manner. O&M and delivery teams can manage technical support services and cloud services in a unified manner.

  - Support

    Technical solutions and resource applications from project teams are reviewed. Project teams are also supported in terms of training, requirement management, comment handling, service/resource application management, and technical guidance.

- Collaboration

  The entire service lifecycle needs to be coordinated with delivery and O&M capabilities. Service delivery and O&M teams cooperate to provide their services.

- Standard compliance

  Project team construction solutions, infrastructure deployment, architecture design, and service management must comply with cloud standards.

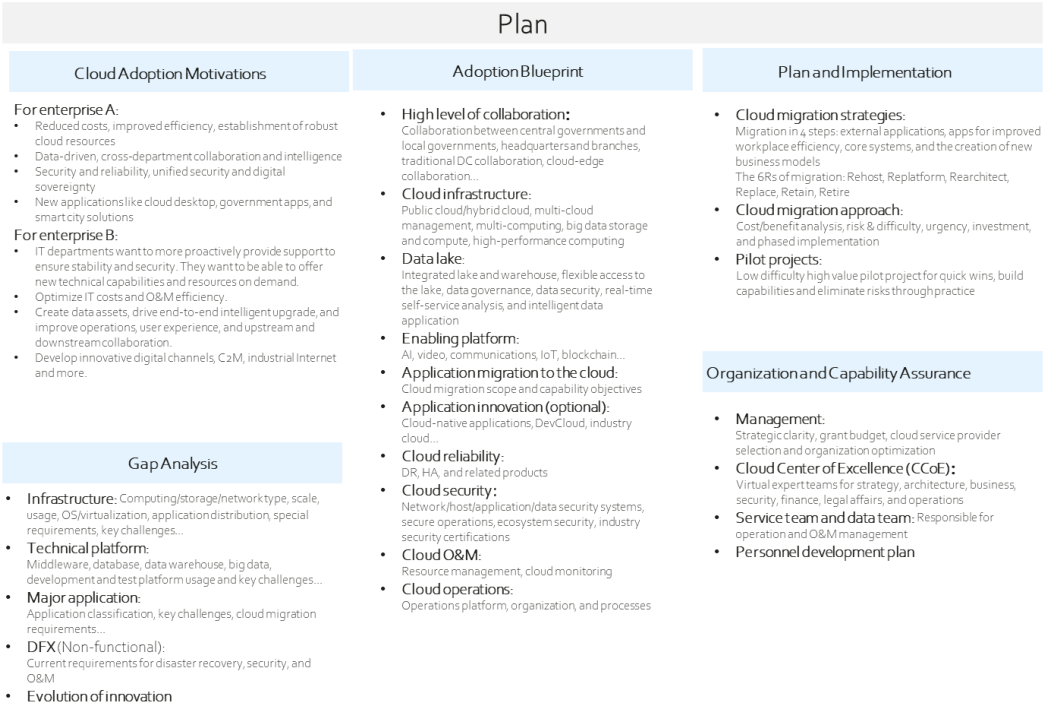
4 Cloud Adoption Framework

4.1 Plan

The purpose of the "Plan" stage is to determine the cloud adoption motivation, IT status quo, gap analysis, and high-level design planning. It includes cloud adoption strategy, cloud adoption blueprint, and cloud migration roadmap.

The cloud plan primarily supports grant budget, organizational optimization, and talent training. It provides an assurance that the people, finances, resources, permission, and security compliance needed for cloud adoption are all in place. In addition, the baseline should be released to relevant organizations after being approved by the management to ensure the right strategy is selected for subsequent cloud adoption. The plan is usually focused on applications with less business critical with high values as pilot projects. Quick wins help build organizational capabilities and reduce risks.

The plan needs to include the following five key tasks.



4.2 Landing Zone Construction

4.2.1 Why Is the Landing Zone Required?

The advantages of the public cloud in terms of security and stability, service quality, execution efficiency, and cost-effectiveness are becoming more widely recognized and accepted by enterprises. More and more enterprises are gradually migrating their application systems to the cloud and preferentially developing future-oriented, cloud native application systems. An era of cloudification is coming. However, in practice, the following challenges are often encountered:

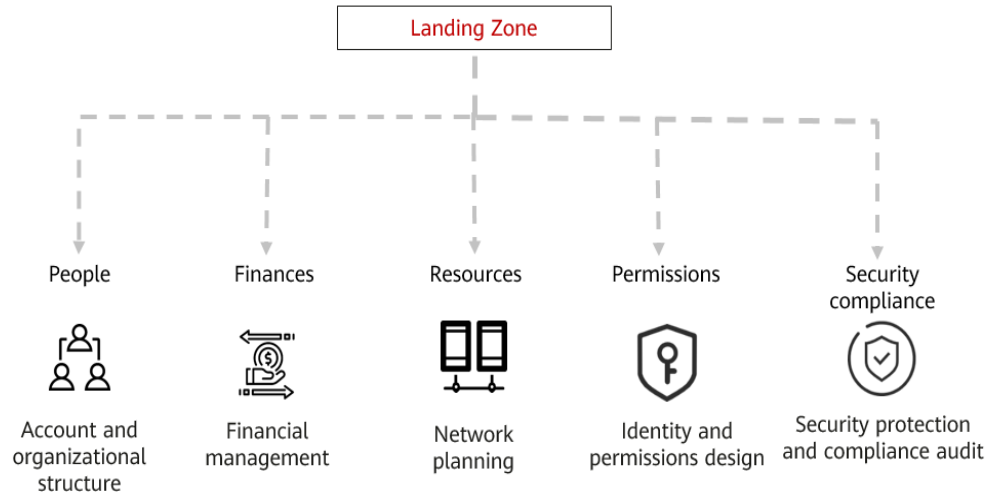
- How to isolate the security and faults of service units (such as BGs, departments, and project teams) to ensure the isolation of cloud resources, applications, and data between service units
- How to flexibly adjust cloud resources
- How to design a network architecture across multiple service units and establish controlled network connection channels
- How to plan the production, development, and test environments
- How to share common resources among multiple service units
- How to centrally manage and control the budgets and costs of each service unit and how to optimize cloud costs
- How to prevent service units from overusing cloud resources
- How to divide user groups and how to set permissions for user groups

To address these challenges, Huawei Cloud has designed a Landing Zone solution to effectively manage service units, personnel, permissions, cloud resources, data, applications, costs, and security. A landing zone is the area where an aircraft, like a helicopter or an airplane, can land safely. Cloud vendors have borrowed this term to describe a place where you can smoothly migrate enterprise service systems on the public cloud. The Huawei Cloud Landing Zone solution helps enterprises build a secure, compliant, and scalable multi-account environment on the cloud where multiple accounts can share resources and there is unified management of the people, finances, resources, permissions, and security compliance.

- People: Unified management of service units, accounts, users, user groups, and roles for multiple accounts
- Finances: Unified management of funds, budgets, costs, invoices, and discounts for multiple accounts
- Resources: Unified O&M, monitoring, and management of cloud resources including computing, storage, network, data, and applications for multiple accounts
- Permissions: Unified management of permissions for cloud resources of multiple accounts based on the principle of least privilege (PoLP)
- Security compliance: Unified management of security compliance in accordance with the security compliance requirements of countries, industries, and enterprises

4.2.2 Landing Zone Reference Architecture

The people, finances, resources, permissions, and security compliance requirements are mapped to the account and organizational structures, financial management, network planning, identity and permissions design, security protection, and compliance audit.



The following describes each module in detail.

4.2.2.1 Account and Organizational Structure

The Landing Zone solution needs a secure, compliant, and scalable multi-account environment on the cloud. It is a best practice to plan the account and organizational structure first.

Huawei Cloud provides a reference structure. It is recommended that you design organizational levels and accounts based on your business architecture, geographic architecture, and IT functions.

- Different organizational levels and organizational units (OUs) are defined on Huawei Cloud based on your service architecture. Independent member accounts can be created for each service OU based on service systems. Independent or shared accounts can be created based on service scales and isolation requirements.
- Different organizational levels and OUs are defined on Huawei Cloud based on the geographic architecture. Independent member accounts can be created under geographic OUs by country or region. On-premises customer relationship management systems and customer service systems can be deployed in these OUs.
- For the central IT department of an enterprise, corresponding OUs are created on Huawei Cloud and the member accounts described in the following table are created based on different IT functions to isolate responsibilities and permissions in the IT management domain and to manage multiple member accounts.

Account	IT Function	Responsible Team	Resource or Cloud Service
Network operations account	Centrally deploy and manage enterprise network resources including network border security resources to unify network resource management and interworking between VPCs under multiple accounts. In particular, centrally manage the Internet-oriented and on-premises IDC equipment room-oriented network ingress and egress.	Network management team and security management team	NAT Gateway EIP VPC Direct Connect Cloud Connect VPN CFW WAF Anti-DDoS
Public service account	Centrally deploy and manage enterprise public resources, services, and application systems, and share them with other member accounts.	Public service management team	NTP server AD server Self-built DNS OBS bucket SWR Collaborative office system
Security	• Serve as the	Security	Services that

Account	IT Function	Responsible Team	Resource or Cloud Service
operations account	enterprise security operations center. • Centrally control the security policies, security rules, and security resources of the entire enterprise. • Set security configuration baselines for other accounts. • Be responsible for the information security of the entire enterprise.	management team	support cross-account security control, such as DEW, SCM, and VSS
O&M and monitoring account	Implement unified monitoring and O&M of resources and applications under each member account, and identify potential issues and send pre-warnings in a timely manner.	O&M team	CBH Grafana Prometheus Third-party O&M and monitoring systems
Log account	Centrally store the run logs and audit logs of other accounts.	Log analysis team and compliance audit team	LTS OBS bucket SIEM system
Data platform account	Centrally deploy the big data platform of the enterprise and collect service data of other accounts to the data platform for storage, processing, and analysis.	Data processing team and business analysis team	Data lake Big data analysis platform Data access service Data governance platform
DevOps account	Centrally manage CI/CD pipelines for the entire enterprise and deploy them across accounts.	Software R&D team	DevCloud Self-built DevOps pipeline
Sandbox account	Test functions and security policies of cloud services.	Test team	On-demand resources and services that need to be tested

In addition to these member accounts, you can create more IT functional accounts such as application integration accounts or collaborative office account as needed.

By default, a master account is created under the root of the organization. It is recommended that no cloud resources be deployed under this master account. You can use this master account to do the following:

- Centrally manage organizations and accounts: Create and manage organizational structures and OUs, create member accounts for OUs, or invite other existing accounts to the organization as member accounts.
- Centrally manage finances: Collect and analyze statistics on the costs of the entire enterprise spent on Huawei Cloud; top up accounts, apply for credit limits, activate coupons on Huawei Cloud and allocate them to member accounts; regularly review the usage of funds, credit limits, and coupons of member accounts and reclaim them in a timely manner.
- Centrally manage organizational policies: Set organizational policies for OUs and member accounts, forcibly restrict user permissions (also for account administrators) under member accounts to prevent security risks caused by excessive permissions. If you apply an organizational policy to a specific OU, the policy applies to all member accounts and lower-level OUs in that OU.

You can use enterprise projects or tags to group resources at a fine granularity under member accounts. For example, you can group application subsystems or sub-products into an enterprise project or tag them on Huawei Cloud. You can also perform cost allocation and find-grained permissions control based on these groupings.

4.2.2.2 Accounting Management

Accounting management allows multiple Huawei Cloud accounts to be associated with each other for accounting purposes. You can create a hierarchical organization and a master account, add member accounts to this organization and associate them with the master account, and use the master account to perform accounting management of associated member accounts.

- **Association between master and member accounts**

  A new Huawei Cloud account can be directly associated with a master account or an existing Huawei Cloud account can be invited for association. On Huawei Cloud, a master account can create organizations that match your organizational structure and create new accounts for or invite existing ones to the organizations.

- **Funds management**

  A master account can allocate its balance and cash coupons to its member accounts for resource provisioning. A member account can use its own balance for resource provisioning.

- **Commercial discount inheritance**

  After member accounts are associated with a master account, they can use the commercial discounts of the master account in their expenditures.

- **Expenditure query**

  A master account and its associated member accounts can log in to Huawei cloud to view their expenditures. The master account can view the expenditure data of its member accounts after being approved.

- **Invoices**

  A master account and its associated member accounts can separately request Huawei Cloud to issue invoices for their expenditures. A master account can request invoices for member accounts.

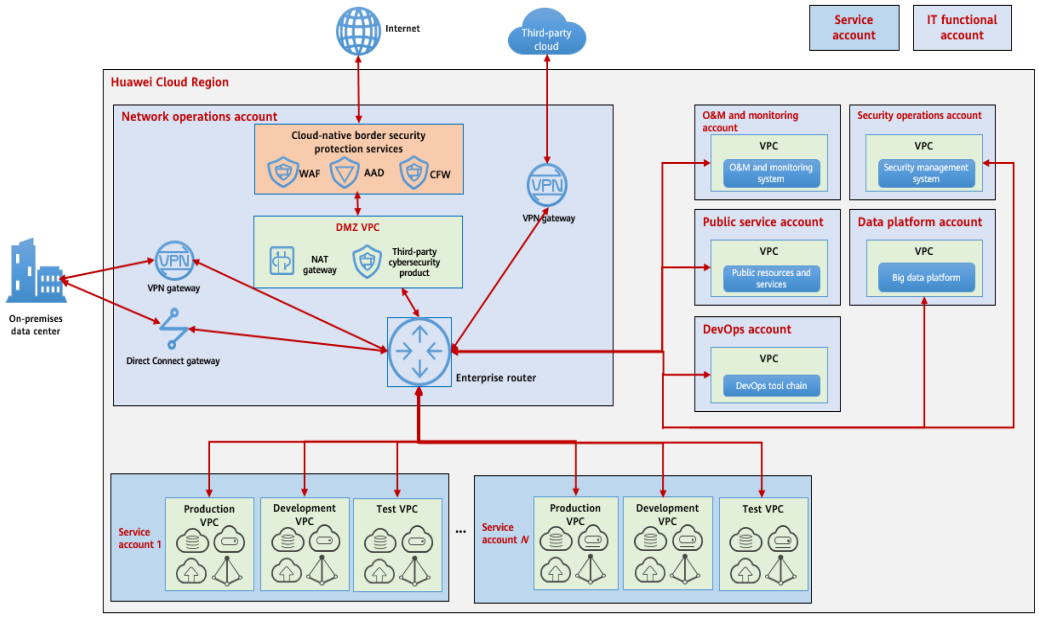
4.2.2.3 Network Planning

The network plan is based on the cloud accounts, there are three scenarios to consider:

- One account and one VPC: Account management and O&M are easy as there is a small team but security is low.
- One account, multiple VPCs: Account management is easy and there is better security. The VPCs are designed for different functions and have different access control rules.
- Multiple accounts or primary account and subaccounts, multiple VPCs: Account management is more challenging as there is a large team with multiple branches and complex services, and meeting security requirements can be challenging too. If multiple VPCs are deployed in the same region, enterprise router should be used to connect them. If multiple VPCs are deployed in different regions, you can use Cloud Connect.

The first and second scenarios are simple. They mainly involve VPC and subnet planning. The following section will be focused on the third scenario.

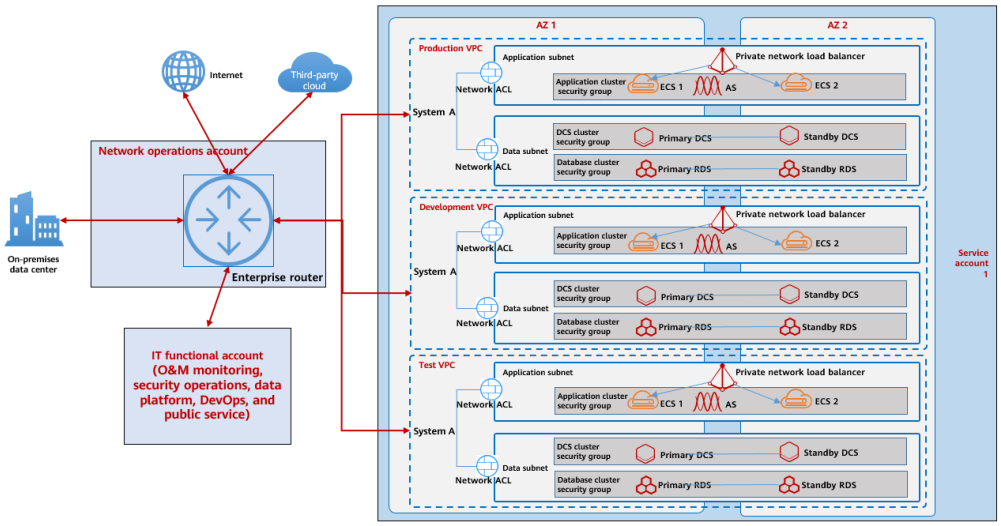
The following figure shows a multi-account, multi-VPC network architecture.



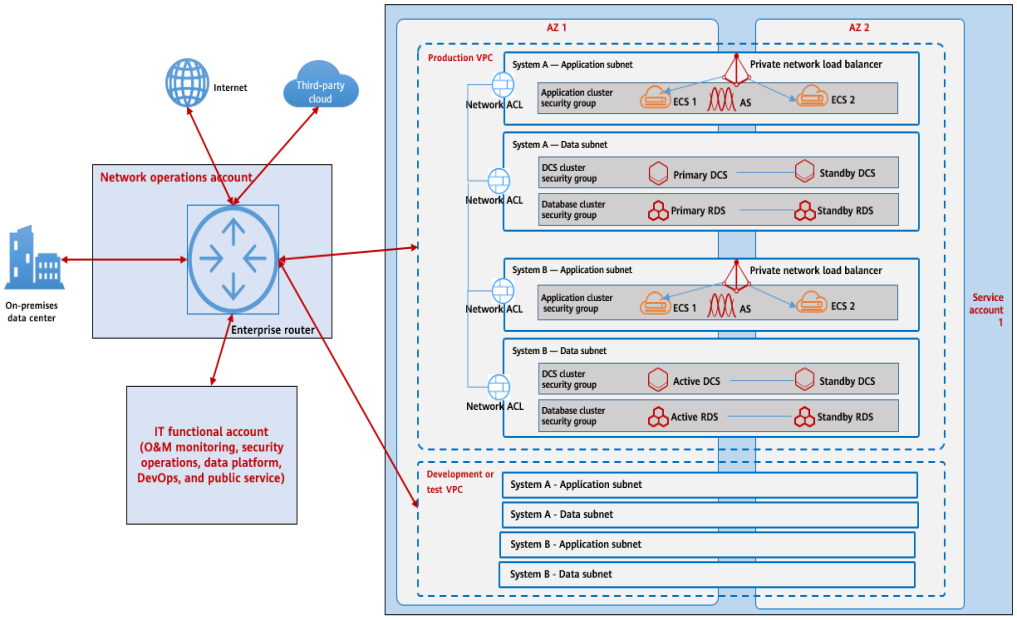
At the core of this architecture is the network operations account, which serves as a hub to connect different accounts. Its enterprise router allows communication between resources of different accounts. The enterprise router lets you configure routes to determine which VPCs can communicate with each other.

- A large service system usually has independent subaccounts. You are advised to create three VPCs: a production VPC, a development VPC, and a test VPC. These VPCs are kept isolated from each other for the service system. Each VPC must have at least two subnets: an application subnet and a data subnet. These correspond to, respectively, the application layer and data layer of the service system. Network ACLs are then used for access control between subnets. Resources such as ECSs, DCS instances, and RDS instances can be added to security groups for instance-level access control based on security group rules.

Application clusters of service systems can be deployed across AZs to ensure high availability at the application layer. The active/standby database clusters and DCS clusters across AZs on Huawei Cloud are used to ensure high availability at the data layer.



- Multiple small-scale service systems that do not have strict security requirements can share the same subaccount. You are advised to create three VPCs: a production VPC, a development VPC, and a test VPC. These VPCs are kept isolated from each other for these small-scale service systems. Different service systems are kept isolated by deploying them in different subnets. Each service system has its own independent application subnet and data subnet. These subnets are associated with network ACLs for access control.



Note: It is a best practice to create an enterprise project for each small-scale service system, associate the resources of each system in the production, development, and

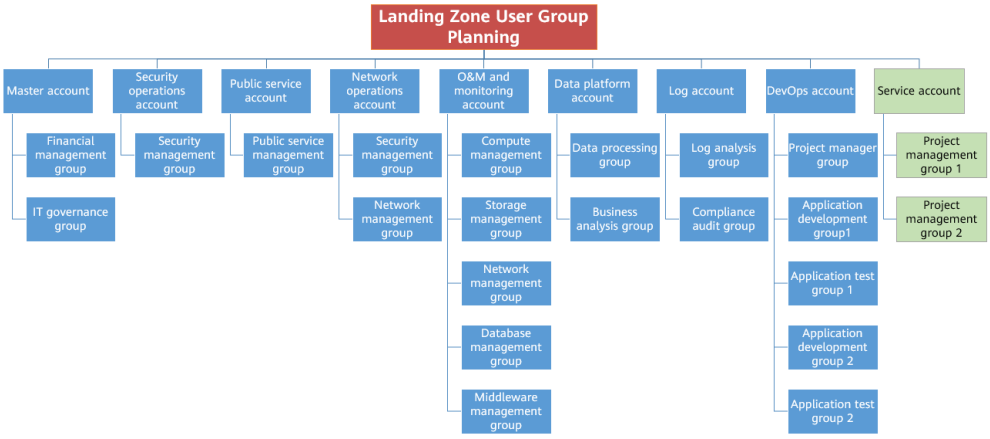
test VPCs with the same enterprise project or tag, and perform cost collection and fine-grained authorization based on the enterprise project or tags.

4.2.2.4 Identity and Permissions Design

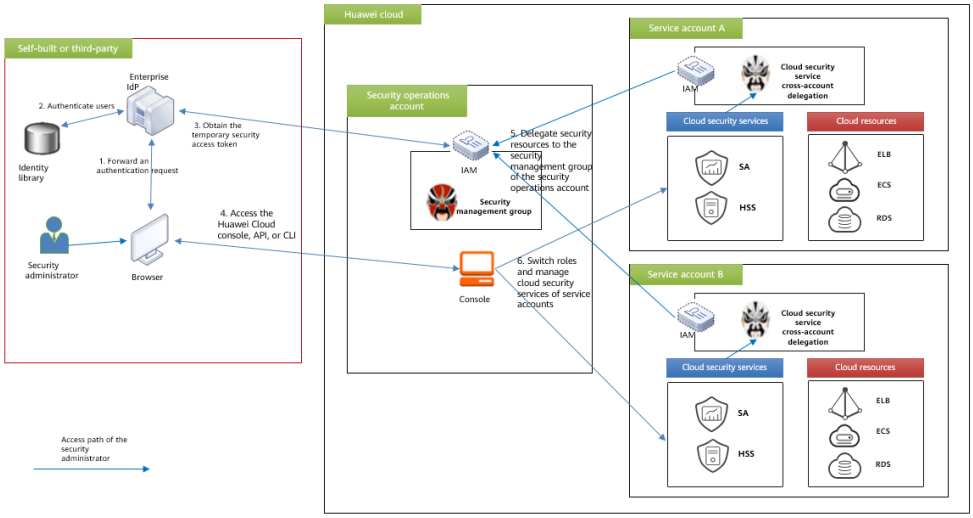
Based on a large number of successfully delivered projects, Huawei Cloud has summarized the following user and permissions management principles:

- Establish trust between the enterprise's identity management system (such as AD) and Huawei Cloud IAM for federated identity authentication so that enterprise employees can use single sign-on (SSO) for the Huawei Cloud console. The enterprise's identity management system can better control the permissions of employees as they are recruited and can revoke permissions in a timely manner for employees who have transferred to different departments or have resigned.
- Do not use Huawei Cloud IAM as the enterprise's user management system. There is no need to create users or user groups on Huawei Cloud IAM for enterprise employees who do not interact with Huawei Cloud.
- Do not share passwords with others. Instead, create an independent user and assign permissions to the user for people who needs to manage or use Huawei Cloud resources. In this way, all operations performed on Huawei Cloud can be tracked and audited.
- Create user groups based on IT functions and add corresponding employees to user groups that match their responsibilities. For example, in terms of resource O&M and management, apply unified management and O&M principles to improve efficiency. Under the O&M and monitoring account, create user groups based on O&M responsibilities. These user groups include a computing management group, storage management group, network management group, and database management group.
- Follow the principle of least privilege (PoLP). Grant only the minimum levels of access or permissions needed to user groups. If the responsibilities of a user group change, adjust the permissions of that user group in a timely manner. To simply operations, grant permissions to user groups rather than individual users.
- The IAM account administrator (with the same name as the IAM user) has high permissions. As such, you should not use this account to access Huawei Cloud directly. Instead, create an IAM user and grant permissions based on the PoLP principle to perform routine management, protecting the security of IAM accounts.

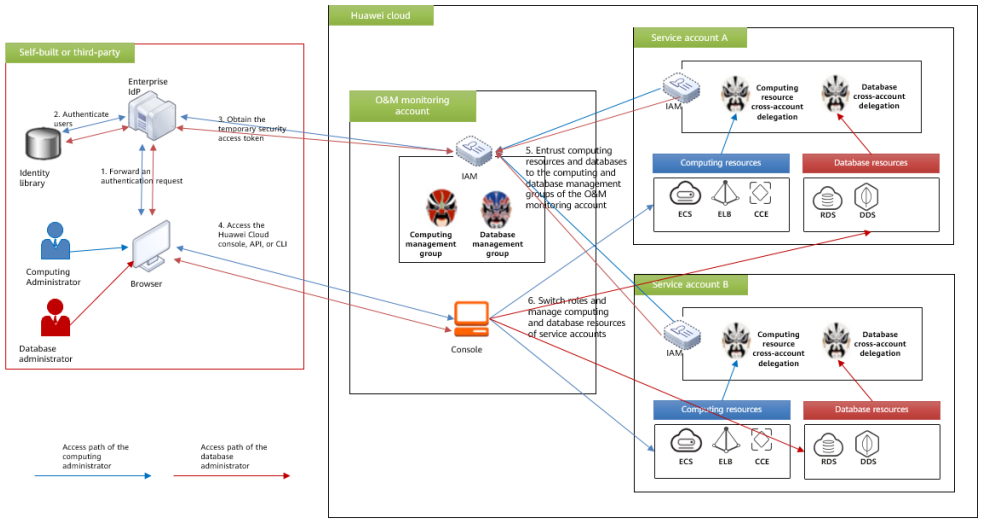
  Based on these principles, user groups are planned for accounts in Landing Zone and corresponding cloud service access permissions are assigned to these user groups based on the PoLP principle on Huawei Cloud. User groups in the enterprise's identity management system are mapped to the user groups on Huawei Cloud so that they have the corresponding cloud service access permissions.



To achieve unified management and control, IT functional accounts of Landing Zone need to access and manage cloud resources under other accounts through cross-account delegation. For example, a security operations account is designed to centrally manage security resources and services (such as SA and HSS) across accounts through federated authentication and cross-account delegation. The security administrator logs in to the console of the security operations account through SSO, switches the role to a service account, and then accesses and manages security cloud services under that service account.



Another similar scenario is that the O&M and monitoring account can access resources in other accounts through federated authentication and cross-account delegation to monitor and manage resources across accounts in an enterprise.



4.2.2.5 Compliance Audit

To ensure that the runtime environment of an enterprise meets the security compliance requirements of countries, industries, and enterprises after cloud migration, Landing Zone provides the following security compliance measures:

- Separation of duty (SOD): A multi-account architecture is used for SOD. Each account is an SOD unit. An enterprise can group accounts based on service units, geographic units, and functional units. The loss of any account does not affect the system as a whole, limiting the impact radius.
- Operation audits: Operation audits are enabled for each account so anytime a resource is accessed by any entity, that activity is logged. All operations can be tracked. Audit logs of all accounts are centrally stored and analyzed.
- Configuration change tracking: Resource configuration recording is enabled for each account to log any resource configuration changes. All resource changes can be tracked. Change logs are centrally stored and analyzed.
- Security guardrail: There are two types of security guardrails. They are security redlines and security baselines. Security redlines, also called preventive security guardrails, forcibly restrict the permissions of member accounts to avoid security risks caused by excessive permissions. The fine-grained authorization mentioned earlier can be used to set security redlines. Security baselines, also called detection security guardrails, require that member accounts meet basic security compliance requirements. For example, MFA must be enabled for the root user and EVS disks must be encrypted. For details about the complete security baseline check items, see the Huawei Cloud official website:

  https://support.huaweicloud.com/intl/en-us/usermanual-sa/sa_01_0022.html
- Unified identity and permissions management: Users, user groups, and permissions in a multi-account environment are centrally managed to enable a user to access resources under multiple accounts. Unified identity and permissions management reduces the workload of permissions management and facilitates the development and implementation of unified permissions standards within an enterprise, reducing security risks caused by improper permissions assignment.
- Unified security control: Security events and risks in a multi-account environment are identified, processed, and analyzed, and events are responded to and

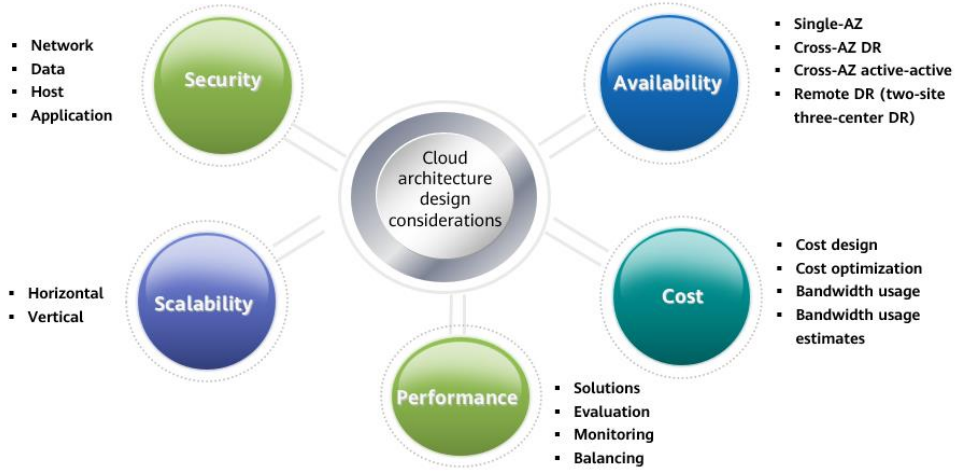
handled centrally. Unified security control reduces the workload involved in security control and facilitates the development and implementation of unified security regulations within an enterprise. The implementation requires cloud security services to support unified security management and control of multiple accounts.

4.3 Cloud Architecture Design

4.3.1 Architecture Design Objectives and Principles

The most important objective of architecture design is to ensure the continuous availability of the system along with the development of enterprise services. Architecture design mainly includes the design of the application architecture and the technical architecture. The application architecture design involves industry-specific features, technology stacks, and enterprise development phases. Designing the technical architecture is more general. In the following sections, we will describe the five aspects of architecture design that affect service continuity the most: high availability (HA), scalability, performance, security, and cost.

- HA: single-AZ availability, cross-AZ DR, multi-active, and remote DR (two-site three-center)
- Scalability: horizontal scaling and vertical scaling
- Performance: performance planning, measurement, detection, and trade-off
- Security: network, data, host, and application security
- Cost: cost design, optimization, and management



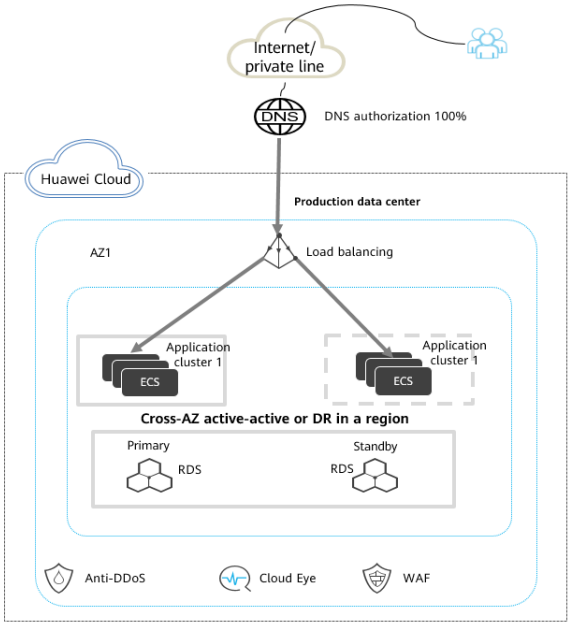
4.3.2 HA

Most Huawei Cloud services have HA deployment options available. They provide you with HA capabilities at multiple layers, including data centers, hardware, data, and self-service. Huawei cloud data centers are deployed around the globe to meet resource requirements for different regions. Each region is divided into multiple availability zones (AZs). Each AZ has independent cooling, fire extinguishing, moisture-control, and electrical facilities, and the failure of one AZ does not affect other AZs. There are four types of HA deployments:

- **Single-AZ HA**: For services that do not require high availability, active/standby or cluster deployment models of cloud services can be used to quickly recover services in the event of a single service node failure. By using automatic fault detection and switchover of nodes in a cluster, single points of failure (SPOFs) are eliminated and service interruptions are prevented.

Key points of the HA design

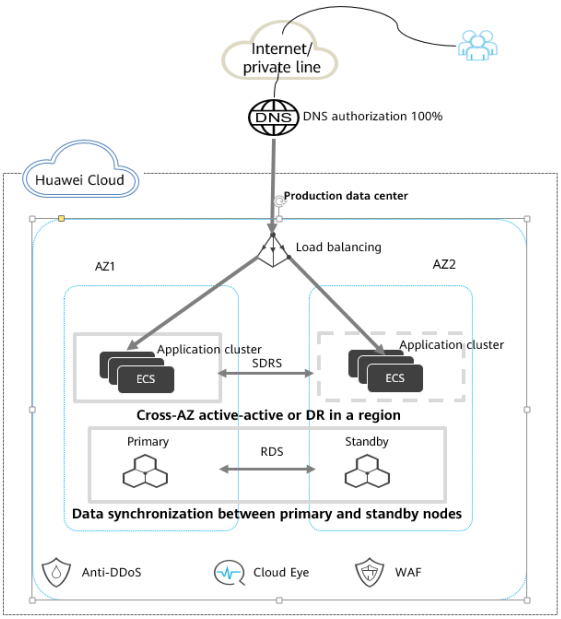
Item	HA Design Focuses
Service HA	Decoupled deployment: Different components are deployed on different ECSs. HA deployment: HA deployment is used for all nodes. If HA (active/standby or cluster) deployment is not supported, an emergency solution must be available, for example, cloud server backup (via CSBS) and an emergency environment. Layered deployment: Services are deployed separately on the web layer, service layer, and database layer. Auto scaling: The Auto Scaling service can be leveraged to adjust compute resources to deal with changes in service volume.
Cloud service HA	Network access layer: Private line: Active-active or active/standby HA VPN: private line (active) + VPN (standby) or VPN (active) + VPN (standby) ELB: Multiple ECSs run at the ELB backend to ensure system availability and scalability. Health check is enabled. ELB is a potential fault point in the system and needs to be monitored by Cloud Eye. NAT gateway: If a large number of ECSs need to access the Internet, use SNAT to prevent too many ECSs from being exposed to the Internet. Selection of cloud service types (RDS, DCS, and more): Production services must use active/standby or cluster deployment. For example, RDS DB instances must be deployed in primary/standby mode, and Redis must be deployed in clusters. Self-built services on ECSs must also meet this requirement, for example, deployed using Redis Cluster.
Data reliability	Data backup and restoration solutions are available. For example, VBS and CSBS are used to back up ECS data, and an RDS backup policy is enabled. Critical data is backed up to other regions or offline IDCs. Data backup and recovery reliability, and emergency drills are periodically verified.



- **Dual-AZ (intra-city) HA**: For services that require high availability, you can deploy services in multiple equipment rooms in the same city. This way, service continuity is guaranteed even if the network, physical device, or power supply of one equipment room fails. Huawei Cloud users can deploy services across AZs. AZs are isolated from each other, so if one AZ fails, you can switch services to another AZ to quickly recover services. Most cloud service products have corresponding capabilities. Select your desired capabilities during purchase and complete deployment.

  Key points of the HA design:

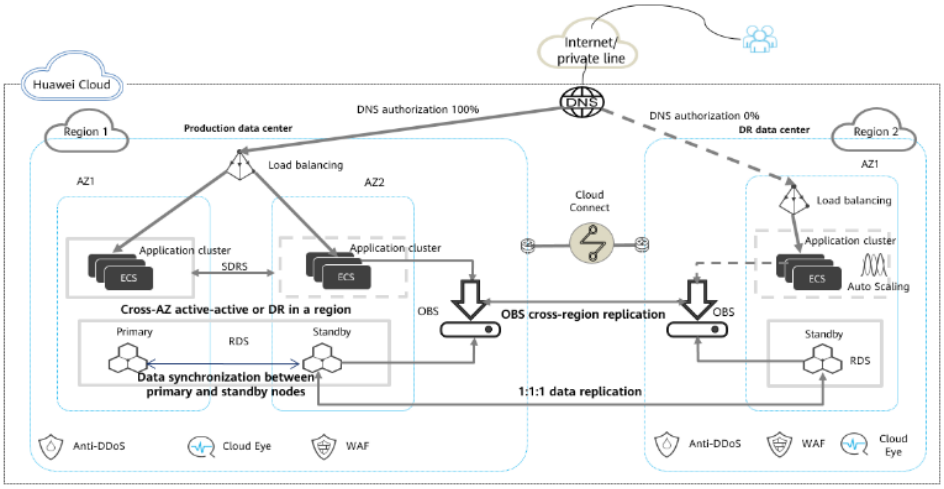
  - Service modules: For services supporting cluster deployment, resources are deployed in two AZs, with the loads balanced using ELB. For single-node ECSs, SDRS is used for AZ-level DR.
  - Cloud service HA: Primary and standby nodes are deployed in different AZs.
  - Database synchronization: RDS is used, and RDS DB instances are deployed in primary/standby mode across AZs and the data kept synchronized.
  - DR switchover: If an AZ fails, RDS database services automatically switch to the standby databases. Application services can be taken over by the DR servers automatically or in just a few clicks.
  - DR drills: Users can perform DR drills with just a few clicks.



- **Two-site three-center HA**: For some ultra-large or commercial systems that demand extra robust protection, the dual-AZ (intra-city) HA solution is still insufficient. It cannot guard against regional disasters, such as an earthquake or a flood. A remote equipment room is required for that. So in a two-site three-center HA solution, you add a remote DR equipment room to the intra-city solution. That way, you are protected even in the event of a regional disaster.

  Key points of the HA design:

  - The production centers and DR center are deployed in two different regions of Huawei Cloud.
  - The production centers are deployed in two AZs, and the DR center is deployed in one AZ.
  - RDS DB instances are deployed in both the production centers and DR center for 1:1:1 primary/standby replication.
  - Configurations, logs, snapshots, and backups generated at the production and DR centers are replicated across regions using OBS.
  - If one production AZ fails, services are switched to the other AZ, and a database switchover is performed.
  - If both production centers fail, a database switchover is performed, and by modifying DNS configuration, 100% of the user traffic is directed to the DR center.
  - After the production centers recover, a database switchover is performed, and DNS directs 100% of the user traffic to the production centers.
  - To improve the utilization of the DR center, some read-only or analysis services can be distributed to the DR center.



4.3.3 Scalability

Compared with traditional IDCs, the cloud has more abundant resources and powerful scalability. There are different types of scaling designed for different service scenarios.

- Vertical scaling: For monolithic applications, independent applications, and stateful applications, hardware sometimes needs to be rapidly upgraded to handle changes in demand as services develop. For example, during promotions, these applications often require many times more resources than normal. In this case, companies can use a UI or open APIs to quickly upgrade resources by adding more vCPUs, memory, bandwidth, and disk space, to handle increasing workloads. After the activities have ended, they can restore the resources to the original specifications to keep costs down.
- Horizontal scaling: For distributed applications, stateless applications, and rapidly changing applications, resources allocated in fixed ratios struggle to keep up with rapid changes in demand. With horizontal scaling, applications can take advantage of abundant resources of the cloud to rapidly scale out or in based on pre-configured scaling policies to handle traffic fluctuations both during and after promotional activities. In addition, these applications can consistently, flexibly get more resources as workloads keep rising.
- Extreme scaling: To respond to unexpected traffic bursts, such as if there are major breaking news, a system needs to support extreme expansion capability. When such events take place, it may need to rapidly add thousands of compute cores. In this case, scaling on the cloud is the best choice.

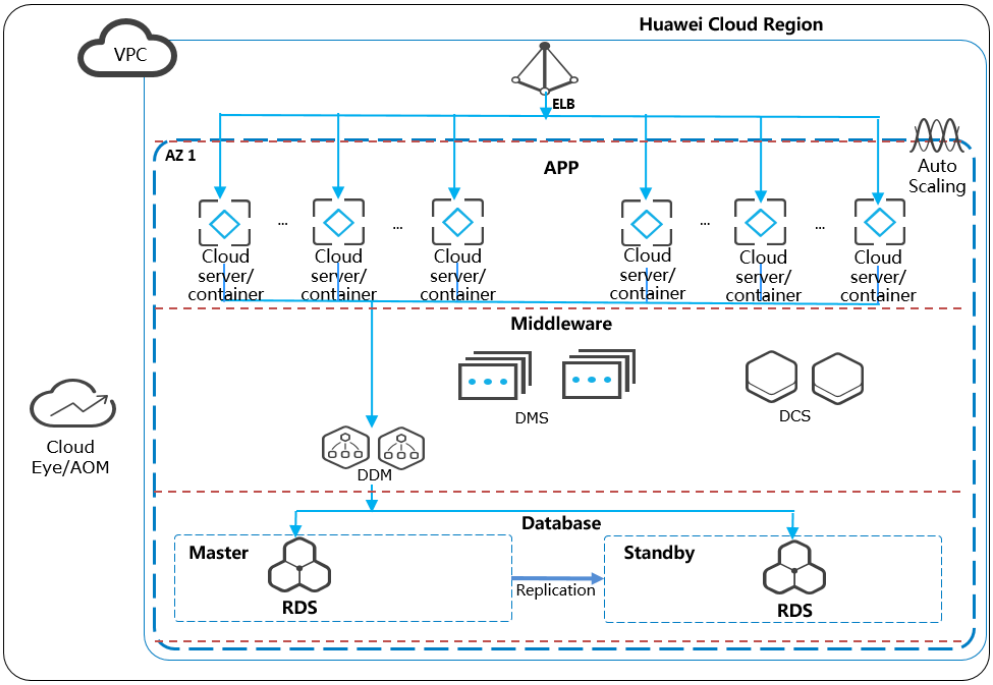
There are several different ways to configure scaling:

- On a schedule: You can create a scheduled task to scale in or out resources at a specified time.
- Using metrics. You can create an alarm-triggered task to monitor resource performance metrics such as CPU usage and average network traffic. When the monitored metric reaches the specified value, an alarm is triggered and resource scaling is performed.
- Based on a configured range. You can configure upper and lower capacity limits. When the number of compute instances is below the lower limit or above the

upper limit, the system automatically adds or removes instances so that the number of instances stay within a certain range.

- Manually. You can add, remove, or delete existing resources manually.

Scalable capabilities can be designed by layer. The figure shows the scalability of Huawei Cloud services at different levels. The scalable design of each layer is as follows.



- Application layer: If this layer uses a microservice architecture and container-based application deployment using Cloud Container Engine (CCE) on Huawei cloud, with the auto scaling capability of CCE, applications can automatically scale out and in on demand. During the auto scaling triggered by alarms from AOM, service pods are automatically added or removed in response to workload fluctuations. If this layer is deployed using Huawei Cloud Elastic Cloud Server (ECS), applications can automatically scale out or in based on scaling policies configured on Auto Scaling.
- Message middleware layer: Huawei Cloud Distributed Message Service (DMS) for RabbitMQ Premium instances are deployed in clusters, and can scale up or down as the message volume and workload changes.
- Cache middleware layer: The master/standby Distributed Cache Service (DCS) for Redis instances from Huawei Cloud can scale up or down as the hot data volume increases or decreases.
- Database middleware layer: The distributed database middleware uses Huawei Cloud Distributed Database Middleware (DDM), which is deployed in a cluster. With the increase of database services, the DDM cluster specifications can be smoothly expanded to cope with more database processing.
- Database layer: Huawei Cloud Relational Database Service (RDS) supports smooth expansion of read-only database instances for scenarios where much more data is read than written. By working with DDM, multiple instances can be scaled out. Data in large tables is horizontally split and evenly distributed to database instances, improving database capacity and performance. In addition,

GaussDB databases use an architecture with decoupled storage and compute to support minute-level horizontal expansion and reduce service interruptions.

4.3.4 Performance

Performance is a key metric for any software system. It is also an important part of cloud design. In performance design, scalability must be considered as it is vital to high performance. In addition, solution selection, performance measurement, performance monitoring, and performance trade-off must also be considered.

The following factors affect the performance of cloud applications:

- Compute latency: the wait time between operations and a direct reflection of cloud computing performance
- Network throughput: the rate at which data is processed
- Transmission throughput (bytes/second or bit/second): a key measure of performance
- Storage input/output operations per second (IOPS): a measure of data transmission
- Data concurrency: the ability to run multiple programs at the same time

- Solution selection
  - Select and combine the solutions that best suit your needs.
  - Upgrade solution selection methods and optimize the selection of resources and configurations through data.
- Performance measurement
  - Configure performance measurement and monitoring metrics.
  - Enable performance tests to be triggered automatically after the fast-running test is complete.
  - Use data visualization to identify performance issues, hot topics, waiting states, or low utilization.
- Performance monitoring
  - Determine the monitoring scope, metrics, and thresholds.
  - Create a full view from multiple dimensions.
- Performance tradeoffs

  Strike a balance in the architecture for better performance, for example by using compression or caching techniques.

4.3.5 Security

Cloud customers need the following security capabilities:

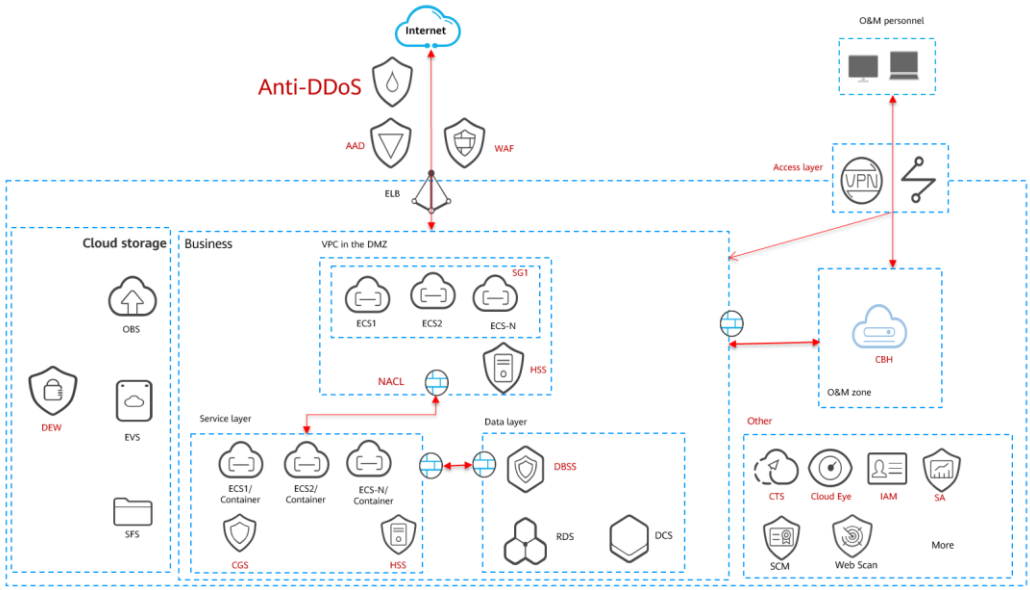
- Service continuity, such as network attack blocking, intrusion prevention, and legal compliance.
- Data confidentiality, such as defense against unauthorized access from external parties, insiders, and cloud service providers.
- Manageable O&M, such as security policies, risk identification and handling, and operation audit and tracing.

The architecture is designed to enhance the following aspects of network security:

- Region boundaries

  - Boundary protection: controlled connections, illegal private connection prevention, illegal external connection prevention, and wireless access restriction
  - Access control policies
  - Intrusion prevention: known threat prevention, unknown threat prevention, and intrusion audits
  - Malicious code prevention: malicious code detection and spam filtering
  - System audits: user behavior audit, security event audit and analysis
- Network communications
  - Network architecture: performance redundancy, link redundancy, device redundancy, and partition isolation
  - Communication and transmission: The encryption technology is used to ensure data confidentiality and integrity during transmission.
- Computing environment
  - Identity identification: identity uniqueness and credential complexity
  - Access control: user permissions management and redundant account clearance
  - Security audit: user behavior audits and audit process protection
  - Intrusion prevention: intrusion detection, closing unused ports, and vulnerability scans
  - Malicious code detection and blocking
  - Image integrity check and snapshot protection
  - Data integrity and confidentiality during transmission and storage.
  - Secure data destruction: When service application data is deleted, all copies in the cloud storage need to be deleted too.
- Management center
  - System management: identity authentication and system configuration for system administrators
  - Audit management: permissions management and operation audits
  - Security management: permissions management and operation audits
  - Centralized management and control: independent secure partitions, network monitoring, centralized log audit, and security event awareness

The following figure shows the security architecture on the cloud.

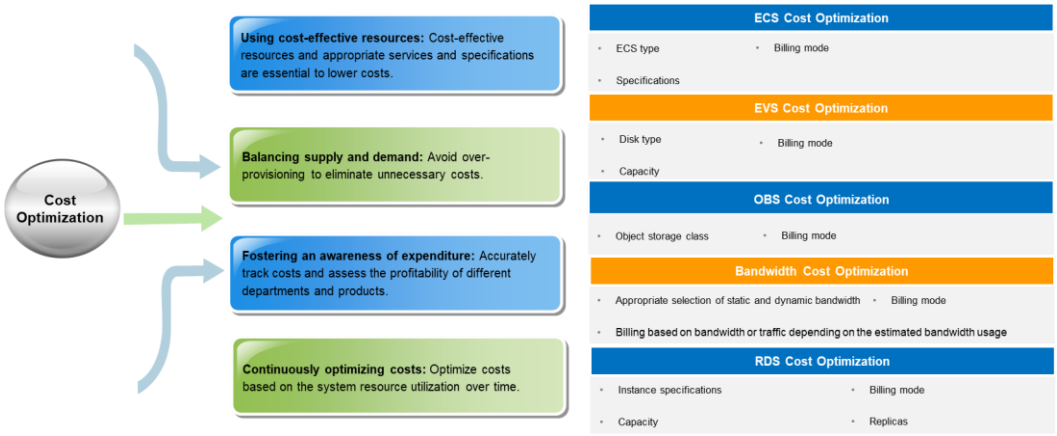


Abbreviations:

- AAD: Advanced Anti-DDoS
- Anti-DDoS: traffic cleaning service
- WAF: Web Application Firewall
- ELB: Elastic Load Balance
- OBS: Object Storage Service
- EVS: Elastic Volume Service
- SFS: Scalable File Service
- SG: security group
- NACL: network access control list
- HSS: Host Security Service
- CGS: Container Guard Service
- DBSS: Database Security Service
- DEW: Data Encryption workshop
- RDS: Relational Database Service
- DCS: Distributed Cache Service
- CBH: Cloud Bastion Host
- CTS: Cloud Trace Service (used for auditing)
- Cloud Eye (used for monitoring)
- IAM: Identity and Access Management (used for unified authentication)
- SA: Situation Awareness
- SCM: SSL Certificate Manager

4.3.6 Cost

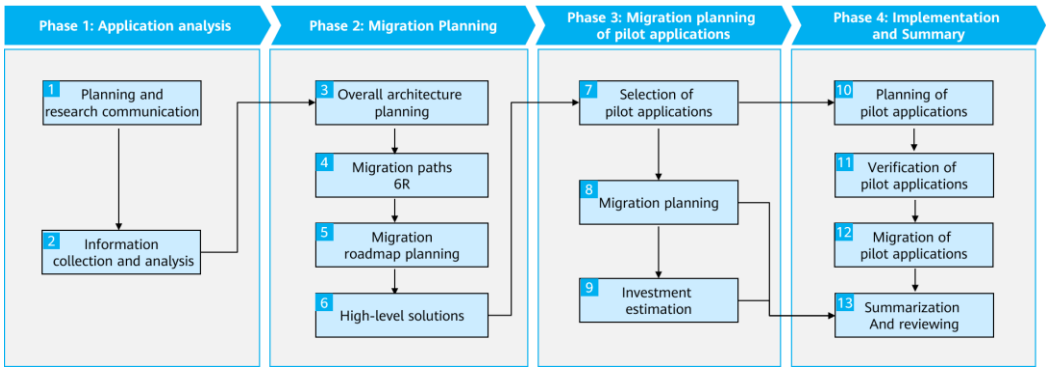
On-demand resource usage, usage-based billing, elastic scaling, and resource utilization determine the costs on cloud. The following figure shows the principles of cost optimization design.



4.4 Adopt

4.4.1 Application Migration to the Cloud

Based on Huawei's successful practices and experience in serving a large number of customers, Huawei Cloud had identified four main phases in the migration process. An understanding of these four phases makes it easier to ensure a smooth migration.



- Phase 1: Application analysis. The services and functions provided need to be clearly identified, along with the technology stacks, deployment modes, SLAs, and dependencies. You need to know where the applications come from and how the O&M is handled.

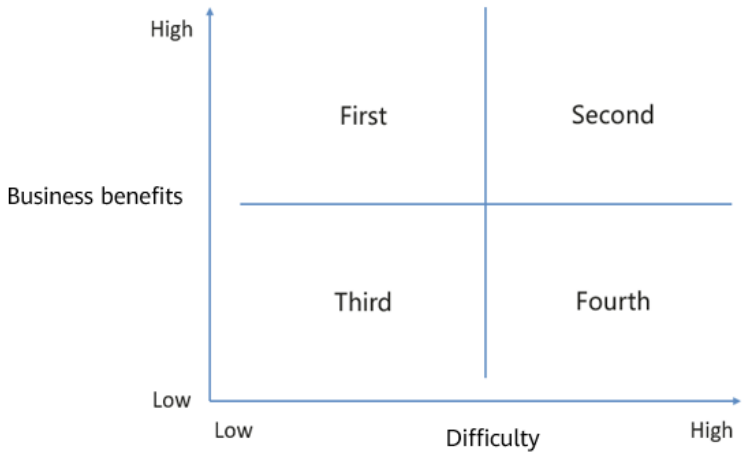
  The research includes, but is not limited to:

  - Application architectures
  - Application modules, internal and external dependencies, and the languages and frameworks used

- Application hosts, including host configurations, specifications, operating systems, total data volume, NICs, HA deployment, as well as DR and backup requirements
- Databases, including the types and versions, how much data needs to be handled, and the performance and HA requirements
- Middleware, including the types (such as message middleware and cache database), versions, as well as cluster scale and capacity

- Phase 2: Migration planning. Based on the information collected in phase 1, assess applications using the 6R model (Rehost, Replatform, Refactor/Rearchitect, Repurchase, Retire, and Retain) and following the general principles described here. Analyze application readiness and benefits and identify proper migration paths.
  - General principles for migrating applications to the cloud
    - For third-party SaaS services, as long as they can keep up with service development, keep them unchanged.
    - For purchased software deployed on hosts, rehost or replatform them to migrate them to the cloud.
    - For applications that are not ready for cloud, for example, applications with incompatible host OSs, applications with outdated, unavailable, or unsupported components, or applications whose benefits from migration seem elusive, keep these applications on premises.
    - For self-developed applications, rehost, replatform, or rearchitect them, depending on the technical stack requirements of customers.
    - For components such as databases and middleware, if there are available cloud services, replatform them, use pay-per-use billing for the best cost-effectiveness, and so you don't have to worry about the O&M.
  - Principles for prioritizing applications for migration

    Prioritize applications to be migrated based on the expected benefits and on their readiness. Migrate applications in the order shown in the following figure. Start with those that are easy to migrate and can benefit much from migration, and give low priority to applications that are more difficult or that will benefit little.



Considerations for business benefits include but not limited to:

- Improving performance, increasing efficiency, and reducing costs

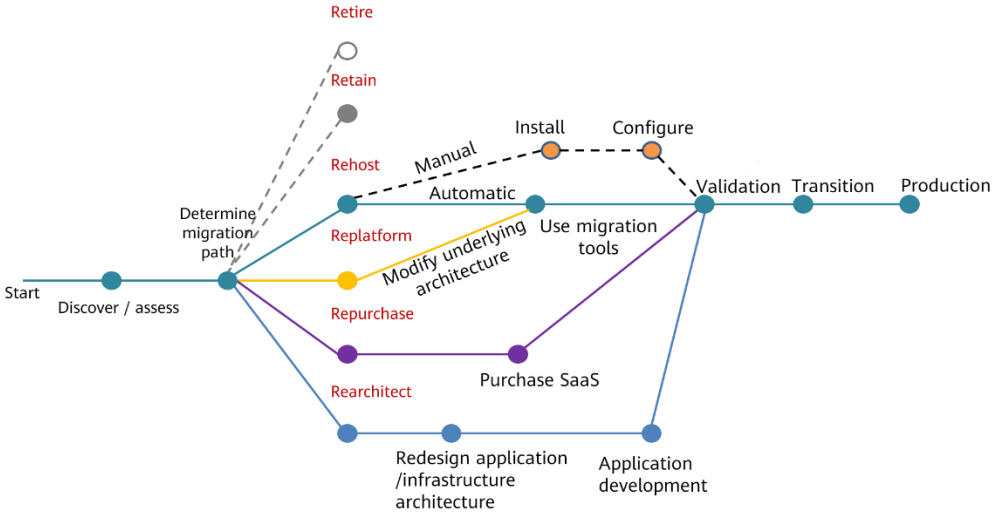
- Changing service requirements
- Enhancing user experience
- Serious architecture issues
- Auto scaling requirements
- Compliance

Considerations for cloud readiness include but are not limited to:

- Service complexity
- Maturity of business and IT design
- Dependency
- Organizations and capabilities

- Phase 3: Migration planning of pilot applications. Select pilot applications based on the overall migration path planned for your applications, plan the migration of the pilot applications, and estimate the costs involved. Proper planning streamlines the architecture and helps ensure a successful experience, which, in turn, supports future migrations.
- Phase 4: Implementation and summary. After developing the migration plan and getting budget approved, migrate the pilot applications to identify what the company should focus on in terms of technology, organization, process, talent, and cost. This phase is critical to accumulating practical experience. This phase helps companies more confidently migrate applications to the cloud and benefit more from the migration.

  When using the 6R methodology, the migration paths are as follows:



The following section covers the Rehost, Replatform, and Rearchitect migration paths.

4.4.1.1 Rehost

Rehost, also known as "lift and shift", is the most common way to migrate applications to the cloud without changing the running environment of applications. It is usually used for Physical to Virtual (P2V) and Virtual to Virtual (V2V) scenarios. It can help companies quickly migrate applications such as SAP, ERP, and CRM to the cloud.

Huawei Cloud provides three rehosting solutions:

- Application redeployment

  In this solution, applications can be redeployed on ECSs or BMSs. This solution is ideal for stateless applications that do not involve data migration. The OSs of cloud servers can be changed as needed, for example, if an old OS is no longer supported. This solution is recommended when a new OS is required, but this means the applications will be offline for a time.

- Image import & export

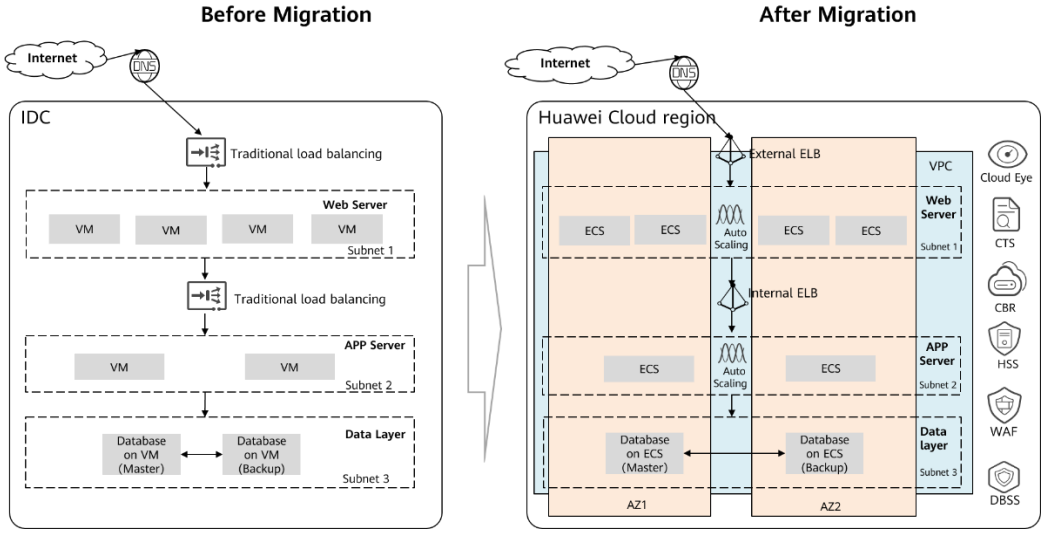
  By exporting system images of source servers and then importing those images to the cloud as private images, you can quickly create cloud servers with the same OSs and other details as your legacy servers. This solution is a good choice when you need to migrate on-premises servers that do not have too much data on them. The servers will have the same OSs before and after the migration, but there will be a fair bit of downtime.

- Server Migration Service (SMS)

  SMS can migrate applications to the cloud and synchronize incremental data to minimize the downtime. However, the OS cannot be upgraded during the migration.

Object	Migration Method	Pros and Cons
Virtual/physical servers	Redeployment	• Easy OS change • Long downtime
	Image import & export	• OS consistency • Long downtime
	SMS	• OS consistency • Long downtime

Take a typical three-layer application architecture as an example. The following figure shows how the architecture is different before and after the migration.



Rehost has the following benefits:

- The application architecture is consistent before and after the migration, so you know the original technology stack still works. Rehosting ensures the migration of your applications can go smoothly.
- If the databases were built using Huawei Cloud ECS, the database licenses can be reused in commercial database scenarios to save money.
- Applications are deployed across AZs, so you can configure DC-level HA.
- With Huawei Cloud ELB and Auto Scaling, services can be flexibly scaled to adapt to workload changes.
- ELB replaces traditional offline hardware load balancing devices and the network ACLs replace traditional hardware firewalls, further reducing the hardware investments required.
- The O&M is simpler. Cloud Eye provides comprehensive O&M monitoring of cloud infrastructure, and LTS provides quick collection and analysis of application logs.
- The reliability is enhanced. CBR backs up cloud servers for restore or other server issues.
- The security is hardened. HSS protects cloud servers, WAF filters web application traffic, and DBSS hardens cloud databases.

4.4.1.2 Replatform

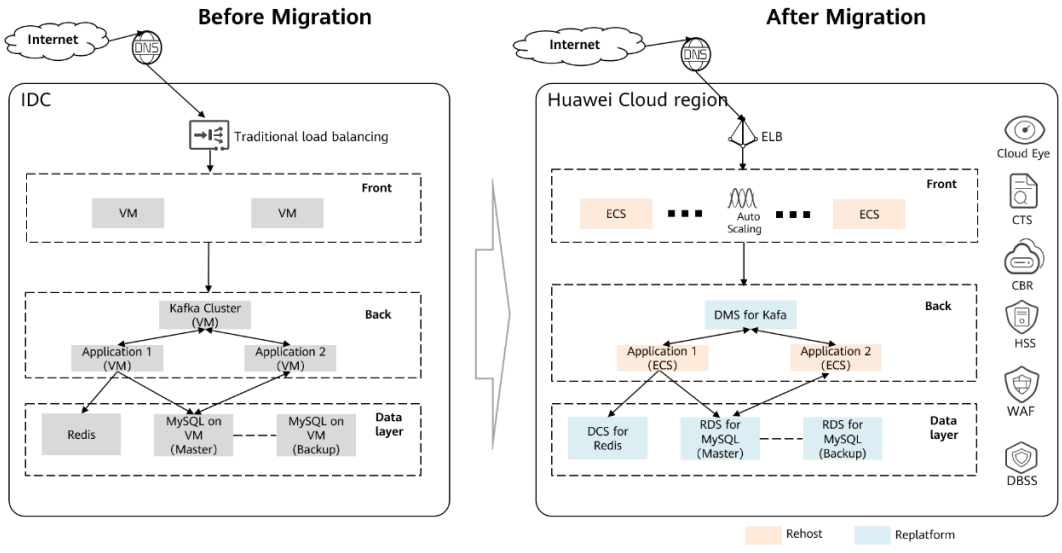
Replatforming involves upgrading an application from its existing legacy platform to a more modern cloud platform. It means replacing traditional application components (such as databases and middleware) with Huawei Cloud services, but without changing the core architecture of applications. For example, you can replace relational databases with cloud database services from Huawei Cloud, replace self-built message middleware with message queue services provided by Huawei Cloud, and replace self-built cache databases with cache database services on Huawei Cloud. This makes management less expensive and makes applications more efficient and scalable.

Huawei Cloud provides the following solutions to migrate customers' self-built databases and middleware or those on third-party cloud platforms:

Object	Type	Source	Target	Migration Method	Pros and Cons
Database	SQL Server	Self-built/DBaaS	Huawei Cloud RDS for SQL Server	Data Replication Service (Huawei Cloud)	The target is RDS, and the downtime is a few minutes.
	MySQL		Huawei Cloud RDS for MySQL	Data Replication Service (Huawei Cloud)	The target is RDS, and the downtime is a few minutes.
	PostgreSQL		Huawei Cloud RDS for PostgreSQL	Data Replication Service (Huawei Cloud)	The target is RDS, and the downtime is a few minutes.
	MongoDB		Huawei Cloud Document	Data Replication Service (Huawei	The target is RDS, and the downtime is a few

Object	Type	Source	Target	Migration Method	Pros and Cons
			Database Service	Cloud)	minutes.
Middleware	Redis	Self-built/Cloud services	Huawei Cloud Distributed Cache Service (DCS) for Redis	DCS-Migration	The target is DCS.
		Self-built/Cloud services	Huawei Cloud Distributed Cache Service for Redis	Redis-port	Offline export and import
	Kafka	Self-built	Huawei Cloud Distributed Message Service (DMS) for Kafka	MirrorMaker	Only data in Kafka clusters can be synchronized. Consumer groups or consumption progress cannot be synchronized.

Consider the following typical architecture. A company uses Kafka message middleware to mask performance inconsistencies between front- and back-end applications. The applications get decoupled. They use a Redis database cache for hot data and MySQL databases for core service data. In a traditional IDC, the company needs to build their own middleware and databases, implement HA deployment, backup and restore solutions, and maintain corresponding components. Deployment can be inefficient, O&M expensive, and capacity expansion difficult.



Huawei Cloud provides cloud services that let companies deploy middleware and database components on the cloud. These services simplify middleware and database deployment and O&M. Companies can enjoy the following benefits:

- Instances can be provisioned in just minutes, so they can take advantages of pay-per-use middleware and database services.

- Cloud services can be deployed in HA configurations. They can deploy active and standby MySQL instances and Kafka/Redis clusters, and use cross-AZ deployment for data center-level HA.
- There is pay-per-use cloud middleware, such as message middleware and cache database, and diverse instance specifications. Easy capacity expansion is easy so they can start small and grow big
- They don't have to worry about O&M of middleware and databases, which saves money on O&M.

4.4.1.3 Rearchitect

Rearchitecting, also called "application refactoring", involves re-imagining how an application is designed and developed, typically using cloud native features, such as transitioning from monolith to microservices. This is typically driven by a strong need to add features, scale, or performance that would otherwise be difficult to achieve in the application's existing environment. This strategy tends to be the most expensive, but will better meet future expansion requirements in the long run. Splitting monolithic applications into microservice requires in-depth involvement of business personnel.

4.4.1.3.1 Traditional Application Architecture Issues

Common issues of traditional monolithic applications include but are not limited to:

- Low resource utilization, and complex deployment and O&M
  - Coarse system granularity, complex scheduling, and poor scalability
  - A lack of systematic application standards, snowflake servers (servers that are fragile and hard to replicate due to environment or component upgrades)
  - A lack of unified or comprehensive application monitoring and O&M, since O&M personnel are busy in maintaining the underlying infrastructure
- Complex application architecture
  - Too many functional modules, an overly complex architecture
  - Coupling between applications and states, which makes expansion difficult
- Slow application iteration
  - Complex application development where developers need to keep track of all the details of an application architecture from service governance (such as rate limiting, circuit breaker, and downgrade) and data access, to message communication
  - Command-based APIs that means developers have to focus too much on the small execution details
  - Manual testing that decelerates application release
  - Applications that iterate too slow, so services cannot be developed fast enough

4.4.1.3.2 Cloud Native Application Architecture

As defined by Cloud Native Computing Foundation (CNCF), cloud native technologies empower organizations build and run scalable applications in dynamic environments like public, private, and hybrid clouds. Containers, service meshes, microservices, immutable infrastructures, and declarative APIs exemplify this approach.

Cloud native application architectures involve:

- Microservices

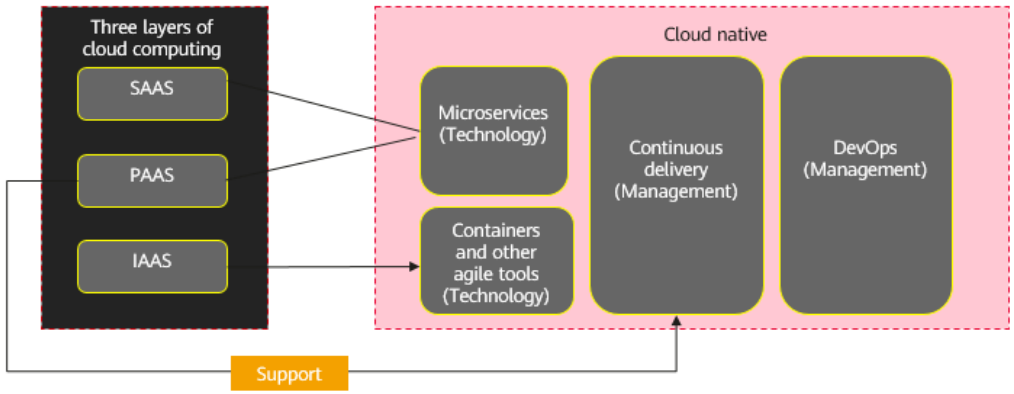
  In a microservice architecture, an application is broken down into various small service components. This process simplifies the complex application architecture.

- Containers

  Containers facilitate the implementation of microservices and help resolve the problem of low resource utilization.

- DevOps

  Containers enhance the efficiency of software development and system O&M, promote the maturity and development of the DevOps system, and help resolve problems such as long application iteration period and complex deployment and O&M.

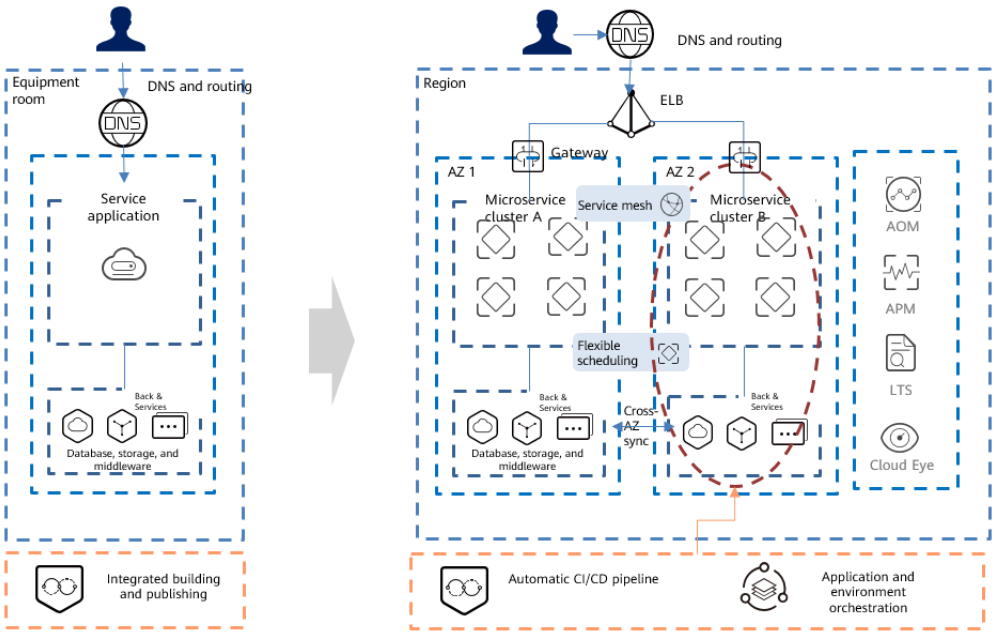


4.4.1.3.3 Application Reconstruction

Application reconstruction involves microservices, containers, and DevOps processes:

- Microservice-based reconstruction includes but is not limited to analyzing the current architecture, in-depth involvement of various business departments, dividing microservices based on required capabilities, defining interfaces between services, formulating development specifications for microservice reconstruction, and managing microservices.
- Container-based reconstruction includes but is not limited to determining the reconstruction scope, identifying application dependencies, creating container images, and orchestrating and managing containers.
- DevOps-based reconstruction includes but is not limited to analyzing R&D process, tools, and peripheral dependencies, identifying gaps, selecting pilot applications, and conducting training and promotion of agile principles.

The following image shows a typical architecture where companies use VMs to carry individual service applications and lack a unified service governance tool and pipeline platform. As a result, resource utilization is low, capacity expansion difficult, R&D inefficient, rollout slow, and O&M expensive.



Application reconstruction can bring the following benefits:

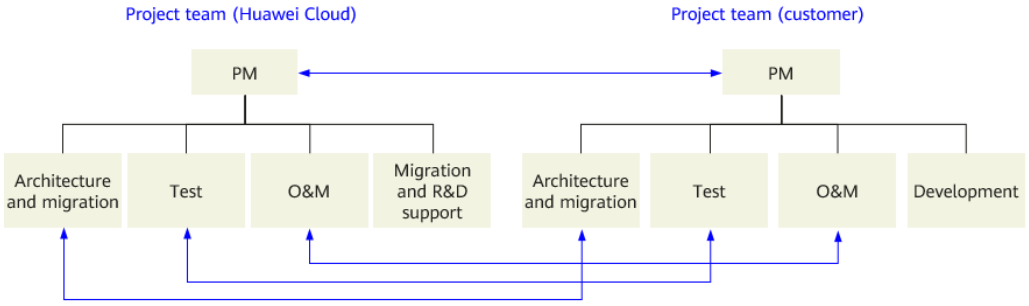
- Microservice-based reconstruction: Monolithic applications can be split into small, independent, light, and loosely coupled microservice clusters based on the AKF Scalability Cube model, front-end and back-end separation, and application and status separation. Service mesh technologies, such as Istio, can be used to manage microservices across programming languages. Applications, databases, and middleware are deployed in multiple AZs. Traffic is distributed through ELB. The underlying data and statuses are synchronized, and applications are deployed in active-active HA mode.
- Container-based reconstruction: Applications are hosted in fast, cost-effective containers that are decoupled from the underlying operating systems. K8S-based CCE provides elastic scheduling for automatic scheduling and self-healing and elastic capacity scaling on demand, ensuring services run smoothly and resources are not wasted.
- DevOps-based reconstruction: Automatic pipelines and application environment orchestration enable resource triggering within seconds and environment deployment within hours. E2E O&M includes faster logging, comprehensive monitoring, efficient alarming, and comprehensive insights in to the entire system.
- Self-organizing teams: Small teams complete service analysis, development, testing, deployment, and O&M. They identify delivery bottlenecks, and quickly verify and optimize services using principles of lean production.

4.4.1.4 Migration Implementation

4.4.1.4.1 Migration Team Establishment

A large-scale migration project usually involves migration of a large number of applications in a short period of time and complex cross-application troubleshooting. The project management offices (PMOs) will take the lead, and organize and arrange personnel from both parties to carry out related work in an orderly and efficient manner based on the project objectives.

Huawei and the customer need to set up their own project teams based on Huawei Cloud experience in migration. The two teams will work together during the migration. The following figure shows the structure of the recommended migration team.



- Project manager (PM): The PM sets up a joint project PMO to manage project progress, identify and manage risks and other issues, and promote the project within the organization.
- Architecture and migration implementation team: The team designs solutions for cloud migration and cutover, manage and implement migration, and control technical risks during the implementation.
- Test team: The team tests solutions and performs function, performance, and joint commissioning tests.
- O&M team: The team creates, manages, and monitors cloud resources.
- Migration and R&D support team (Huawei Cloud): The team processes escalated technical issues.
- Development team (customer): The team is responsible for application development, deployment, and migration.

4.4.1.4.2 Migration Assurance

- Project kick-off meeting: A formal project kick-off meeting is initiated to specify the project scope, objectives, delivery period, and responsibilities.
- Project communication management: Regular communication is important between project members and project stakeholders. Communication exposes and helps address potential issues.
- Project progress management: The execution of project activities and key tasks is monitored to ensure that the project progresses as planned. If there is a deviation from the schedule, corrective measures will be taken and reported to project members and stakeholders in the form of weekly or daily reports.
- Issue and risk management: Project assumptions and risks are continuously monitored to quantify risks. The risk management plan may need to be updated at time and the implementation of risk mitigation measures needs to be ensured. Issues and risks are recorded and tracked, including details such as the issue owners and the times when the issues were identified. These details then need to be reported to project members and other stakeholders in the form of weekly or daily reports.
- Delivery matrix: A joint project delivery matrix covering Huawei Cloud and the customer will be developed.

4.4.1.4.3 Testing and Verification

Before services are cut over to the cloud, functional and performance tests will be performed to verify that the applications are stable in the cloud environment.

- Service function testing

  The Huawei project personnel use the cloud resource list to make sure all the required resources on Huawei Cloud are enabled. They initialize environment configurations, deploy applications, and migrate some data to perform joint commissioning tests. After the environment is deployed, customers can use their own test cases to test applications and confirm that services can run properly on the cloud.

- Performance testing

  Performance testing tests not just performance. There are also load testing, pressure testing, and stability testing. Performance testing means continuously raising the access pressure on the system (in a system test environment, the number of concurrent requests is constantly increasing for the test program) to obtain system performance KPIs and maximum load-bearing capacity.

  For performance testing, you can select Cloud Performance Test Service (CPTS) from Huawei Cloud, which supports CPTS and JMeter test projects, or use third-party pressure test tools.

  CPTS or JMeter (for simulating user traffic) and TCPCopy or GoReplay (for recording or copying real user traffic) are used to perform pressure testing on all-link applications on the cloud to check whether functions and performance meet requirements.

  During the test, test records and reports are generated based on the monitoring metrics of Huawei Cloud monitoring systems (such as Cloud Eye, AOM, APM) and the customer's monitoring systems as well as application logs.

  After the pressure test is complete, the dirty data generated during the test is deleted, and a full and incremental data migration is performed. After the migration is complete, at an appropriate time, service traffic is switched over to the application on the cloud.

4.4.1.4.4 Service Cutover and Rollout

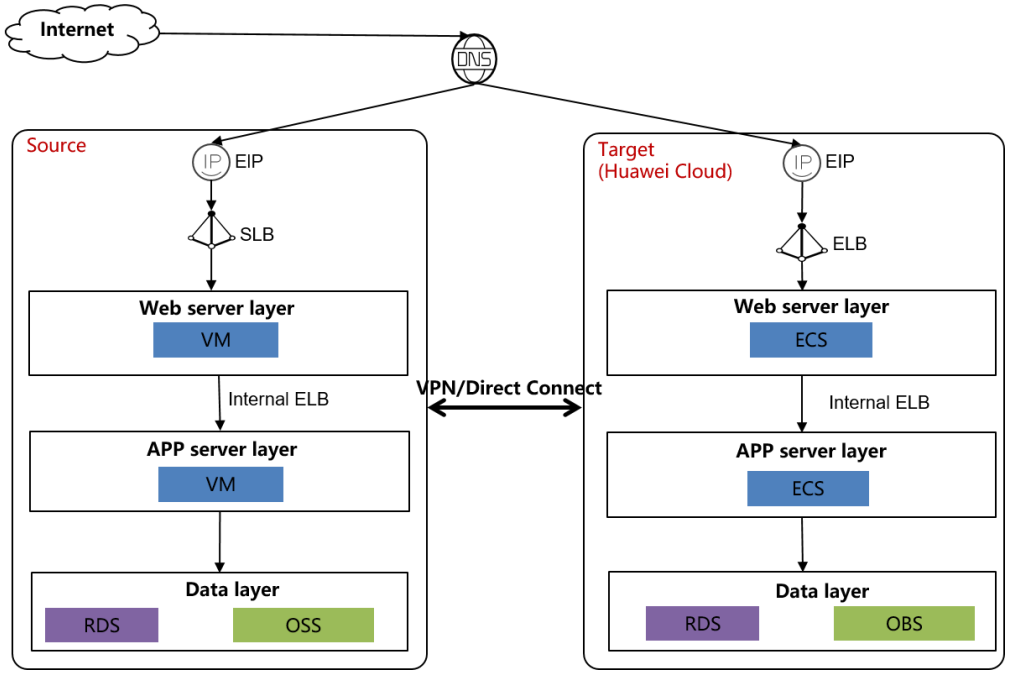
Service cutover and rollout are the most critical steps in a cloud migration. Any issues in this step may lead to major faults. To ensure smooth service cutover and rollout, a detailed pre-cutover checklist must be formulated.

After the service cutover is complete, service running and data are continuously monitored and observed until the services are stable on the cloud.

After the final data synchronization is complete and dirty data generated during testing is deleted from the target environment on Huawei Cloud, service cutovers are started in off-peak hours. Generally, service cutover can be implemented all at once, or layer by layer.

- All at once: For a simple or small-scale system, as long as a full-service verification is performed first, the cutover can be completed at a time. This kind of cutover is fast and has little impact on ongoing services.

  The one-off cutover process is as follows:



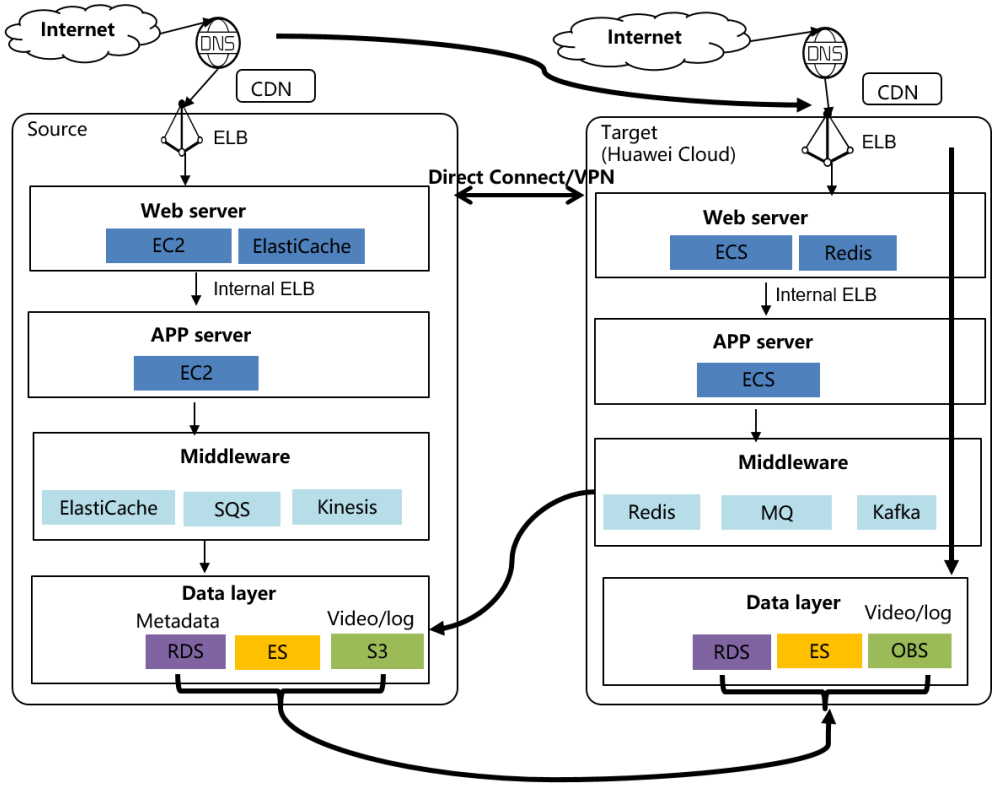
1. Stop the tests on the target and delete the test data.
2. Stop the services on the source.
3. Complete the incremental data synchronization.
4. (Optional) Configure reverse data synchronization to prepare for a rollback in case the cutover fails.
5. Modify the DNS configuration, switch the EIP, and switch traffic to the target.
6. Observe the service stability on the target.
7. Provide continuous assurance.

- Layer by layer: For complex services or large-scale systems, services can be decoupled and cut over layer by layer. If there are any problems, a rollback can be performed for an individual layer, which reduces the risk of impacting services. However, a hierarchical cutover like this requires multiple cutovers, which are time-consuming and involve a heavy workload.

  Layer-by-layer migration consists of two steps. First, the application layer is cut over to Huawei Cloud, but with the service data still being read from and written to on the source. After the application layer cutover is complete, a data migration or application dual-write is performed to synchronize service data from the source to the target on Huawei Cloud in real time. After the incremental data migration is complete, the second step, the data layer cutover, will be executed.

  Layer-by-layer migration involves cross-cloud database access. The network latency needs to be evaluated to see if it meets application requirements.

  The layer-by-layer cutover process is as follows:



Step 1 Application layer cutover

a Stop the tests on the target and delete the test data.

b Modify the configuration of the middleware layer on the target to point to the data layer on the source (through Direct Connect or VPN connections).

c Modify the DNS configuration, switch the EIP, and switch traffic to the target.

d Observe the service stability on the target.

Step 2 Database layer cutover

a Stop the services on the target.

b Complete the incremental data synchronization.

c (Optional) Configure reverse data synchronization to prepare for rollback upon migration failure.

d Modify the configuration of the middleware layer on the source to point to the data layer on the target.

e Start services on the target.

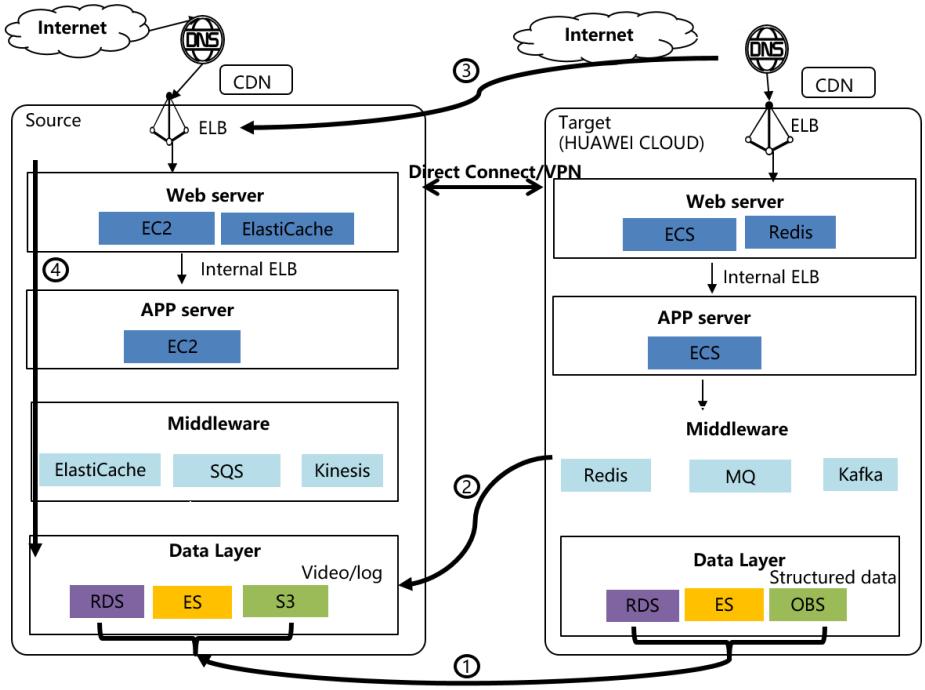
f Observe the service stability on the target.

g Provide continuous assurance.

----End

Cutover Risks and Service Rollback

The rollback solution in the all-at-once cutover scenario is simple. This section describes the rollback solution in the layer-by-layer cutover scenario.



Rollback solutions in the hierarchical cutover scenario are different if:

- The data layer has not been cut over:

  You can directly switch the DNS to switch traffic back to the source system.
- The data layer has been cut over:

  1) Roll back the data layer. The procedure is as follows:

     a Stop the services on the target.

     b Complete the reverse data synchronization.

     c Modify the configuration of the middleware layer on the target to point to the data layer created after the reverse synchronization on the source.

     d Start services on the target.

     e Observe the service stability.

  2) Roll back the application layer. The procedure is as follows:

     a Modify the configuration of the middleware layer on the source to point to the data layer created after the reverse synchronization on the source.

     b Modify the DNS configuration, switch the EIP, and switch traffic to the source.

     c Start services on the source.

     d Observe the service stability.

Assurance Solution for Seamless DNS Switchover

After a public domain name record is modified, it needs to be delivered to all DNS servers around the world. Due to certain restrictions inherent to international basic network infrastructure, the modification takes effect within 2 hours in the Chinese mainland, but takes 48 hours outside the Chinese

mainland. During this time, the domain name may be resolved to the original IP address, resulting in access exceptions.

If the domain name is resolved to the original IP address because the modification to the domain name resolution record has not taken effect, you can deploy Nginx or iptables on the source and use the Nginx HTTP proxy or the iptables NAT to forward traffic to Huawei Cloud, achieving seamless switchover of domain name records.

The following uses Nginx HTTP proxy as an example to describe the procedure for configuring traffic forwarding:

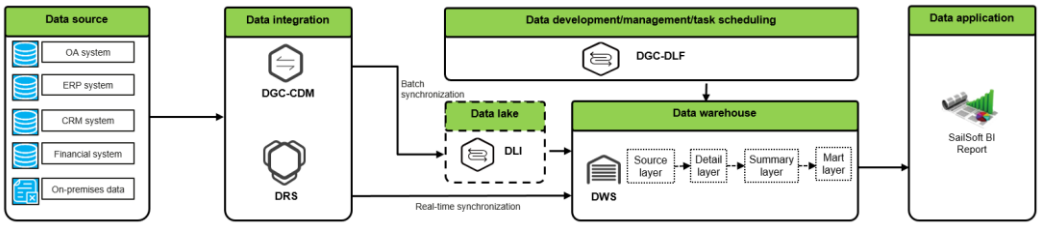
1. Before the switchover, deploy an Nginx reverse proxy server at the back of the load balancer on the source to forward traffic to Huawei Cloud ELB.
2. After the domain name is switched, bind the source domain name and IP address to the load balancer.
3. The load balancer forwards the traffic to the backend Nginx server. The backend Nginx server forward the traffic to the EIP of Huawei Cloud ELB load balancer over the public network.
4. (Optional) Deploy traffic monitoring software (for example, ntop) on the source Nginx server and view Nginx access logs. If no traffic is generated and access logs are not updated, DNS resolution has been switched over to Huawei Cloud.
5. Delete the load balancer and Nginx proxy server deployed on the source.

4.4.2 Data Migration to the Cloud

4.4.2.1 Typical Data Lake Scenarios

4.4.2.1.1 Data Warehouse and Report Analysis

This is a traditional data warehouse mode in which a real-time data warehouse is available and data mainly comes from databases. The data warehouse aggregates data from different business systems (such as the ERP, CRM, OA, and financial systems) and on-premises data centers, and processes, governs, and gains insight into the data by layer. The data warehouse eliminates data silos between departments, helps build a decision-making and analytics system, and provides data for business analysis and decision-making.



This solution mainly uses GaussDB(DWS). It has the following characteristics:

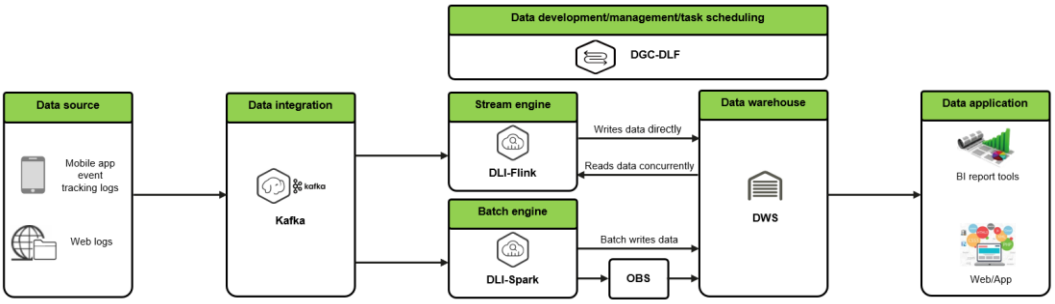
- One-stop big data BI platform: a comprehensive and efficient data collection, analysis, and BI platform that streamlines data from multiple business systems and provides full-stack technical capabilities.
- Cost-effective data analytics foundation: Huawei Cloud high-performance GaussDB(DWS) service and data synchronization solution can be used to analyze massive amounts of data quickly to unleash data value.

- Efficient development and simple management: DGC provides visualized, easy-to-use, flexible, and efficient development, management, and scheduling of data ETL tasks.
- Mature and reliable BI tools: Huawei Cloud cooperates with top BI vendors to provide mature, reliable, flexible, and efficient visualized BI tools, making operational analysis much easier and accelerating value monetization.

4.4.2.1.2 Integrated Solution for Streaming and Batch Processing and Query

This solution mainly applies to event tracking logs, which are commonly used in user behavior analysis and content and product recommendations. Event tracking logs are transferred to Flink through Kafka for real-time processing. The dimension tables required by Flink are stored in GaussDB(DWS) which allows millions of queries per second (QPS) for Flink, stores the analysis results from Flink, and enables upper-layer services to query the results.

Spark is a batch processing engine. The report data processed by Spark is stored in GaussDB(DWS) or OBS, and read using GaussDB(DWS) foreign tables.



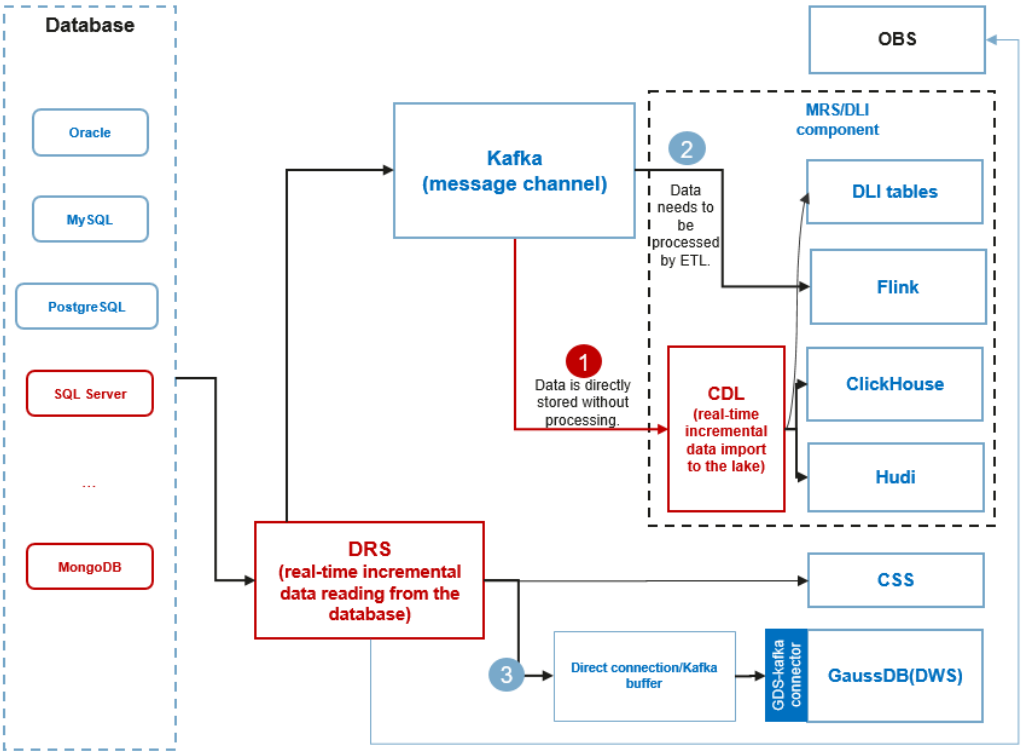
This solution mainly uses GaussDB(DWS). It has the following characteristics:

- DLI-Flink supports stream processing and can efficiently process event tracking logs.
- GaussDB(DWS) supports efficient indexing and clustering and millions of QPS for Flink dimension tables.
- GaussDB(DWS) can read OBS data so that the data processed by Spark can be queried using GaussDB(DWS) foreign tables, avoiding redundant data storage.
- GaussDB(DWS) enables query of fixed reports and self-service analysis, and provides high-performance, high-concurrency, and flexible report query capabilities.

4.4.2.1.3 Real-Time Import of Data from Databases to a Data Lake

Incremental data in databases can be imported to the data lake in real time.

DRS provides the CDC capability for TP/HTAP databases and can push real-time incremental data to a standard Kafka cluster for consumption in the data lake.



1. The built-in Kafka cluster of Change Data Lake (CDL) can directly connect to DRS and write data to Hudi, ClickHouse, and DLI in real time. By default, DRS can dump data to CSS and OBS (low priority) in real time for real-time search and AI training. CDL establishes a link to DLI. DLI and MRS share CDL's capability to import data to the data lake.

   CDL has embedded Spark scripts. You can import data directly to Hudi or ClickHouse without writing SQL scripts.

2. CDC imports data to the data lake through DRS and Kafka.

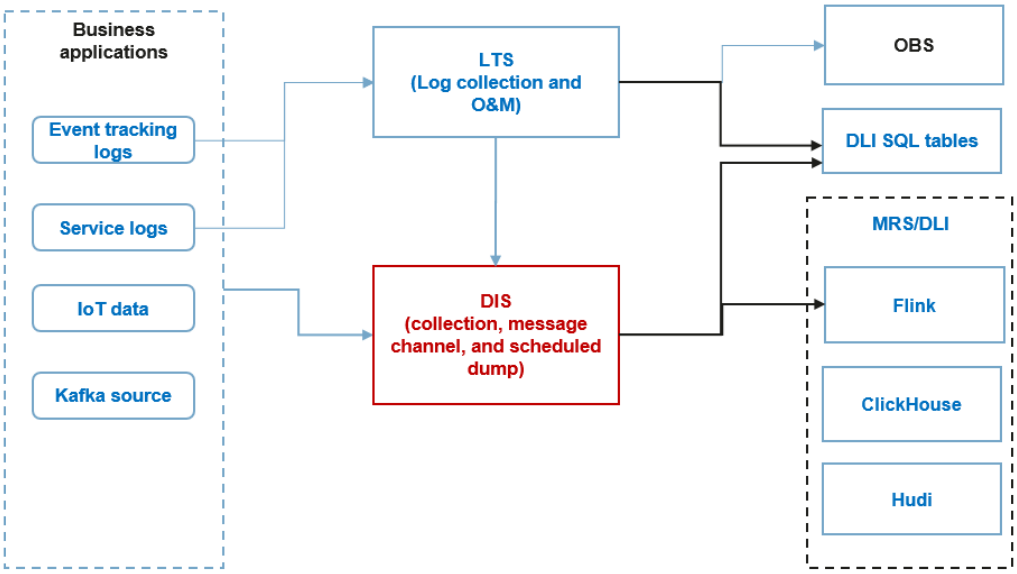
   DIS/DMS provides a standard Kafka message pipe for big data services such as MRS and DLI. DIS must support standard Kafka APIs.

3. The database CDC imports data to GaussDB(DWS) in real time through DRS and Kafka.
   - Direct import: This mode applies to scenarios where less than or equal to 3,000 lines of data need to be synchronized per second. DRS parses the real-time incremental data in the source database and directly writes it to GaussDB(DWS).
   - Import through a buffer: This mode applies to scenarios where more than 3,000 lines of data need to be synchronized per second. DRS obtains real-time incremental data from the source database and pushes the data to the backend Kafka message cluster. Then the built-in GDS-kafka Connector of GaussDB(DWS) writes the data to GaussDB(DWS) tables.

4.4.2.1.4 Real-Time Import of Messages and Logs to a Data Lake

DIS allows you to build Serverless message clusters and collect and dump data to OBS, DLI tables, and CloudTable.

LTS provides application operations and O&M capabilities, such as log collection, query, and analysis. In addition, LTS can dump log data to data lake components such as OBS, DIS, and DLI tables for further analysis.



Real-time import of streaming data such as messages to the data lake through DIS

- DIS provides Serverless message clusters with message pipelines for dumping messages to big data cloud services such as MRS and DLI.
- DIS can also collect and dump data to big data ecosystem services such as OBS, DLI, and CloudTable at scheduled time.

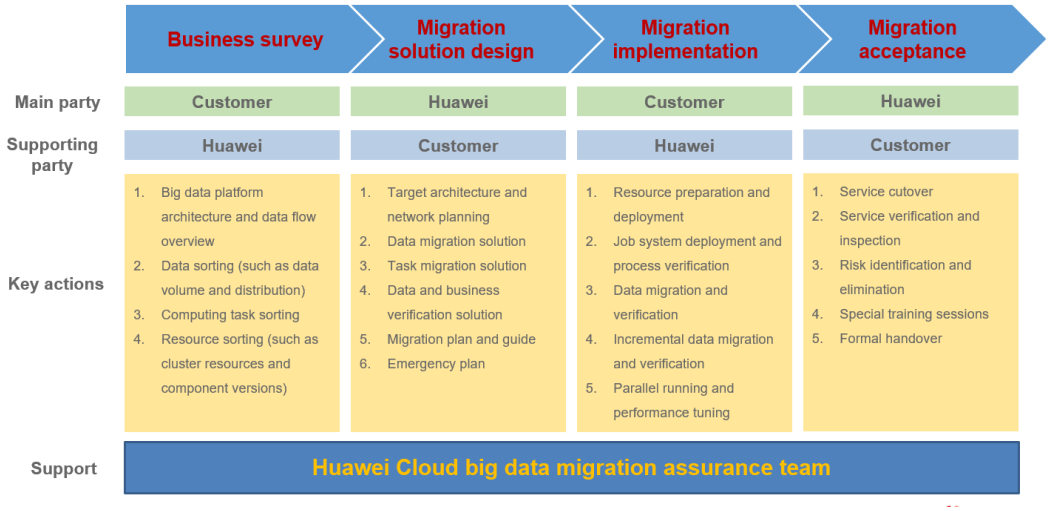
Import of streaming data such as application logs to the data lake through LTS

LTS provides log collection and O&M functions and can dump logs to DIS/DMS, OBS, and DLI for further analysis.

4.4.2.2 Big Data Migration

4.4.2.2.1 Process

Big data migration is a part of data migration and complies with the cloud migration theory and project management logic. In terms of project management, big data migration contains four phases: business survey, migration solution design, migration implementation, and migration assurance.



4.4.2.2.2 Migration Solution Design

In the **business survey phase**, the customer is the main party, and Huawei is the supporting party. They should work together to conduct the survey and determine the customer's business status, including:

- Customer's big data platform and services
- Physical deployment and data flows of the customer's big data platform
- Big data assets, including resources, data, and permission configurations

In the **migration solution design phase**, Huawei is the main party, and the customer is the supporting party. They should work together to design the migration solution and determine the migration content, including:

- The reconstruction and optimization needed to adapt to the destination platform and how to ensure a smooth migration
- Whether to perform migration in different phases and the objectives for each phase
- Cloud services involved in a platform migration
- Volume of the data to be migrated and the migration method to be used
- The customer's task scheduling system is the core for task migration
- The traffic switchover solution involves real-time data flows and basic service capabilities such as Nginx and ELB.
- How to ensure a successful rollback

In addition, required cloud resources should be planned and evaluated, including:

- Platform building resources
- Migration network bandwidth. This is closely related to the entire cloud migration solution and involves cost estimation. In addition, the networking on the cloud should also be planned, including how to enable network connectivity and protect network security.

In the **migration implementation phase**, the customer is the main party in resource preparation and deployment, job system deployment and process verification, data

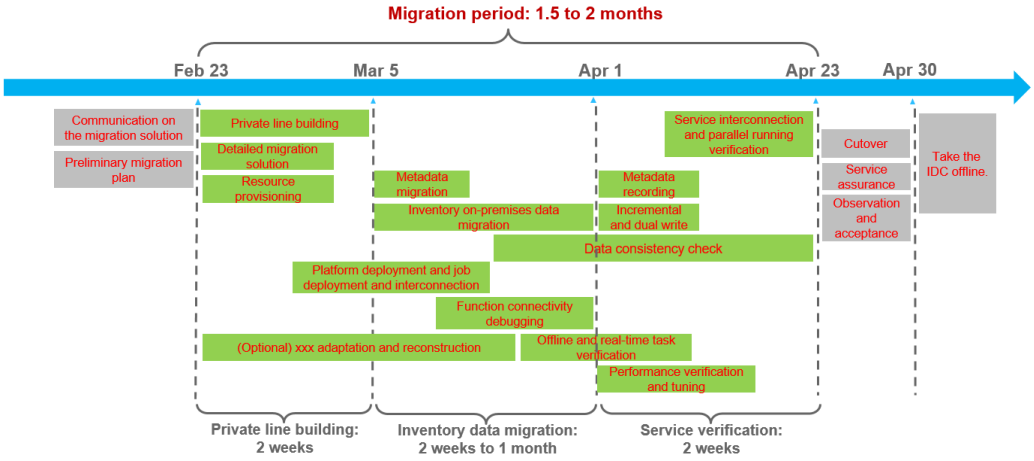
migration and verification, incremental data migration and verification, parallel running, and performance tuning.

The **migration acceptance phase** covers service cutover, verification, and inspection, risk identification and handling, special training and enablement, and formal handover.

4.4.2.2.3 Example Migration Plan and Period

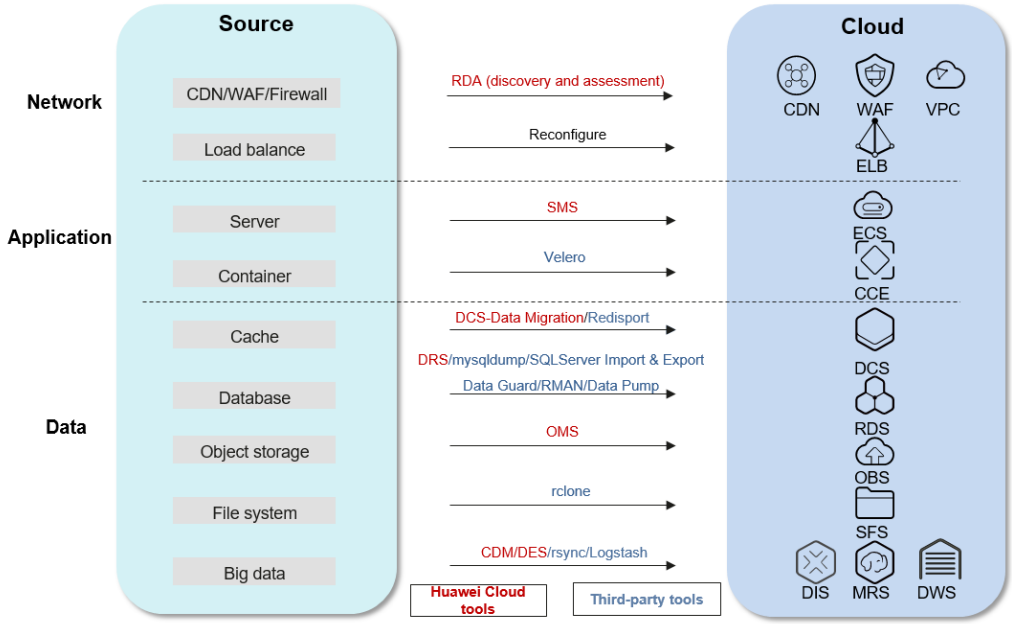
The migration plan and period vary depending on the data volume, number of tasks, components used, scheduling system, private line bandwidth, and time period when data can be migrated.

The following figure shows an example migration plan and period for a medium-sized customer (data volume: x PB; computing resources: x thousand cores).



4.4.2.2.4 Big Data Migration Tools

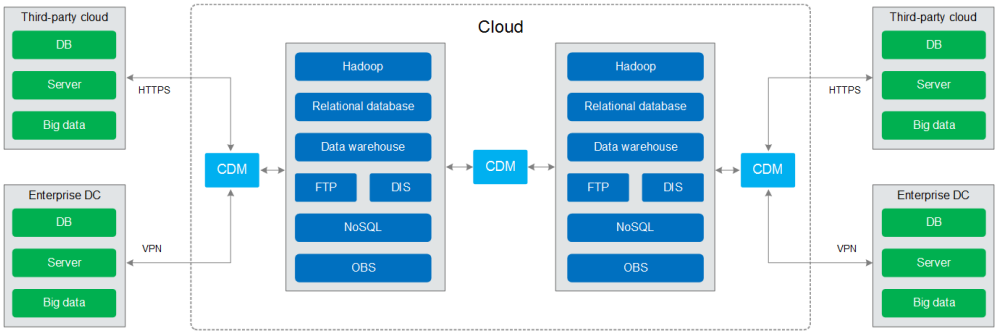
Huawei Cloud provides various migration tools for customers to choose based on the scenario, data source, data volume, and requirements for applications' responses to data. Customers can also choose open-source or third-party migration tools.



Huawei Cloud provides the following big data migration tools:

- CDM

  CDM is an efficient and easy-to-use batch data migration service. It provides easy-to-use migration capabilities and can integrate a wide range of data sources into the data lake, reducing the data migration and integration complexity and improving the efficiency. For details, visit https://www.huaweicloud.com/intl/en-us/product/cdm.html.



- Kafka MirrorMaker

  It forwards streaming data to MRS-Kafka on Huawei Cloud in real time. It is applicable to sequential messages as the consumer group has strict requirements on the message sequence.

  Generally, the MirrorMaker process needs to be started on the destination. The metadata of the source and destination Kafka clusters must be the same and manually configured.

- DRS

  DRS aims to enable database migration to the cloud without downtime. It supports migration between homogeneous, heterogeneous, distributed, and sharded databases. It also enables data integration and transmission from

databases to databases, data warehouses, and big data clusters in seconds, laying a solid foundation for enterprise digital transformation. For details, visit https://www.huaweicloud.com/intl/en-us/product/drs.html.

- OMS

  OMS helps you migrate data from the object storage on other clouds to OBS on Huawei Cloud.

  OMS applies to the following scenarios:

  - Object migration: When you migrate typical web applications to Huawei Cloud, OMS helps you easily migrate objects to OBS buckets on Huawei Cloud.
  - Cloud disaster recovery: OMS allows you to replicate your objects to OBS buckets on Huawei Cloud for disaster recovery and backup.
  - Object restoration: OMS allows you to use data backups from other cloud service providers to quickly restore data on Huawei Cloud.

  For details, visit https://www.huaweicloud.com/intl/en-us/product/oms.html.

4.5 Govern and Manage

4.5.1 Cost Management

Cost Center is a cost management service provided by Huawei Cloud free of charge. Cost Center can help collect information about Huawei Cloud costs and usage, explore and analyze Huawei Cloud cost usage, and monitor and track Huawei Cloud costs.

- Cost Analysis: presents your cost data in graphs and tables. You can view your Huawei Cloud costs and usage of different billing cycles and dive deeper by service type, region, linked account, billing mode, and cost tag.
- Cost and Usage Forecasting: forecasts your future costs and usage based on your historical costs and usage on Huawei Cloud.
- Budget Management: allows you to configure budgets and stay informed of how your costs and usage progress.
- Report Management: allows you to save the cost analyses as reports so that you can share with the various member accounts under your master account. On the **Reports** > **Analysis Reports** page, you can use some commonly used analysis reports preconfigured by Huawei Cloud. It also allows you to track your budgeting on a regular basis after you create budget reports on the **Reports** > **Budget Reports** page.
- Cost Anomaly Detection: monitors the costs of your pay-per-use resources to detect cost anomalies and reduce unnecessary expenditures.
- Billing Modes:
  - Yearly/Monthly: Cost Center analyzes the usage of your pay-per-use ECS, EVS, and RDS resources, and provides optimization options (changing from pay-per-use to yearly/monthly) to find opportunities to reduce costs.
  - Resource Packages: Cost Center allows you to view the analyses of your resource packages to determine whether they are properly used and whether costs can be reduced.

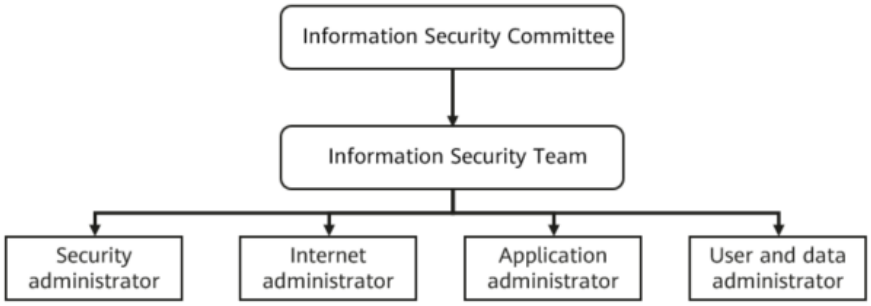
- Cost Tag: identifies your Huawei Cloud resources. You can use the cost tags to manage your resources and activate them to track your costs. Cost tags can be used for cost analysis and budget management.

4.5.2 Security Compliance and Governance

4.5.2.1 Security Management Organization

Set up a security management organization and define its working processes and the roles and responsibilities of its members.

A security management organization consists of an information security committee and an information security team. The information security team consists of a security administrator, a network administrator, an application administrator, a data administrator, and the users. The following figure shows a sample organization structure.



- **Responsibilities of the information security committee**
  - Implement information security guidelines and policies, and review and approve the information system security construction plan.
  - Make decisions on major issues related to information system security.
  - Study and review systems, standards, and related policies in the construction and management of information system security, and coordinate relevant departments to supervise the implementation of systems and policies.
  - Arrange and guide activities to promote information security.
- **Information security team responsibilities**
  - Implement government requirements and regulations on information security.
  - Manage information security for business departments.
  - Be responsible for the security management system of the information system, including system construction, technical support, and operational regulations.
  - Handle the daily management of, and regularly improve, the disaster recovery system. Formulate and regularly revise standards for evaluation of a disaster recovery system.
  - Monitor the status of the disaster recovery system, organize drills, audit and evaluate the system, and provide suggestions for improvement.

- If a major problem occurs in an information system, help related departments identify the cause and take immediate measures, such as initiating an emergency procedure.
- Raise awareness of information system security through training or publicity.
- Manage, inspect, and report information system security.
- Prevent system-wide malicious code and manage network security.
- Work with external security agencies. Obtain support from external security agencies in case of major security incidents.
- Develop and improve information system security plans.
- Manage, arrange, and supervise the security measures taken on information systems.
- Build, operate, and manage the O&M system and technical support platform.
- Arrange the formulation and implementation of the information system security O&M regulations.
- Arrange responses to faults and emergencies.
- Manage data utilization, query, and visualization.
- Manage operations in the DR center.
- Follow the instructions of the information security committee.
- Develop information system security regulations, and technical assurance and operation regulations.
- Regularly check information security status. If there are security incidents, report them in a timely manner, and work with related personnel in investigations and for troubleshooting.
- Conduct training to raise security awareness among internal employees.

- **Security administrator responsibilities**
  - Be responsible for the security-related affairs of information systems, and assist in supervising the implementation and modification of the security system.
  - Be responsible for the implementation of the security regulations.
  - Develop and improve information system security plans.
  - Make a list of security equipment, systems, or other related assets. The list should include the departments, categories, IDs, names, importance, and locations of the assets.
  - Standardize security equipment or system management. Record asset owners and users, security equipment or system identifiers, and actions taken on the assets, such as purchase, use, change, and retirement.
  - Develop O&M procedures and related documentation for security equipment or systems, so that relevant personnel can operate the system in a standardized manner, avoiding information security incidents caused by misoperations.
  - Manage O&M of security equipment or systems, including regular inspections, maintenance, troubleshooting, and change management.
  - Assist managers in security training. Develop annual security training plans. Raise security awareness through training and publicity.
  - Develop a security inspection plan. Arrange inspections at least once a quarter. The plan should describe inspector responsibilities, inspection frequency, inspection scope and content, problem rectification, and inspection

report. Conduct security inspections and report results. Collect, sort, and archive related materials.

- Report security incidents in time. Assist with troubleshooting and rectification.
- When a major problem occurs in the operation of the information system, assist the relevant departments in correctly determining the cause, and immediately take safety measures to initiate relevant processing procedures according to the instructions.
- Work with external security agencies. Obtain their support when there are incidents.
- Summarize security documents and forms semi-annually, and report the progress made in implementing security regulations.

- **Network administrator responsibilities**
  - Be responsible for network-related O&M, review, and changes, and archive related documents and data.
  - Manage network security development. Check whether the owners are fulfilling their responsibilities, catch up with the plan, and meet task requirements and quality standards.
  - Standardize the network management process. Record asset owners and users, security equipment or system identifiers, and actions taken on the assets, such as purchase, use, change, and retirement.
  - Adopt technical and management measures to enhance the security control over information systems. Continuously improve network security and stability, and logically isolate the external network of the information system from the Internet. Review and record network access. Arrange training about information system security.
  - Develop O&M procedures and related documentation for security equipment or systems, so that relevant personnel can operate the system in a standardized manner, avoiding information security incidents caused by misoperations.
  - Assist security administrators in deploying network security products. Manage the O&M of security equipment or systems, including regular inspections, maintenance, troubleshooting, and change management.
  - Implement multi-level approval regulations for network system changes, important operations, and access. Review routine applications. Report major changes to the Information Security Department.
  - Arrange emergency responses to network faults or incidents.
- **Application administrator responsibilities**
  - Manage O&M, application review, and changes in the application system. Archive the documents and data related to the application system.
  - Manage software development. Check whether the owners fulfill their responsibilities, catch up with the plan, and meet task requirements and quality standards. Set information security objectives during project establishment and approval to enhance the security of the system design and development process. Enhance code security. For outsourced software development, sign a confidentiality agreement with the service provider. During delivery acceptance, invite a third-party security organization to evaluate product security.
  - Standardize the management process of the information system. Record asset owners and users, security equipment or system identifiers, and actions taken on the assets, such as purchase, use, change, and retirement.

  - Develop O&M procedures and related documentation for important applications, so that relevant personnel can operate the system in a standardized manner, avoiding information security incidents caused by misoperations.
  - Standardize the asset management process of the application system. Record asset owners and users, security equipment or system identifiers, and actions taken on the assets, such as purchase, use, change, and retirement.
  - Manage routine security measures regarding password management, authorization, and approvals in the application system, so that employees can better comply with information security regulations.
  - Implement multi-level approval regulations for application system changes, important operations, and access. Review routine applications. Report major changes to the Information Security Department.
  - Arrange emergency responses to network faults or incidents.
- **User and data administrator responsibilities**
  - Provide technical support of large-scale application systems. Monitor and analyze daily operational data. Identify opportunities and pain points in business operations.
  - Conduct in-depth analysis and mining of customer demand models and data. Generate metrics-based warnings.
  - Manage and monitor accounts.
  - Ensure the database is running properly.
  - Identify and eliminate background problems and risks in a timely manner.
  - Fine-tune the system.
  - Plan and implement backup strategies.

4.5.2.2 Personnel Security Management

This section describes how to enhance security through personnel management.

- **Recruitment**
  - Recruited personnel need a basic understanding about the technologies involved and of security management. Arrange a test on this knowledge before formal appointment.
  - Check the identities, backgrounds, and qualifications of recruited personnel. Archive related materials.
  - Test recruited personnel's technical skills.
  - Introduce the roles and responsibilities to recruited personnel and arrange job training.
  - Sign a confidentiality agreement with recruited personnel to forbid information leaks.
- **Staff position change or resignation**
  - When an employee leaves their position, revoke their system permissions.
  - When an employee resigns, they need to submit any passwords they have been using for company devices. After their resignation, the administrator shall change the passwords as soon as possible.
  - Resigning employees shall go follow a strict resignation procedure in the HR department and sign a non-disclosure agreement.
- **Staff appraisal**

- Regularly conduct training focused on improving security skills and awareness. Managers, regular employees, and third-party managers and users all participate if necessary. Ensure they can identify information security threats and risks, and can comply with security policies.
- Conduct strict, comprehensive security audits for personnel in key positions. For system administrators, review their identities and how they are fulfilling job responsibilities. Review their system permissions and job responsibilities, and their compliance with non-disclosure regulations. For other key personnel, check their identities and evaluate their fulfillment of job responsibilities.
- Punish those who violate the security policies or regulations. For minor violations, a warning and a requirement to write a formal apology are appropriate. For major violations, relevant departments need to investigate the person's legal responsibilities. If the person responsible for a minor violation comes from a third party, ask them to correct their mistakes and inform their company.

- **Security awareness training**
  - Regularly conduct security training. Ask employees to learn about network security and ensure they understand relevant policies and regulations. Ensure employees understand that they will be held responsible even if they meant no harm and even if they violated security regulations unintentionally.
  - Accurately define the security responsibilities of every position and the punishments for different violations. Arrange meetings to explain important or complex regulations if necessary.
  - Arrange publicity activities to help employees learn about network security community operations and common violations, for example, by making short videos.
  - Develop security training plans. Conduct training on basic information security knowledge and job procedures.
  - Record security training in detail and archive the records.
  - Ask employees to sign the Information Security Commitment Letter and promise to abide by the company's network security policies and regulations.
- **Security capability training**
  - Establish a network security training system based on the industry best practices. Arrange security capability training during different stages, for example, during new employee orientation, during their regular work week, or before promotions. Ensure employees are capable of delivering secure products, solutions and services.
  - Basic network security training: Develop training plans for different roles and positions. For example, new employees must pass on-the-job training and exams on network security and privacy protection before they become regular employees. On-the-job employees need to take courses as required by their positions. Managers must participate in network security training and seminars.
  - Precise training: Identify typical security problems in the product development process and those responsible for the problems. Recommend security training programs (including cases, training courses, exercises, etc.) to them.
  - Drills: Adopt industry best practices, develop a platform for network security drills, and arrange confrontational role playing exercises. Improve employees' security capabilities through practices.

- Incorporate network security requirements in the acceptable criteria for jobs and promotion.

4.5.2.3 Security PMI

Develop Preventive Maintenance Inspection (PMI) standards. Define the monitoring and alarm handling process for services or systems after they are migrated to the cloud. Describe the key monitoring points, remediation actions, and issue escalation process to minimize the impact of security incidents on services.

- **Network security PMI**
  - Check for ports accessible from the Internet. For example, check the configurations of ELB and NAT Gateway. Ensure there are no unsafe ports accessible from the Internet, and no resources directly bound to an EIP. Unsafe ports include internal management ports (for example, 22 and 3389) and high-risk ports (135, 139, and 445 for TCP; 137 and 138 for UDP; or other ports that can be exploited by EternalBlue or other ransomware).
  - Check intranet connections. Ensure the network ACL, security group, or other network security configurations are clearly defined, only necessary ports are open, and arbitrary access to ports is denied.
- **System security PMI**
  - Vulnerability management: Scan for vulnerabilities, fix them in a timely manner, and track the vulnerabilities until they are fixed. Our vulnerability management functions can help you detect vulnerabilities in Linux, Windows, and Web-CMS.
  - Baseline checks: Detect unsafe settings based on security baselines, including weak password policies, common weak passwords, and unsafe configuration items.
  - Asset management: Inventory all the assets on your servers, including accounts, open ports, processes, web directories, software, and auto-started items.
  - Intrusion detection: Use HSS to check for malicious programs, suspicious processes, web shells, abnormal shells, unsafe accounts, and rootkits. Work with related personnel to eliminate risks in a timely manner.
- **Application security PMI**
  - Defense against common attacks: Ensure your business systems have configured web firewall policies to monitor and block common web application attacks, such as SQL injections, XSS, file inclusion, directory traversal, sensitive file access, command/code injection, Trojans, and web shells.
  - Service access control: Meet with business personnel and determine the scope of allowed access sources. Minimize this scope by configuring WAF, VPN, and CBH; blocking IP addresses in certain regions; and creating an IP address blacklist and whitelist.
  - Common vulnerability monitoring: Scan for and fix vulnerabilities in a timely manner.
  - Web page tamper-proofing: Configure static Web Tamper Protection (WTP) in HSS to protect website root directories, preventing attackers from uploading scripts and launching reverse shells.
  - Penetration testing: Manually check for and fix logical vulnerabilities in a timely manner.

  - Account management: Work with business experts to review all the application accounts and revoke unnecessary permissions, especially those for authentication, authorization, account, and audit (4A) systems. Enable multi-factor authentication (MFA) and delete zombie accounts.
- **Data security PMI**
  - Data access permission control: Review all the database accounts and permissions, delete test accounts and other unnecessary accounts, and retain only the least required permissions for accounts.
  - Database protection: Deploy database audit and protection equipment to filter, monitor, and audit database operations in real time.
  - O&M plane checks: Prevent insiders from using O&M accounts to access the internal network. Ensure CBH, VPN, and other systems used for O&M access require two-factor authentication (2FA) and strong passwords. This can prevent account cracking, which may result in intrusions into a large number of servers.
  - Account security checks: Ensure that the length and complexity of account passwords meet security requirements. Manage permissions at a fine granularity. Allow administrators to access only the resources related to their jobs.
  - Backups: Configure backup mechanisms for important business systems. Use data backups, server backups, and snapshots to restore services if there is an emergency.

4.5.2.4 Account Security Management

This section describes common account types and management measures.

- **Account types**
  - CBH accounts
  - OS account
  - Database accounts
  - Cloud management platform account
  - Business system account
  - Tenant account provided by the cloud provider
  - Other business-related accounts
- **Roles and responsibilities**

Role	Description
Account applicant (Maintenance, development, or audit engineer)	The one who is applying for an account or permission (for themselves or others)
Account user (Maintenance, development, or audit engineer)	User of an account or permission
Account administrator (System	Responsible for routine account and permissions management, such as reviewing account creation and deregistration requests, managing passwords, and

Role	Description
administrator)	changing the owner of an account.
Account approver (O&M manager)	Responsible for reviewing and approving account permissions applications. Their work includes: • Checking whether the applications for accounts and permissions are authentic, necessary, and appropriate. • Regularly reviewing accounts and permissions. Revoke unnecessary accounts or permissions in a timely manner. • Ensuring the overall security of accounts and permissions.

- **Constraints**
  - Employees can apply for permissions to access the production environment only when necessary.
  - For employees who have worked in your company for one month or longer, but have not completed their probationary periods, their supervisors shall assume responsibilities when they apply for system permissions.
  - Employees who have violated O&M regulations shall not be granted long-term change permissions. Only one-time change permissions can be granted.
- **Validity period**

  The validity period of a permission depends on the role. If a permission expires, you need to apply for it again.

Permission Validity Period	Usage
1 year	Long-term network O&M. Common applicants include: • O&M personnel • O&M manager • O&M platform developers responsible for long-term O&M tool development, system interconnection, and release
3 months	Project development. Permissions shall be revoked after the project delivery.
1 week	Short-term project delivery, acceptance, and issue tracking.
1 day	Short-term access for incident/problem handling, issue escalation, and security check.

- **Application process**

1. The applicant fills out the application form, sends it to the account approver, and CC's it to the account administrator.
2. The account approver checks whether the application is necessary and appropriate, and then approves it or provides review comments.
3. If the application is approved, the account administrator creates an account and send it to the account user.
4. The account user changes the initial password before using the account.

NOTE

To change the permissions of an account, you do not need the account approver to approve the account again. Submit an online request or send an email to obtain approval for the change. Then, send the proof of approval to the account administrator. The account administrator will assign permissions accordingly.

- **Deregistration process**

  1. The account applicant sends an email or submits an online application for deregistration and explains the reason (such as resignation, job transfer, or job responsibility change).
  2. The account administrator deregisters the account and revokes its permissions.
  3. The account administrator records the deregistration results and sends them to the account applicant, owner, and user.

  The account administrator regularly checks account validity periods, reports expired accounts, and deletes accounts as needed after conferring with related business managers and users.

NOTE

- Accounts that are valid for only one week or less will automatically be deregistered upon expiry.
- For employees who have submitted resignation applications, ensure their permissions are revoked before they leave the company.

- **Precautions**

  - Configure appropriate password complexity rules. A password must be at least 8 characters long and must contain digits, letters, and special characters. Do not allow common weak passwords or easy-to-guess passwords.
  - Do not allow users to lend their accounts to others. A user is responsible for all the operations performed through their account.
  - For accounts that have not been logged in for 30 days, the account administrator can deregister them after notifying the account users by email.
  - If the account administrator finds that the password of an account was leaked or compromised, they must immediately report the problem to the account owner, reset the password and notify the account user, and send a detailed report about this incident to the account owner.
  - Passwords need to be changed regularly. Password cannot be used for longer than 90 consecutive days.

4.5.3 Cloud-based O&M

4.5.3.1 Trends and Challenges

Cloud-based O&M differs greatly from traditional O&M.

- **Model**
  - Traditional O&M: Enterprises need to build most capabilities on their own. They directly manage computing, network, and storage, and other resources. They take all the responsibility and full control of their resources. However, capability construction is slow and difficult.
  - Cloud-based O&M: Cloud services provide standard basic capabilities. Based on these services, enterprises build required architectures and O&M capabilities. They use software interfaces or open APIs to manage abstract resources. Enterprises share responsibility with cloud vendors and do not need full awareness of underlying resources. However, capability construction is fast and flexible.
- **Scope**
  - Traditional O&M focuses on maintaining IDC equipment rooms. O&M environments are rarely changed and therefore familiar to enterprises.
  - Cloud-based O&M varies by IT development and industry requirement. Environments include offline IDC, private cloud, public cloud, or even multi-cloud. Increasingly complex O&M environments require higher O&M control and skills.
- **Overall capability**

  The cloud provides many more products and functions than traditional IDCs. Although O&M personnel are not tasked with underlying maintenance, they need to be familiar with the various cloud products for better selection and usage.
- **Security risks**
  - Cloud resources are mostly logically isolated, which is riskier than physical isolation. Therefore, proper security planning is crucial. For example, cloud storage must be encrypted to protect data.
  - The cloud architecture solution is flexible. Most management is based on open APIs, and resource access is by running single commands. Therefore, security needs to be controlled during external access design. In addition, strict audit is recommended for risky commands.
- **Troubleshooting**
  - Cloud service faults are more difficult to locate. Once microservices are deployed, the relationship between services becomes more complex. Therefore, cloud-based O&M is key to locating faults.
  - Some faults need to be located with the cloud service support team. Therefore, in-depth communication with the cloud service support team is needed.

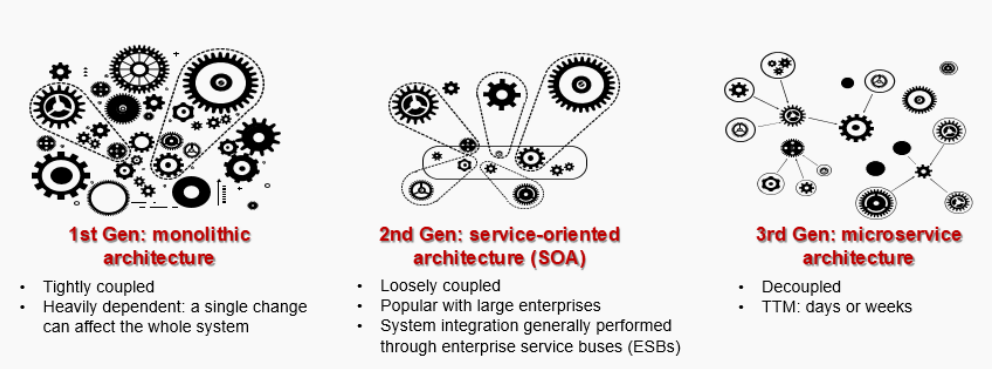
- **Skillset**

Cloud-based O&M requires a different skillset than traditional IT operations. Organizations need to ensure that they have the right talent in place to manage and maintain their cloud-based infrastructure, applications, and services.

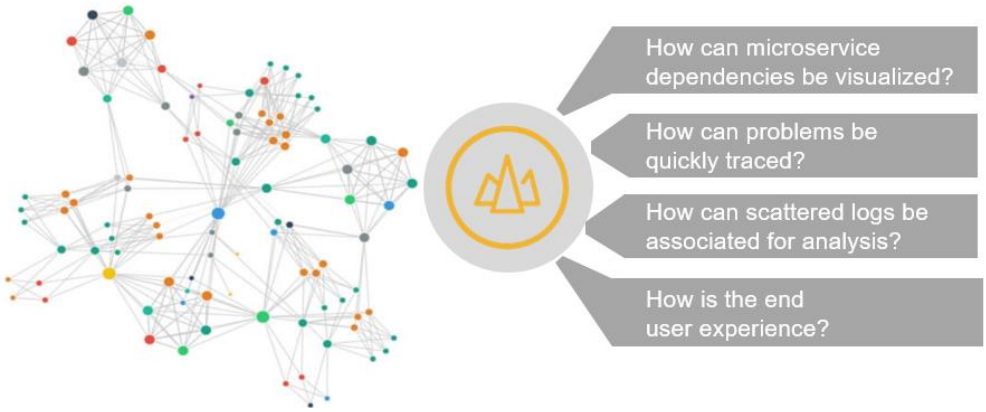
4.5.3.2 Multi-Dimensional O&M

The O&M system is designed to provide high-quality IT services. It monitors underlying resources, applications, user experience, as well as overall system running. It enables personnel to quickly respond to issues, ensuring stable service running. With the changes in service formats, architectures, resource models, and call relationships, the cloud-based O&M system must provide refined and large-scale management. There are several typical O&M requirements.

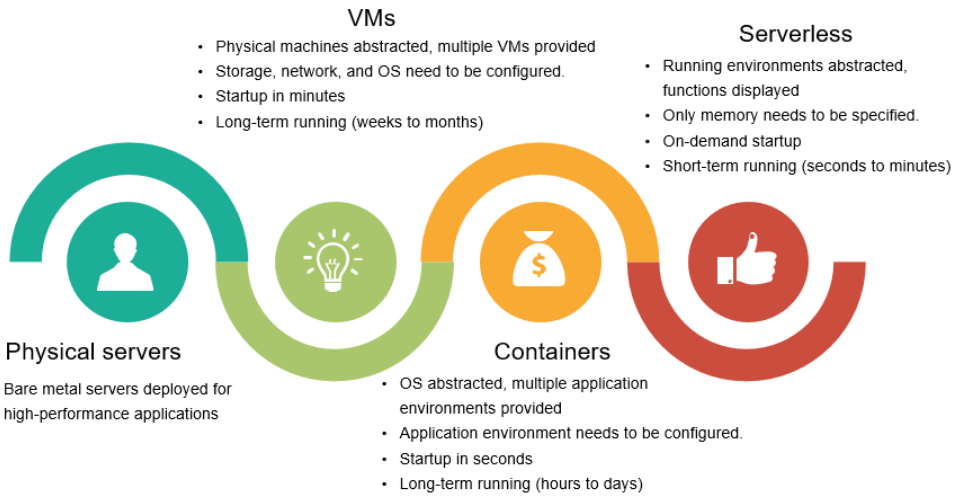
1. As service architecture evolves from monolithic and service-oriented to microservice-based, the number of services increases exponentially. Finer-granularity services are faster to iterate and more refined to manage. However, many microservices also pose great challenges to O&M in terms of management scale and timeliness.



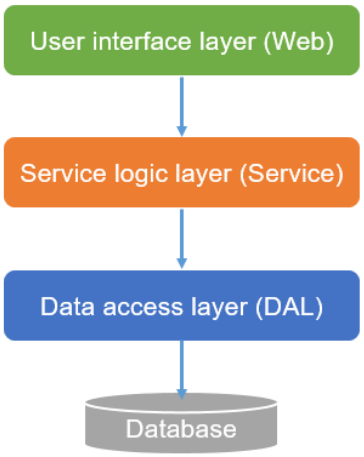
2. In a large cloud-based distributed application system, service call relationships are complex. The call relationships, call quality, and latency in each phase need to be visualized so that problems can be detected and resolved in time.



3. After applications undergo microservice reconstruction, underlying resources are no longer hosted in physical machines and VMs; instead, they are hosted in containers or even run in serverless mode. Rapidly increasing service refinement and scale also mean exponential resource growth. When different types of resources are combined in services, O&M is even more complex. Personnel need refined and large-scale management capabilities to monitor resource usage in time.

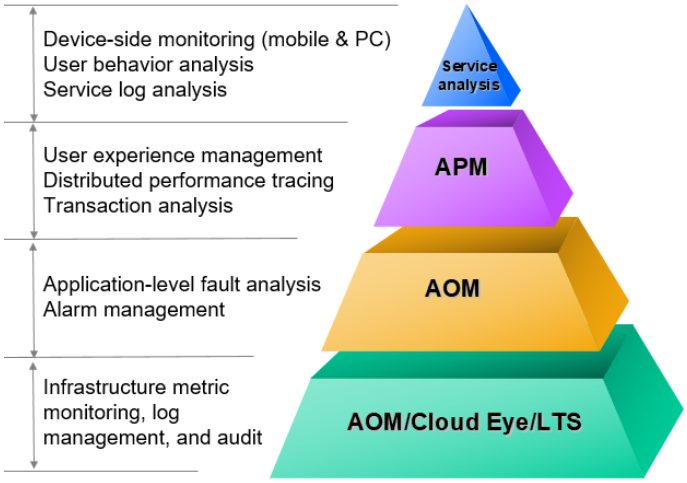


4. Currently, applications use layered architectures, typically consisting of the web, service logic, data access, and database layers. To obtain required data, a request may go through multiple layers. O&M personnel need to be able to drill down links to monitor inter-layer access quality and serial/parallel relationships during troubleshooting.



5. Complex cloud resources and services are complicated by many accounts and operators, as well as frequent service changes. To ensure system stability, O&M personnel must have strong resource management, control, and audit capabilities to quickly detect abnormal operations and locate causes.

   Huawei Cloud has built a multi-dimensional O&M system for cloud applications. It integrates AOM, APM, log collection, and monitoring. It monitors VMs, storage devices, networks, databases, and applications in real time, and uses technologies such as application and resource alarm correlation, log analysis, intelligent threshold, distributed tracing, and mobile app exception analysis to quickly diagnose and rectify faults within minutes, ensuring long-term and stable running of cloud applications.



Multi-dimensional O&M monitors the infrastructure, application, and user experience layers, and also provides logs and auditing.

- **Cloud Eye**

  Cloud Eye is a multi-dimensional resource monitoring service. Monitor resources, set alarm rules, and improve resource utilization and performance.

- **AOM**

  AOM is a one-stop cloud operations management platform for problem management, monitoring, security, and performance. Streamline cloud operational processes and effectively manage cloud hardware, software, services, and networks.

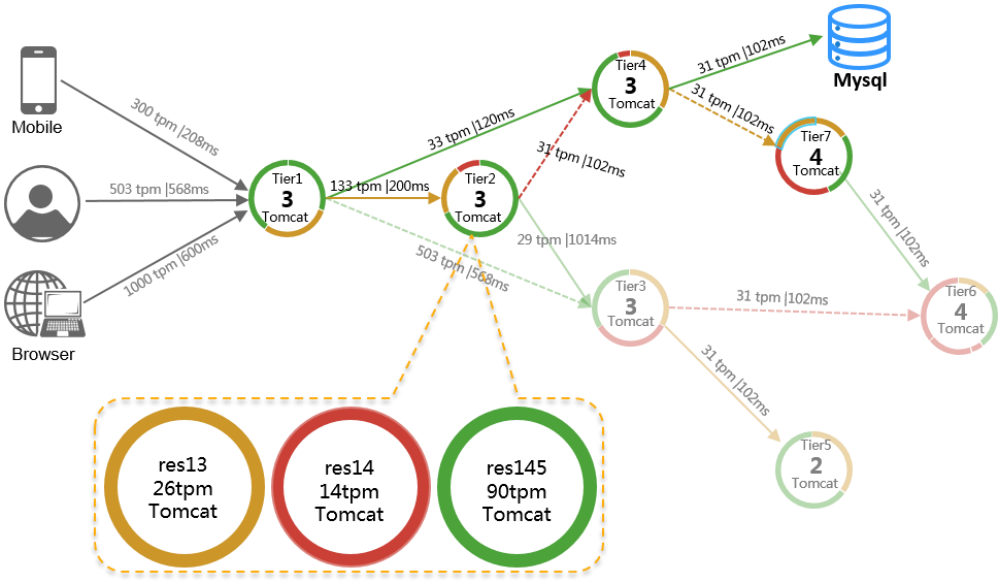
- **APM**

  APM consists of application and frontend monitoring. Manage the performance of your distributed applications, container environments, browsers, applets, and apps. With full-stack performance monitoring and E2E full-link tracing and diagnosis, application management is easy and efficient.

  APM has several features.

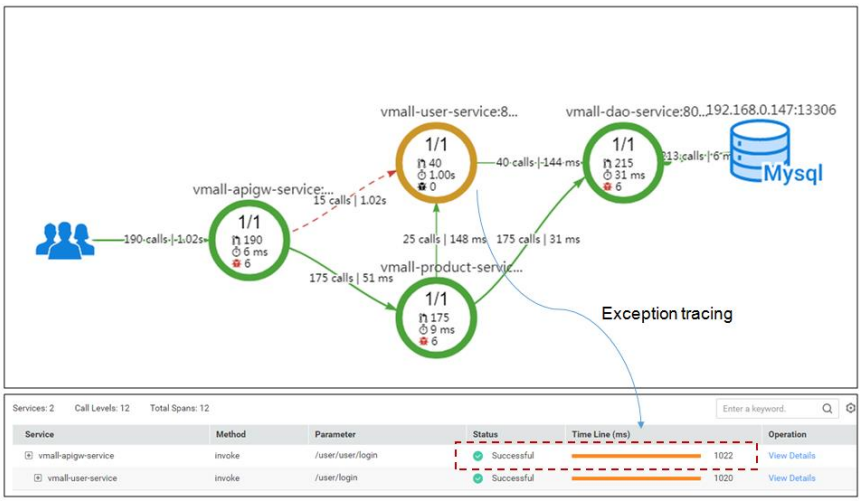
  - **Full-link Topology:** Clear Call Relationships, Easy Error Identification, and Convenient Resource Drill-down

    The full-link topology displays the call relationships and dependencies between applications, including application status, latency, errors, and loads. View multiple open-source components, such as databases, caches, message middleware, and NoSQL, filter information by time, service, transaction, top SQL statistics, or other metrics, and drill down resources to locate faults.



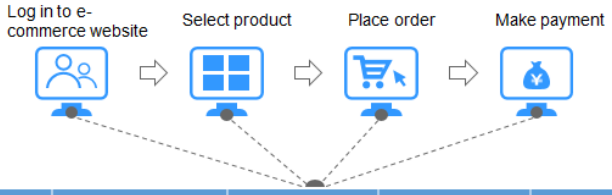
- **Tracing:** Bottleneck Identification and Fault Locating in Minutes

  APM traces and records service calls, and displays the distributed system's request execution tracks and statuses. When a service method is called, the caller, detailed stack, and parameters of the method are automatically captured for fast fault locating.



- **Transaction Analysis:** Displaying Key Metrics

  APM analyzes service flows on the service side in real time and displays key metrics such as throughput, error rate, and latency. Abnormal transactions trigger alarm reporting. Application Performance Index (Apdex) is used to evaluate user satisfaction with applications. You can view and trace the topology of any transaction with poor user experience to locate the cause. An e-commerce application is used as an example here.



Transaction	Throughput (Calls per Minute)	Average Latency	Health	Error Rate	Performance Analysis
Login	52	323 ms	Normal	0%	Tracing analysis
Product search	234	721 ms	Slow	1%	Tracing analysis
Procurement	3	1.32s	Normal	0%	Tracing analysis
Payment	1	2.1s	Abnormal	100%	Tracing analysis

- **LTS**

  LTS collects log data from hosts and cloud services. By processing massive amounts of logs efficiently, securely, and in real time, LTS provides useful insights for you to optimize the availability and performance of your services and applications. It also helps you make quick decisions, better manage and maintain devices, and analyze service trends.

  The following table shows open-source O&M products and their corresponding Huawei Cloud products.

Category	Common Tool	Huawei Cloud Service	Function
Infrastructure O&M	Zabbix	Cloud Eye	Monitors and generates alarms for infrastructure and instances.
	ELK	LTS	Collects and analyzes logs, generates alarms, and archives logs.
Application-layer O&M	Prometheus+Grafana	AOM	Monitors applications and cloud resources in real time, analyzes application health, and displays visual data, so that O&M personnel can quickly detect faults and learn statuses of applications, resources, and services.
Application performance	Zipkin/Pinpoint/ SkyWalking	APM	Monitors and manages

Category	Common Tool	Huawei Cloud Service	Function
management			application performance and faults in real time, and analyzes the performance of distributed applications, helping O&M personnel quickly locate faults and performance bottlenecks.
Service monitoring	Service system O&M functions and service logs	LTS	Monitors and analyzes behavior and logs on the service side.

4.5.3.3 Backup and Restoration

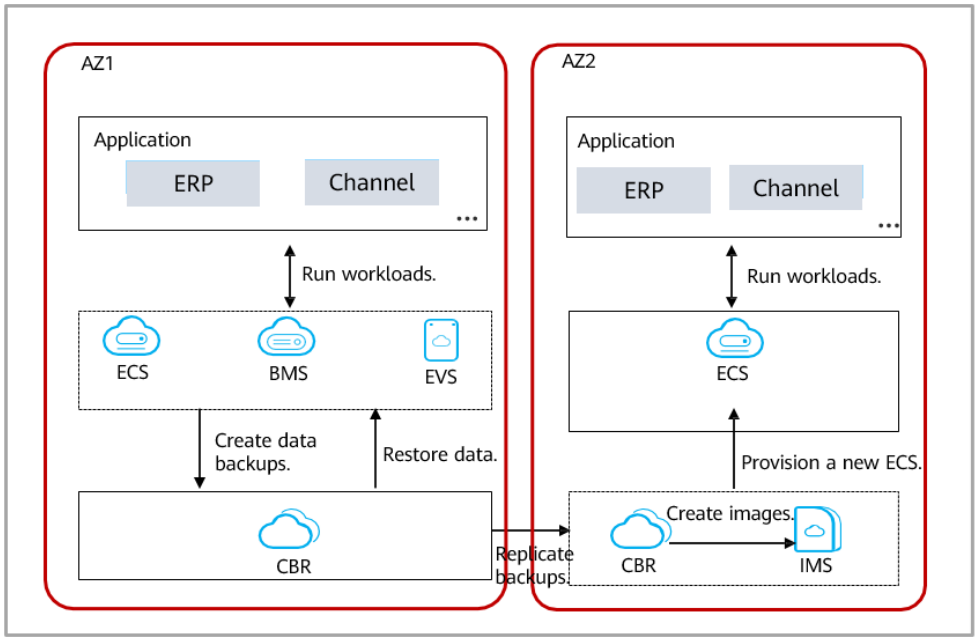
Data security is a critical concern for most enterprises today. When migrating services to the cloud, cloud data security is a major concern. Huawei Cloud provides a comprehensive backup and restoration system to guarantee data security.

- **Overview**

Object	Backup Policy
ECS	Cloud Backup and Recovery (CBR) uses snapshots to back up entire ECSs, or you can configure it to back up the system disk and data disks separately if needed. You can configure a backup policy with whatever execution time and backup frequency you need.
Database	• After the **automated backup policy** is enabled, a full backup is automatically triggered. Then, a full backup will be performed based on the time window and backup cycle specified in the backup policy. • When an instance status is **backed up**, data is copied from the instance, compressed, and uploaded to an OBS bucket, and stored in OBS for as long as you specified in the backup policy. Amount of time required to complete the backup depends on the amount of data to be backed up. • After the automated backup policy is enabled, an **incremental backup** is automatically performed every 5 minutes to ensure data reliability.
Key service data (such as configuration files and key documents)	You can set the backup cycle and retention period for key service data based on the **file importance** and **change frequency**. You are advised to back up service data to OBS buckets.

- **Cloud Server Backup and Restoration**

  The following figure shows the backup capabilities provided by Huawei Cloud CBR for ECSs, BMSs, and EVS disks.

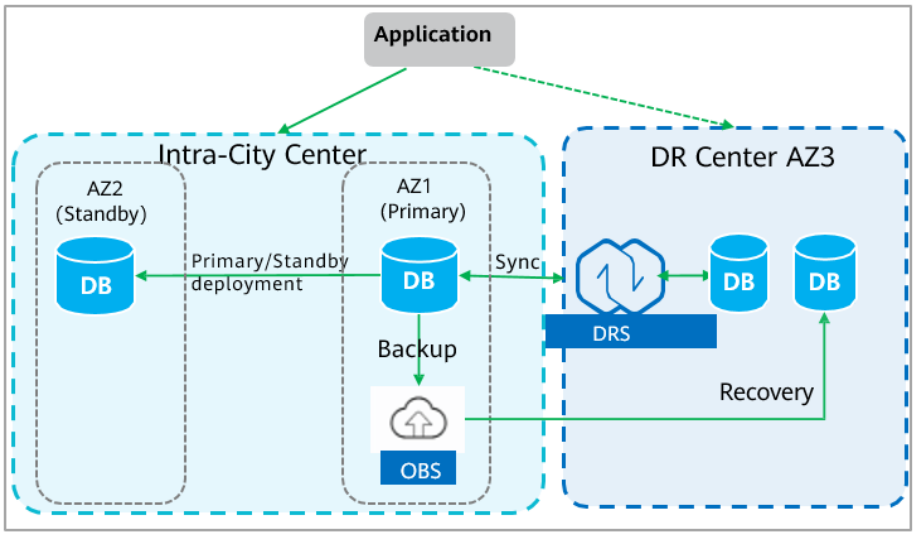


Category	Function	Feature
Backup	Cloud server backup	Backs up ECSs including their EVS disks using the snapshot technology.
	BMS backup	Backs up BMSs including their EVS disks using the snapshot technology. (The BMSs must use EVS disks.)
	Cloud disk backup	Backs up cloud disks.
Recovery	Cloud disk restoration using backups	Restores backup data to specified points in time.
	Cloud server restoration using backups	Restores servers to specified points in time.
	Cloud disk creation using backups	Creates disks from backups.
	Image creation using backups	Uses backups to create images and then provision ECSs.

- **Database Backup and Restoration**

  For service scenarios that demand HA, OBS (cold standby) or DRS (hot standby) are deployed in different AZs in the same city to establish a remote DR

center and ensure geo-redundant data security. If both AZ1 and AZ2 fails, the DR center can still ensure data security.



If a database or table is maliciously or mistakenly deleted, the standby database is also deleted and cannot be restored. Backups can protect you from this sort of malicious or accidental operations.

- Automated backup: Huawei Cloud database service creates automated backups for the DB instance during the backup period and saves the backups for a length of time determined by the configured backup retention period.
- Manual backup: Manual backups of DB instances are saved based on the configured backup retention period. You can restore data to any point in time within the backup retention period if needed.
- Full backup: The system backs up all selected data even if no changes were made to the data after the last backup.
- Incremental backup: RDS automatically backs up data every 5 minutes.

Backup data restoration is implemented by restoring the data backup files stored in OBS to databases. You can use full and incremental backups to restore data to a required point in time. In addition, the Huawei Cloud GaussDB-based new database architecture (such as GaussDB for MySQL), with decoupled storage and compute, provides backup and restoration capabilities using snapshots. The backup and restoration is significantly faster than with traditional logical or physical backups (see the official website for more information). These databases can ensure zero RPO for customer databases.

Cloud Backup and Recovery (CBR) enables you to easily back up ECSs, BMSs, EVS disks, SFS Turbo file systems, local files, and on-premises VMware virtual environments. If there is a virus attack, accidental deletion, or software or hardware fault, you can restore data to any point in the past when the data was backed up. For details about CBR, visit https://www.huaweicloud.com/intl/en-us/product/cbr.html.

4.5.3.4 Change Management

Change management aims to standardize change activities (such as change request, review, scheduling, implementation, and verification) in the production environment. With clear requirements, make full preparation to reduce impacts on services, ensure

successful change activities and safe, stable production environment, maximize system availability, and meet SLAs.

4.5.3.4.1 Roles and Responsibilities

Role	Responsibility
Change applicant (Maintenance engineer)	Submits change applications according to standards, checks the submitted information, verifies changes, and confirms results.
Change manager (O&M owner)	Reviews and assigns change applications, conducts communication, schedules resources, and follows up with the change implementer to track the change progress and close the change applications.
Change review team (O&M owner and development owner)	Reviews and approves common change solutions.
Change implementer (Maintenance engineer)	Implements changes based on approved solutions.

4.5.3.4.2 Change Classification

- **Emergency Changes**

  These changes are proposed when the production environment is or about to be unavailable. There is not enough time to evaluate or approve such changes by the usual process. For example, an emergency change is required when there is a new version defect that directly affects user operations.

- **Major Changes**

  These changes are made during version upgrade and O&M of the service system. If service interruption lasts longer than 30 minutes during the change, notify affected users in advance.

- **Normal Changes**

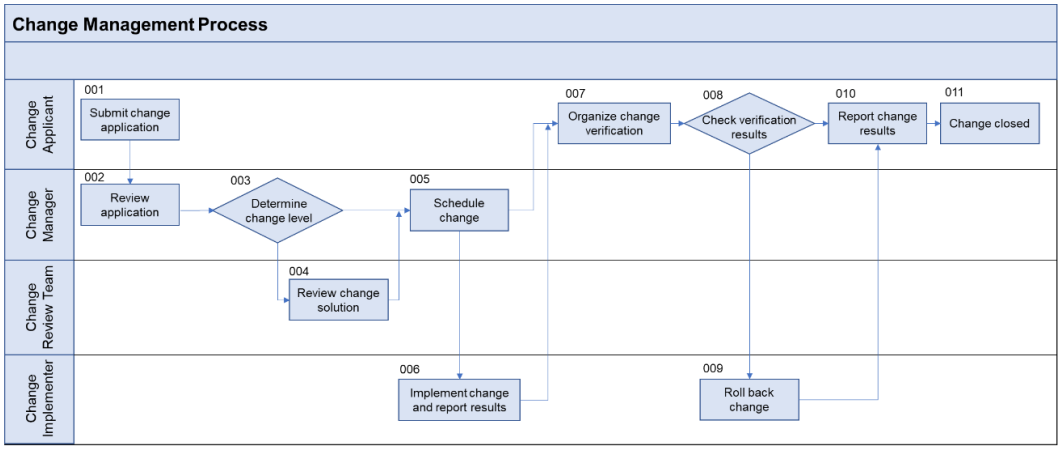
  These changes are made during version upgrade and O&M of the service system. If service interruption lasts less than 30 minutes during the change, notify affected users in advance.

- **Routine Changes**

  These are production environment changes that are low-risk. The change impact has been evaluated by O&M and development departments in advance, procedures are standard, implementers have been determined, and change effectiveness has been proven more than three times. This type of change is performed according to standard operation guides and will not affect services. Customers are unaware of such changes during upgrades.

4.5.3.4.3 Change Process

- **Flowchart**



- **An applicant submits a change application.**

  There are factors to determine before the change. If any are missing, provide reasonable explanation.

  - Reason and purpose: Explain why the change is required and what the change will do. If necessary, provide outline proof that the change is feasible.
  - Risks: Describe the potential problems brought about by the change and how to avoid them. Assess both current and related systems and domains for risks.
  - Impact: Describe how the change will affect customers and other domains.
  - Test report: If a test was previously conducted, attach its report to the application.

  The change guide should detail the backup, implementation, rollback, and verification schemes. If the change requires special implementation times, the applicant should specify them in the application.

- **The change manager reviews the application.**

  This review includes the change time window, plan, and type.

  - Check that the change time fits the overall planning. If not, advise the applicant to postpone the change or assist them in initiating the emergency change process.
  - Check that the change type (emergency, major, or normal) is reasonable and the change level is valid.
  - Determine whether the implementation time is reasonable, and tell the applicant to adjust it if necessary.

- **The change review team reviews the change solution.**

  - Check that the reason and purpose of the change are clear.
  - Check that the risks and measures are clear.
  - Check that the impact of the change is clear. For example, how users and the production environment will be affected during a specific period.
  - Check that the implementation process is valid, the procedure is reasonable and clear, and the purpose is achievable.

  - Determine whether a test report is required for the change. If so, check for an attached test report.
  - Check for any temporary backup, rollback, and verification before the change.
- **The change manager schedules the change.**
  - Check for any conflicts with other changes in time, content, and components, and assist the applicant in checking whether the change affects routine operations, shutdown, and backup.
  - Approve the change.
- **The change implementer implements the change.**
  - Wait for approval before proceeding.
  - Use the approved solution for the change.
  - If an unexpected service degradation, interruption, or data loss occurs during implementation, contact the change manager. Record and handle the exception.
  - After implementing and verifying the change, report the actual results and time taken.
- **The change applicant verifies the change.**
  - Verify the change according to the plan and check whether the purpose has been achieved. If it has not, report this to the change implementer.
  - Complete the verification report.
  - If an unexpected service degradation, interruption, or data loss occurs during verification, ask the change manager how to handle the exception. Record the exception.
- **The change implementer rolls back the change.**
  - If the change result does not meet expectations, roll back the change as planned. If services cannot be recovered within the expected time window, escalate the problem by following the incident management process.
  - If the change passes the verification or cannot be verified in the time window, and needs a rollback, submit an application.
- **The applicant reports the change results.**
  - Report the results promptly (ideally within two working days) after the change is completed.
  - Close the work order if the change is successful.
  - If the change fails, analyze why and provide the causes and improvement measures.

4.5.3.4.4 KPIs

- **Change success rate**

  The percentage of changes in a month that were successful. Only changes whose status is Successful when closed count.

  Formula: (Number of successful changes/Total number of changes) x 100%
- **Service interruptions**

  The number of service interruptions caused by changes in a month.

  Formula: Number of service interruptions caused by all changes in a month
- **Service interruption change rate**

  The percentage of changes in a month that interrupted services.

Formula: (Number of service interruption changes/Total number of changes) x 100%

- **Emergency change rate**

  The percentage of changes in a month that were emergencies.

  Formula: (Number of emergency changes/Total number of changes) x 100%

4.5.3.5 Emergency Handling

The process of receiving, handling, and escalating incidents must be standardized to ensure that customer issues are handled at the promised service level. The incident handling responsibilities, time requirements, and notification mechanism must be clearly defined. Services must be quickly recovered to ensure the promised quality and availability.

- **Roles and responsibilities**

Role	Responsibility
O&M engineer	• Receive customer incidents reported by hotline or email. • Record all information about received incidents, including contact methods of reporters, incident features, details, and time. • Diagnose and analyze incidents and provide solutions by referring to relevant documentation. For incidents that cannot be resolved, transfer them to developers or seek help from O&M leaders. • Hold first accountability to track incidents, record handling progress, and keep customers updated on the progress as required. • Demarcate incidents and provide solutions within a specified period. Transfer unresolvable incidents to developers within the specified time. • Close incidents that the reporters have confirmed resolved with the provided solutions. • Transfer reoccurring incidents and those with unknown causes or known defects to the issue management process. • Summarize lessons learned from typical and general incidents.
Developer	• Locate and analyze the causes of incidents transferred from O&M engineers, and resolve the incidents. • Follow up with incident owners to confirm resolution. • Locate causes of bugs in the production environment, and provide and implement comprehensive solutions.
Incident manager (O&M leader)	• Coordinate and monitor the incident handling process. • Coordinate resources for major incidents. • Review and approve solutions for major incidents. • Trace the output of major incident reports.

- **Incident severity and response**

  Level 1: incidents that have major impacts on services, such as serious damage, data loss, service data or function errors (which cause multiple customer complaints), and system faults recurring within a short period of time.

  Level 2: incidents that have minor impacts on services, such as unavailability of a few functions (service degradation), function impairment on some users, data inconsistency (no financial loss), and common system faults.

  Level 3: incidents that have no impact on services, such as data query and consulting. Services are normal but experience is affected.

  Response and resolution time requirements vary by incident level. The response timing is 24/7 and starts once an incident is reported.

Incident Level	Response Time (minutes)	Recovery Time (hours)	Resolution Time (days)
1	10	2	7
2	30	6	20
3	60	24	60

  Remarks:

  - The incident levels and response time above are only examples. Adjust them as required.
  - O&M engineers can transfer out incidents that they cannot resolve by referring to relevant documentation within the specified time.
  - The response time is the maximum delay before an incident handler starts handling an incident after receiving it.
  - The recovery time is the maximum duration needed to recover services after an incident occurs.
  - The resolution time is the maximum duration taken by an O&M engineer to resolve or transfer out an incident.

- **Incident escalation and notification**

  The notification mechanism defined in the following table is for incidents that are Level 1 and 2.

Notification	Method	Recipients
Initial notification (30 minutes)	SMS and email	(Send to) Service owner (CC) O&M team
	Phone call	R&D leader
Handling progress (1 hour)	SMS and email	(Send to) Service owner (CC) O&M team
Fault rectification	SMS and email	(Send to) Service owner (CC) O&M team
Escalation of overdue incidents	SMS and email	(Send to) Service owner (CC) O&M team, R&D

Notification	Method	Recipients
		leader
	Phone call	R&D leader

If a Level-2 incident's impact on services and users worsens, escalate it to Level 1 and then handle it with Level-1 standards.

- **Precautions during incident management**
  - Record all incidents on the live network in the unified event management system for analysis.
  - Respond to and resolve incidents within the specified time.
  - Analyze incidents handled every month.

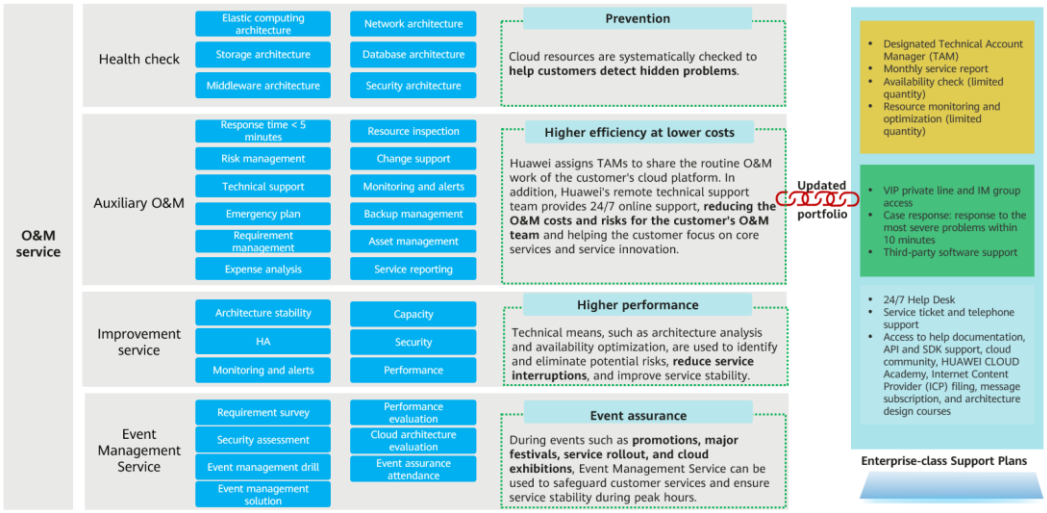
4.5.3.6 O&M Service

What are common maintenance problems faced by enterprises after they migrate their services to the cloud? Some typical challenges are listed below:

- Lack of troubleshooting capabilities: Assistance is required for fault locating and troubleshooting during the use of cloud products.
- Insufficient cloud O&M experience: Lack of cloud O&M personnel and related experience.
- Unstable services: The system frequently breaks down.
- Lack of special assurance: Stable and continuous service running must be ensured during major events.
- Lack of risk predication capabilities: Cloud resources need to be periodically checked to keep abreast of the overall health status of services.

To address these problems, Huawei Cloud provides four O&M services that are based on Huawei Cloud customers' best practices: health check, auxiliary O&M, improvement service, and Event Management Service.

- Health check: Helps customers analyze metrics and evaluate system efficiency, robustness, and security.
- Auxiliary O&M: Huawei provides customers with onsite O&M support based on Huawei Cloud best practices to ensure service continuity.
- Improvement service: Suggestions for improving customer service stability are provided based on full-link analysis of service status and Huawei Cloud best practices.
- Event Management Service: During major events, such as promotion activities, important festivals, service rollout and provisioning, and cloud exhibitions, Event
- Management Service can be used to ensure service stability during peak hours.



Service process:

- **Service request**
  1) A customer submits a service request.
  2) The delivery manager evaluates customer requirements.
  3) The delivery manager confirms the service target and scope with the customer.
  4) The customer places an order.
- **Service implementation**
  1) Delivery personnel deliver the service based on the service content.
  2) Delivery personnel provide deliverables to the customer.
- **Customer acceptance**

  The customer accepts and evaluates the service.

Best Practices

5.1 Customer Background

As a leading Internet platform enterprise in China, Company X has deployed services in many countries and regions, including China, Southeast Asia, Africa, Middle East, Latin America, and Europe. It provides information services such as social networking, news, and entertainment for global Internet users.

As a mobile payment subsidiary of Company X, XXpay has expanded from Africa to emerging blue ocean markets such as the Middle East, Southeast Asia, and Latin America. So far, XXpay has gained XX million users and XXX thousand merchants in the mobile payment market. They manage payment in excess of USDX billion, and the estimated market value has reached USDX billion. XXpay has become a FinTech unicorn in Africa and Middle East.

Nigeria, Africa's largest economy, occupies an important strategic position in FinTech. This makes Nigeria a key country where XXpay can develop FinTech services. XXpay's services have been growing too fast for the service nodes in Europe to maintain requirements for low latency.

XXpay's service and risk control data is not centrally managed. As a result, data value cannot be effectively explored and utilized. As XXpay's services grow, these problems have been getting more severe.

5.2 Solutions

5.2.1 One Global Network

Huawei Cloud KooVerse provides one global network to reduce network latency and provide global service coverage.

KooVerse network and nodes cover multiple 60-100 ms latency circles in countries and regions such as Asia Pacific, Europe, and Latin America. To provide better experience for Nigerian users, XXpay chose to migrate its services from another cloud in Europe to the Southern Africa region (located in Johannesburg) of Huawei Cloud. The Nigeria POP site in this region provides excellent private line networks for XXpay services.

5.2.2 Cloud Native

To better meet the performance requirements of FinTech services, Huawei Cloud upgraded its architecture to cloud native.

As a populous country in Africa, Nigeria's payment transactions can reach X million per day. The peak hours are in the morning and evening every day. Huawei Cloud infrastructure such as CCE containers and VMs are elastically scaled to properly adapt to XXpay's service volume changes in peak and off-peak hours. This will also greatly improve user experience in concurrent access.

A cloud native architecture will greatly improve service iteration, version release, O&M assurance, and resource utilization, which is critical to agile development and service innovation for XXpay in the future.

5.2.3 One-Stop Big Data

Huawei Cloud's one-stop big data solution integrates service and risk control data to help XXpay with precision risk control, real-time reports, and other data-related objectives. With a decoupled storage and compute architecture, 2 PB data can be properly processed, greatly improving the cost-effectiveness.

5.2.4 Migration Team and Solutions

Huawei Cloud organized a team with eight senior technical experts and engineers to assist XXpay in service migration. Migration solutions such as gray testing for core services and big data running in two environments were formulated to ensure smooth and lossless migration.

5.3 Customer Benefits

- **Low latency**: Huawei Cloud's Nigeria POP site in the Southern Africa region provides excellent private line networks for XXpay services. The latency is reduced from 130 ms to 80 ms.
- **One-stop big data, decoupled storage and compute**: Huawei Cloud's one-stop big data solution integrates service and risk control data to help provide XXpay with precision risk control, real-time reports, and other data-related objectives. With a decoupled storage and compute architecture, more than X PB data can be properly processed, improving the cost-effectiveness by 30%.
- **High concurrency**: Huawei Cloud infrastructure such as CCE containers and VMs can be elastically scaled to handle up to X million transactions per day. A single database can handle up to 4,000 TPS.
- **Quick service innovation**: Cloud native greatly accelerates service iteration and version release, and improves O&M efficiency and resource utilization.
- **Optimized interconnection with third-party systems**: Third-party systems include settlement systems with Nigeria Inter-Bank Settlement System (NIBSS) and merchant payment systems.
- **Smooth migration and efficient delivery**: A migration team is jointly established with customers to migrate services (xxxxx vCPUs included) within two months. The migration process is smooth and lossless, ensuring timely project delivery.

6 About Huawei Cloud

Founded in 1987, Huawei is a leading global provider of information and communications technology (ICT) infrastructure and smart devices. Huawei has approximately 205,000 employees and operate in over 170 countries and regions, serving more than three billion people around the world.

Huawei's mission is to bring digital to every person, home and organization for a fully connected, intelligent world. To this end, they focus on driving ubiquitous connectivity and promote equal access to networks to lay the foundation for the intelligent world; provide diversified computing power to deliver ubiquitous cloud and intelligence; build powerful digital platforms to help all industries and organizations become more agile, efficient, and dynamic; redefine user experience with AI, offering consumers a more personalized and intelligent experience across all scenarios, including home, travel, office, entertainment, and fitness & health.

In 2017, Huawei officially launched their Huawei Cloud brand, which opens Huawei's more than 30 years of technical prowess and product solutions in ICT to customers. Through Infrastructure as a Service, Technology as a Service, and Expertise as a Service, Huawei Cloud provides Everything as a Service to deliver reliable, secure, trustworthy, and sustainable cloud services for customers, partners, and developers.

Over the past few years, Huawei Cloud has adhered to technological innovation to lead industry development, strengthened software, and met customers' diversified service requirements through system and architecture innovation. They have proposed leading technical concepts and best practices such as Cloud Native 2.0, distributed cloud, and application modernization. Huawei Cloud is trusted by customers globally, as it provides neutral, secure, and reliable cloud services based on their in-depth understanding of enterprises.

Huawei Cloud made breakthroughs in Garter's Magic Quadrant for Cloud Infrastructure and Platform Services (CIPS), as well as Gartner's Magic Quadrant for Security Information and Event Management (SIEM), Enterprise Integration Platform as a Service (EiPaaS), and Network Firewall. Huawei Cloud has launched more than in excess of 248 cloud services and more than 78,000 APIs, attracted more than 41,000 partners and 4 million developers around the world, and released more than 10,000 applications in the KooGallery marketplace.

FB: https://www.facebook.com/profile.php?id=100090157621463

LN: https://www.linkedin.com/feed/update/urn:li:activity:7048579339952472064/

TW: https://twitter.com/4AfricaCloud/status/1642814162953748480

7 FINTECH ASSOCIATION OF NIGERIA

FinTech Association of Nigeria (FintechNGR), the premier national Fintech Association in Africa, is a self-regulatory, not-for-profit and non-political organization. The Association's vision is to make Nigeria one of the world's leading markets for FinTech Innovation and Investment and its mission is to foster an ecosystem that supports all stakeholders to achieve a thriving and growing Nigerian FinTech industry. Its objectives are to: ACCELERATE fintech growth and impacts, CONNECT fintech stakeholders within and outside Nigeria and ADVOCATE for progressive regulatory reforms/support, and an enabling innovation environment attracting investment.

Fintech NGR is governed by a Board of Trustees and a Governing Council both comprising of professionals of impeccable character.

The Association currently has 350 members across 21 sectors of the economy.

Further information can be obtained on the website of the Association and our social media handles:

www.fintechng.org

Twitter: @fintechngr

Instagram: fintechngr

LinkedIn: fintechngr

Facebook: fintechngr