



Sangfor SCP

Technical White Paper

Product Version	6.11.3
Document Version	02
Released On	May. 25, 2026



Copyright © Sangfor Technologies 2024. All rights reserved.

Unless otherwise stated or authorized, Sangfor Technologies Inc. (hereinafter referred to as "Sangfor") and its affiliates reserve all intellectual property rights, including but not limited to copyrights, trademarks, patents, trade secrets, and related rights to text, images, pictures, photographs, audio, videos, charts, colors, and layouts as presented in or concerning this document and content therein. Without the prior written consent of Sangfor, this document and content therein must not be reproduced, forwarded, adapted, modified, displayed, or distributed by any other means for any purpose.

Disclaimer

Products, services, or features described in this document, whether wholly or in part, may be not within your purchase scope or usage scope. The products, services, or features you purchase must be subject to the commercial contract and terms as agreed by you and Sangfor. Unless otherwise provided in the contract, Sangfor disclaims warranties of any kind, either express or implied, for the content of this document.

Due to product version upgrades or other reasons, the content of this document will be updated from time to time. Unless otherwise agreed, this document is used for reference only, and all statements, information, and recommendations therein do not constitute any express or implied warranties. Sangfor shall not be held liable for any loss or damage arising from omissions, changes, or errors in this document.

Contact Us

Pre-sales consulting: +86 (755) 8662-7888

After-sales service: +60 12 711 7129 (7511)

Official website: www.sangfor.com

About This Document

This document describes the architecture, features, and application scenarios of Sangfor Cloud Platform (SCP) to help sales, pre-sales, and technical support representatives become familiar with SCP and its application scenarios.

Intended Audience

This document is intended for:

- Product Manager
- Pre-Sale Manager
- Partner Manager

Note Icons

English Icon	Description
	Indicates an imminently hazardous situation which, if not avoided, will result in death or serious injury.
	Indicates a potentially hazardous situation which, if not avoided, could result in death or serious injury.
	Indicates a hazardous situation, which if not avoided, could result in minor or moderate injury.
	Indicates a hazardous situation, which if not avoided, could result in settings failing to take effect, equipment damage, or data loss. NOTICE addresses practices not related to personal injury.
	Calls attention to important information, best practices, and tips. NOTE addresses information not related to personal injury or equipment damage.

Contents

Contact Us.....	1
1 Overview	13
1.1 Operations Management Requirements for Cloud Platforms	13
1.2 Sangfor SCP	14
1.3 SCP Benefits.....	14
2 Application Scenarios.....	16
2.1 General Management	16
2.2 Hierarchical Management	16
3 Product Architectures	18
3.1 Functional Architecture.....	18
3.2 Deployment Architecture	20
3.2.1 General Deployment	20
3.2.1.1 Standalone Mode	20
3.2.1.2 Cluster Mode.....	22
3.2.1.3 Active/Standby Mode.....	24
3.2.2 Hierarchical Deployment	25
3.2.2.1 Sangfor Cloud with SCP Standalone Mode.....	25
3.2.2.2 Sangfor Cloud with SCP Cluster Mode	26
3.3 Technical Architecture.....	27
3.3.1 Microservice Architecture	27
3.3.2 Distributed Architecture.....	29
3.3.3 Cloud-Native Architecture.....	31
3.3.4 Large-Scale Management and Small-Specification Optimization	32
3.4 HA Architecture	35
3.4.1 Principles.....	35
3.4.2 Components	36
3.4.2.1 Cluster Management Components	37
3.4.2.2 Distributed Management Components	38
3.4.2.3 Database Management Components	38
3.4.2.4 Trend Data Management Components	39
3.5 Hierarchical Management Architecture.....	41
3.5.1 Design	41
3.5.2 Architecture	42
3.5.3 Capacities	43
3.5.4 User System	44

4 IT Infrastructure	46
5 Virtualized Resource Pools	47
6 Resource and System Management	50
6.1 Resource Management	50
6.1.1 Clusters	50
6.1.2 Resource Pools	50
6.1.3 Dedicated Server Groups	51
6.1.4 Heterogeneous Resource Management	51
6.1.4.1 VMware Clusters	51
6.1.4.2 Physical Hosts	51
6.1.4.3 AWS Cloud	53
6.1.5 Quotas	54
6.1.6 Application Center	54
6.1.7 Homogeneous Hybrid Cloud	57
6.2 System Management	57
6.2.1 Users	57
6.2.2 Policies	58
6.2.3 Certificates	58
6.2.4 Platform Security	58
6.2.4.1 Login Protection	58
6.2.4.2 Authentication	59
6.2.4.3 SSO	59
6.2.4.4 Key Pairs	59
6.2.4.5 SM Encryption Mode	59
6.2.5 System	60
6.2.5.1 Services	60
6.2.5.2 Date and Time	60
6.2.5.3 Platform	61
6.2.5.4 Data Center	61
6.2.5.5 SMS and Email	61
6.2.5.6 Components	61
6.2.6 Licensing	61
6.2.6.1 Unified SCP Licensing	61
6.2.6.2 Individual HCI Licensing	62
6.2.6.3 Licensing for SCP	63
6.2.6.4 Licensing for HCI Clusters	63
6.2.6.5 Licensing for Security Components and Advanced Components	63

6.2.6.6 vKEY Licensing	64
6.2.6.7 Grant License Based on CPU Cores	64
7 IaaS-Based Services	66
7.1 Compute	66
7.1.1 VM	66
7.1.1.1 Introduction	66
7.1.1.2 Application Scenarios	66
7.1.1.3 Supported Features	67
7.1.2 Bare Metal	67
7.1.2.1 Introduction	67
7.1.2.2 Application Scenarios	68
7.1.2.3 Supported Features	68
7.1.3 GPU-Enabled VM	69
7.1.3.1 Introduction	69
7.1.3.2 Supported Features	70
7.1.4 Image	72
7.1.4.1 Introduction	72
7.1.4.2 Supported Features	73
7.1.5 Auto Scaling	75
7.1.5.1 Introduction	75
7.1.5.2 Supported Features	75
7.1.6 aDesk VDI	76
7.1.6.1 Introduction	76
7.1.6.2 Supported Features	77
7.1.7 VM Customization Specifications	77
7.1.7.1 Introduction	77
7.1.7.2 Application Scenarios	79
7.1.7.3 Supported Features	79
7.2 Network	79
7.2.1 VPC	79
7.2.1.1 Introduction	79
7.2.1.2 Supported Features	81
7.2.2 EIP	85
7.2.2.1 Introduction	85
7.2.2.2 Elastic IP Pool	86
7.2.3 Shared Service Network	88
7.2.3.1 Introduction	88

7.2.3.2 Network Flow Direction	89
7.2.4 Network Orchestration	90
7.2.4.1 Classic Network	90
7.2.4.2 Tenant Network.....	91
7.2.5 Communication Domain.....	92
7.2.5.1 Introduction	93
7.2.5.2 Application Scenarios	95
7.2.6 Traffic Mirroring	97
7.2.6.1 Application Scenarios	97
7.2.6.2 Supported Features	98
7.2.6.3 Policy Configuration.....	98
7.2.7 QoS.....	100
7.2.7.1 Introduction	100
7.2.7.2 Application Scenarios	101
7.3 Storage	102
7.3.1 Block Storage	102
7.3.1.1 Introduction	102
7.3.1.2 Application Scenarios	102
7.3.1.3 Supported Features	103
7.3.2 File Storage.....	104
7.3.2.1 Introduction	104
7.3.2.2 Application Scenarios	105
7.3.2.3 Supported Features	105
7.3.3 Object Storage	107
7.3.3.1 Introduction	107
7.3.3.2 Application Scenarios	107
7.3.3.3 Supported Features	108
7.4 Cloud Security	109
7.4.1 Distributed Firewall.....	109
7.4.1.1 Introduction	109
7.4.1.2 Application Scenarios	110
7.4.1.3 Supported Features	110
7.4.1.4 Policy Configuration.....	111
7.4.1.4.1 Create Policy.....	111
7.4.1.4.2 Import Policy	112
7.4.2 VPC Firewall.....	114
7.4.3 aSecurity.....	115

7.4.3.1 Technical Architecture	115
7.4.3.2 Dashboard.....	116
7.4.3.3 VM Security.....	117
7.4.3.4 Security Capabilities.....	118
7.4.3.5 Logs	119
7.4.3.6 Security Settings	119
7.4.4 Security Components	120
7.4.4.1 vNGAF	120
7.4.4.2 aTrust	121
7.4.4.3 vIAG	123
7.4.4.4 vSSL VPN	125
7.4.4.5 vADC.....	126
7.4.4.6 vOSM.....	126
7.4.4.7 vDAS	127
7.4.4.8 Sangfor Host Security	127
7.4.4.9 vCyber Command-Logger	127
7.4.4.10 aSI	128
7.4.5 Data Encryption.....	131
7.5 Disaster Recovery	132
7.5.1 Data Protection Overview	132
7.5.2 Scheduled Backup.....	133
7.5.2.1 Introduction	133
7.5.2.2 Application Scenarios	133
7.5.2.3 Supported Features	133
7.5.3 CDP.....	134
7.5.3.1 Introduction	134
7.5.3.2 Application Scenarios	134
7.5.3.3 Supported Features	135
7.5.4 HotAdd-based Agentless Backup.....	135
7.5.4.1 Introduction	135
7.5.4.2 Supported Features	136
7.5.5 Remote Disaster Recovery.....	136
7.5.5.1 Introduction	136
7.5.5.2 Application Scenarios	136
7.5.5.3 Supported Features	137
7.5.6 Cloud Disaster Recovery	142
7.5.6.1 Introduction	142

7.5.6.2 Application Scenarios	143
7.5.6.3 Supported Features	143
7.6 Memory Tiering Service	143
7.6.1 Background.....	143
7.6.2 Technical Principles	144
7.6.2.1 Configuring Node Memory Tiering.....	147
7.6.2.2 Clearing Node Memory Tiering Configuration.....	151
7.6.2.3 Memory Calculation Algorithm Description	152
1.1.1. Enabling/Disabling Memory Tiering for VM Creation/Cloning	153
7.6.2.4 Enabling/Disabling Memory Tiering for VMs.....	154
7.6.2.5 Powering on VM	157
7.6.3 Identifying Memory Hot and Cold Pages.....	160
7.6.4 Migrating Hot and Cold Pages.....	161
7.6.5 Monitoring Hot and Cold Data in VM Memory	161
7.6.6 Disabling Memory Tiering for VMs	162
7.6.7 Replacing NVMe SSDs.....	163
7.6.8 Deployment Assessment	164
7.6.8.1 Hybrid Deployment Mode.....	164
7.6.8.2 Workload Assessment and Tiering Ratio	164
7.6.8.3 Server Hardware Configuration	164
7.6.8.4 Selecting NVMe SSD	165
7.6.8.5 NVMe SSD Failure Scenarios.....	165
8 PaaS-Based Services.....	166
8.1 Database	166
8.1.1 Introduction	166
8.1.2 Application Scenarios	167
8.1.3 Supported Features	170
8.1.3.1 Automated Deployment.....	170
8.1.3.2 Lifecycle Management.....	171
8.1.3.3 Self-Healing	172
8.1.3.4 Data Protection.....	173
8.1.3.5 Oracle Database Disaster Recovery	173
8.1.3.6 DTS	173
8.1.3.7 Account Management	174
8.1.3.8 DMC.....	175
8.1.3.9 Operations Management.....	175
8.1.3.9.1 Multi-Tenant Isolation.....	175

8.1.3.9.2 Quota Management.....	176
8.1.3.9.3 Work Orders.....	177
8.2 Middleware.....	178
8.2.1 Introduction.....	178
8.2.2 Application Scenarios.....	179
8.2.3 Supported Features.....	179
9 O&M Management.....	181
9.1 Monitoring Center.....	181
9.1.1 Monitoring Dashboard.....	181
9.1.2 Reports.....	185
9.1.3 Alerts.....	185
9.2 SkyOPS.....	186
9.2.1 Introduction.....	186
9.2.2 Supported Features.....	187
9.2.2.1 Disk Failure Prediction.....	188
9.2.2.2 Memory Disk Prediction.....	189
9.2.2.3 Resource Prediction.....	190
9.2.2.4 Alert Convergence.....	191
9.2.2.5 Root Cause Analysis.....	191
9.2.2.6 Full Link Failure Detection and Location.....	192
9.2.3 Benefits.....	194
9.3 Reliability.....	195
9.3.1 Holistic View.....	195
9.3.2 HA.....	196
9.3.2.1 Introduction.....	196
9.3.2.2 Application Scenarios.....	196
9.3.2.3 Supported Features.....	196
9.3.3 Auto Scaling.....	197
9.3.3.1 Introduction.....	197
9.3.3.2 Application Scenarios.....	198
9.3.3.3 Supported Features.....	198
9.3.4 DRS.....	198
9.3.4.1 Introduction.....	198
9.3.4.2 Application Scenarios.....	199
9.3.4.3 Supported Features.....	199
9.3.5 Resource Reservation.....	200
9.3.5.1 Introduction.....	200

9.3.5.2 Application Scenarios	200
9.3.5.3 Supported Features	200
9.3.6 VM Resource Limits	200
9.3.6.1 Introduction	200
9.3.6.2 Application Scenarios	201
9.3.6.3 Supported Features	201
9.3.7 VM Scheduling	201
9.3.7.1 Introduction	201
9.3.7.2 Application Scenarios	202
9.3.7.3 Supported Features	202
9.3.8 SCP Status Check.....	202
9.3.8.1 Introduction	202
9.3.8.2 Application Scenarios	203
9.3.8.3 Supported Features	203
9.3.9 Hardware Health Check	204
9.3.9.1 Introduction	204
9.3.9.2 Application Scenarios	204
9.3.9.3 Supported Features	204
9.4 Security.....	204
9.4.1 Login Policies	204
9.4.2 Disk Encryption.....	205
9.4.3 Data Erasure	205
9.4.4 Key Pairs	205
9.4.5 Key Management Service.....	205
10 Operations Management	207
10.1 Multi-Level Roles.....	207
10.1.1 Platform Administrator	207
10.1.2 Tenant.....	209
10.1.3 Tenant User	210
10.2 Multi-Tenant Management.....	211
10.3 Quota Management	212
10.4 Work Orders	213
10.5 Billing	215
10.6 Tasks	216
10.7 Recycle Bin	216
10.8 Federal Identity Authentication	217
10.9 Custom Attributes	217

10.10 Portal Customization.....	217
11 Access Layer	219
12 SCP and Sangfor Cloud Portals.....	220
13 Open APIs	222
14 Glossary	223
15 Abbreviations	225

1 Overview

1.1 Operations Management Requirements for Cloud Platforms

With the development of digital transformation, more and more organizations and enterprises are turning to cloud platforms as the IT infrastructure for hosting digital applications. Their business services impose increasing requirements on the cloud platform on aspects such as service directory, reliability, usability, and maintainability to meet the rapidly changing market demands. The following are the characteristics related to operations and maintenance management required for cloud data centers:

1. Multi-Tenant Support: Isolates resources and services of each tenant from each other for security so that each tenant has its resource needs and quota limits.
2. Quick Response: Quickly adjusts services and limits to meet changes in business requirements.
3. High Scalability: Supports rapid elastic scaling to meet business needs.
4. High Security: Applies effective security protection by periodically updating and maintaining security policies and monitoring and responding to security events as the importance of the data and applications.
5. High Compliance: Ensures security and compliance of services, including data security, access control, backup, and recovery.
6. Automation: Uses automation tools and technologies to simplify processes for operations and maintenance (O&M) management to eliminate the difficulties of manual management and maintenance due to the scale and complexity of cloud platforms.
7. Monitoring and Alerting: Provides real-time monitoring and alerts to guarantee the stability and availability of services.
8. Cost Optimization: Optimizes O&M costs by allocating resources to prevent them from being idle.

1.2 Sangfor SCP

Sangfor provides a comprehensive cloud data center solution based on Sangfor Hyper-Converged Infrastructure (HCI) to help users build the cloud platform base for digital business needs. Users only need to deploy common x86 servers, ARM servers, and layer 2/3 switches to implement a software-defined cloud infrastructure with resource pools of compute, storage, network, and security through the operating system (OS) of cloud platforms. Then Sangfor Cloud Platform (SCP) can provide unified management of resource pools on cloud platforms and a wide range of cloud services, covering management scenarios of tenants, resources, services, security, monitoring, and O&M.

SCP can not only manage HCI clusters, but also VMware clusters to provide full lifecycle management for VMs on VMware vCenter, and can be integrated with third-party public cloud platforms such as AWS Cloud to act as a unified management platform for heterogeneous multi-cloud resource pools.

For hierarchical management and O&M, Sangfor offers homogeneous multi-cloud unified management platforms: Sangfor Cloud with SCP, which disperse cloud platform computing resources based on business needs and unify management and O&M according to the organizational structure.

1.3 SCP Benefits

Based on over 10 years of construction experience and thorough analysis of O&M management of cloud platforms, Sangfor launched SCP, a cloud management platform, to help customers build hyper-converged technology-based cloud platforms for continuous updates to meet business needs. SCP has the following benefits:

1. Large-Scale: Supports distributed cluster architecture for high-availability systems, for example, an SCP cluster with three nodes can manage up to 2,000 tenants, 128 HCI clusters, and more than 60,000 VMs; also supports deployment and management of small and medium cloud platforms through a single SCP node.

2. **Highly Reliability:** Built on cloud-native container architecture with capabilities such as health checks, failure self-healing, and fast-rolling upgrades.
3. **Easy to Use:** Provides a virtualized management UI that enables administrators to directly understand the status and usage of cloud resources and services to improve efficiency and reduce risks.
4. **High Compatibility:** Supports unified management of physical hosts, VMware vCenter, and AWS Cloud and maximizing the use of their resources with minimal operations.
5. **High Openness:** Offers numerous APIs to quickly integrate with users' informationization platforms and O&M management systems.
6. **Business Innovation:** Supports quick deployment, expansion, and coming online of cloud services to help enterprises innovate their business services.
7. **Cost Control:** Helps users improve the utilization of cloud resources and services and control relevant costs through the billing feature.
8. **Security Management:** Provides comprehensive security services for business services as well as built-in security management capabilities such as security audit, security policy management, and security event response.

2 Application Scenarios

SCP is widely used in the general or hierarchical management of cloud data centers in various industries to support business services and resource management.

2.1 General Management

General management refers to the management of cloud platforms in data centers through a SCP cluster. A cloud platform can contain multiple HCI clusters deployed in physical data center rooms in the same or different locations. When HCI clusters on a cloud platform are located in different data center rooms, it is recommended to deploy SCP in the primary data center room and connect the primary room with other rooms where HCI clusters are deployed through direct connections or bare fibers to ensure normal network communication between SCP and HCI clusters and prevent alerts on SCP. An SCP cluster supports the management of up to 128 HCI clusters.

In this scenario, all clusters on a cloud platform are usually managed centrally by an IT O&M team.

2.2 Hierarchical Management

With the continuous progress of digital transformation, the current group enterprises, multi-factory enterprises, multi-branch enterprises, and multi-campus schools all have many data centers and cloud platforms. Cloud platforms are not only required to host business services or data production but also required to be managed hierarchically. Sangfor offers a hierarchical management and O&M solution for multi-cloud platforms by deploying both Sangfor Cloud and SCP, which disperses cloud platform computing resources based on business needs and unifies management and O&M according to the organizational structure.

In this scenario, it is recommended to deploy Sangfor Cloud in the primary data center room, deploy SCP nodes in data center rooms with HCI clusters as needed, and connect the primary room to these other rooms through direct

connections. When the communication between Sangfor Cloud and SCP is interrupted due to network failure, administrators can manage the corresponding HCI clusters through SCP for regional autonomy.

3 Product Architectures

3.1 Functional Architecture

Software-defined cloud data center architecture can be layered by features. The advantage of layered design is that even though the services and architecture of each layer depend on the functional modules of the next layer, they all interact with each other using common or standardized interfaces, making them dependent on and independent of each other at the same time. This loosely coupled IT architecture gives greater flexibility in updates and eliminates any incompatibility with the underlying layers.

Sangfor SCP, a unified O&M management platform for cloud data centers, is designed with a layered architecture and mainly organized into the following modules in terms of logical functions: IT infrastructure, cloud resource pools, platform management, cloud services, cloud service management, O&M management, operations management, access layer, and cloud service portals. The overall functional architecture and modules of SCP are described below:




Figure 1: SCP Functional Architecture

- **IT Infrastructure:** IT hardware for building cloud platforms, including standard x86 and ARM servers, standard layer 2/3 switches, hardware firewalls, distributed and centralized storage, and bare metal servers.

- **Cloud Resource Pool:** Based on the underlying IT infrastructure, hosts, storage devices, network devices, and other hardware are connected to pool resources to enable high availability, scalability, and stability of cloud platforms; includes resource pools for computing, network, storage, security, and backup.
- **Resource Management:** Used to manage the underlying HCI resource pools, VMware resource pools, and AWS Cloud resources; manages quotas; provides the Application Store to flexibly add custom cloud services.
- **System Management:** The basic management module of cloud platforms, providing upward management of policies, users, platform security, systems, and certificates.
- **Cloud Services:** Cloud resources and services provided by cloud platforms to tenants' business services. SCP provides various IaaS and PaaS capabilities to support tenants' business needs for digital transformation. Typical IaaS-based services cover computing, network, storage, security, and disaster recovery; typical PaaS-based services include containers, microservices, databases, middleware, and edge computing platforms. Depending on the systems on which users' cloud platforms are running business services, SCP also has the flexibility to extend and enable additional services such as communication domains, Network Insight, Application Store, and aSecurity.
- **O&M Management:** Guarantees the stability of cloud platforms and the continuity of the services running on them. SCP provides cloud platforms with monitoring and O&M capabilities covering hardware, cloud platform components and resources, and business applications, with O&M features such as security management, operation upgrades, recycle bins, monitoring platforms, and alert centers.

- **Operations Management:** Builds granular operations capabilities from the perspective of serving tenants for the sustainable development of cloud platforms, including tenant management, expense management, quota management, and work order services.
- **Access Layer:** Provides a series of access and control policies for tenants and administrators to access cloud platforms, mainly through gateways of Sangfor-WAF, traffic, business services, and OpenAPI.
- **SCP Portal:** The unified registration and login portal of cloud platforms for tenants and administrators.

3.2 Deployment Architecture

The deployment architecture of SCP varies with different application scenarios. When only SCP is used as a cloud data center management platform, it is deployed in a general architecture; when SCP and Sangfor Cloud are used to build a cloud data center with hierarchical management, they are deployed in a hierarchical architecture.

3.2.1 General Deployment

Different cloud platforms have different numbers of physical servers, HCI clusters, and tenants. To conform and manage cloud platforms of various sizes, SCP supports standalone and cluster deployments and switching from standalone mode to cluster mode to scale out and obtain high availability (HA).

3.2.1.1 Standalone Mode

In standalone mode, SCP is recommended to be deployed on a single NIC, so that administrators and tenants of cloud platforms access SCP through the management interface. It is recommended that the SCP VM resides on the same network segment as that of the managed clusters and communicates with them through an edge.




Figure 2: Standalone Mode

A single SCP VM with 16 CPU cores and 26 GB of memory can manage 32 cloud platforms and nearly 3,000 VMs. By increasing the SCP VM specifications, up to 256 cloud platforms and nearly 30,000 VMs can be managed.

SCP can be switched from standalone mode to cluster mode. With the wizard-based switching feature on the SCP management UI, cloud platform administrators can quickly switch SCP during cloud scale-up.




Figure 3: Switch SCP from Standalone Mode to Cluster Mode

3.2.1.2 Cluster Mode

In cluster mode, the transaction processing performance of the cloud management platform can be dramatically improved to enable high concurrency and high availability for large-scale cloud platform delivery scenarios. It is recommended to deploy an SCP cluster consisting of three nodes with the following architecture:




Figure 4: Cluster Mode

As shown in the figure above, an SCP cluster is under a dual-NIC deployment. One NIC is used for the management network access of administrators and the other NIC is used for the business network access of tenants, thus splitting the management network from the business network.

The management network must be on the same network segment as the HCI cluster for the management and configuration of resources, accounts, and quotas. Administrators can manage the SCP cluster by logging in to the controller node through the cluster IP address.

The business network is used for tenants or tenant users to individually access SCP to use the allocated resources. Tenants can access the nodes in the SCP cluster with the load balancer (Active Directory) provided by the business network frontend for large concurrent business access traffic.

The following are descriptions of the main features of each module of this deployment architecture:

- **Active Directory (AD):** A load balancer that adds SCP nodes to the node pool and schedules external requests to available nodes through real-time health detection and load balancing algorithms.
- **Cluster IP:** Bound to the controller node as the management IP address, which will automatically float to a new node when the controller node goes offline.
- **Cluster Management Components:** Used to perform management operations on nodes in the cluster, including adding, deleting, and editing; detect and maintain node status by using available nodes to reorganize the cluster when a node goes offline, and re-adding it to the cluster when it comes online.
- **Distributed Management Components:** Used for real-time operations of services and scheduling and mutual exclusion of scheduled tasks on the nodes in the cluster.
- **Database Management Components:** Used to manage the database nodes in the cluster and select multiple nodes to form a database cluster. When a node goes offline, the cluster will be reorganized with available nodes to ensure HA.

3.2.1.3 Active/Standby Mode

SCP also can be deployed in active/standby mode to enhance the reliability of the cloud management platform. This architecture is typically applicable to scenarios such as remote disaster recovery. For example, the active SCP node is deployed in a cloud platform cluster in the active server room, and the standby SCP node is deployed in a cloud platform cluster in the disaster recovery server room. When the active node is unavailable due to a failure in the active server room, the standby node can be manually switched to the active node to guarantee the HA of the cloud management platform.




Figure 5: Active/Standby Mode

As shown in the figure above, the active/standby mode can be formed by two SCP VMs through the active/standby management services, where the system configurations and databases of the active node can be synchronized to the standby node in real time. The main modules and their features of this architecture are described as follows:

- **Active/Standby Management Services:** Used to add, delete, and edit active and standby nodes; manages data synchronization.
- **Configurations:** Used to save platform configuration information, such as secret keys.
- **Databases:** Used to save resource information, account configuration information, etc.

3.2.2 Hierarchical Deployment

To manage cloud data centers in a hierarchical architecture, both Sangfor Cloud and SCP need to be deployed. As SCP can be deployed on a single node or in a cluster, the hierarchical architecture is also deployed in the following two ways.

3.2.2.1 Sangfor Cloud with SCP Standalone Mode

When SCP is deployed in standalone mode, administrators and tenants on Sangfor Cloud access SCP through the SCP management interface, from which

Sangfor Cloud collects monitoring and logging data and uploads them to its data storage layer through the data event bus.




Figure 6: Sangfor Cloud with SCP Standalone Mode

3.2.2.2 Sangfor Cloud with SCP Cluster Mode

When the SCP is deployed in cluster mode, administrators on Sangfor Cloud access SCP through the SCP management interface, while tenants access SCP through the SCP business interface. Sangfor Cloud collects monitoring and logging data and uploads them to its data storage layer from SCP management interface through the data event bus.




Figure 7: Sangfor Cloud with SCP Cluster Mode

3.3 Technical Architecture

Designed in an advanced layered, distributed, and microservice architecture and developed on the Kubernetes container platform, SCP is a cloud management platform with high reliability, multiple deployment architectures, and large-scale cloud platform management capabilities.

3.3.1 Microservice Architecture

Based on the domain model of cloud services, SCP microservice architecture can be abstractly split into multiple modules that communicate with each other through standard REST APIs and decouple services internally through RPC APIs. SCP has built-in microservice features and management capabilities because each internal microservice module is developed on a framework with unified

engineering and business service bases.

SCP provides external standard REST APIs for secondary development. API standards include the specifications of OpenStack APIs and private OpenAPIs.




Figure 8: Microservice Architecture

The main microservice components of SCP are unified for authentication and zero-trust permission management through one API gateway to meet the role-based access control (RBAC) requirements for cloud multi-tenancy. The internal microservices access each other through service registration, authentication, and proxy for distributed scaling. The internal services provide capabilities such as distributed link tracking, anomaly information collection, and service topology during business service requests through a unified management framework. Services are based on stateful and stateless layers, and stateless layer services can be elastically scaled for large-scale cloud platform management.




Figure 9: Microservice Workflow

3.3.2 Distributed Architecture

SCP is designed on an architecture that distributes stateless services and management components on different nodes, and proxies forwarding through loads to share the access pressure from business services and enable large scale-out for cluster deployment.




Figure 10: Distributed Architecture

The main components of the distributed architecture are as follows:

- **Pre-Scheduler:** Provides access portal to VM IP addresses and loads external requests to nodes in the cluster according to policies.
- **Business IP:** The actual business network interface IP address of the node, which is required to be configured to the node pool of the pre-scheduler.
- **Stateful Services:** Used to manipulate public data, such as clearing resources, generating O&M report information, etc.
- **Stateless Services:** Node-based services for processing business service requests and not for data storage.
- **Data Storage Services:** Used to store the public data of the cluster in databases such as MySQL, Redis, and MongoDB.

- **Proxy Client:** Used to forward requests from stateless services to access public data to the controller node to access the data storage services.
- **Cluster Management Client:** Used to reset the destination IP address of the proxy client when the controller node is switched so that it can connect to the new controller node.
- **Cluster Management Components:** Used to maintain the nodes in the stateful node pool by notifying other nodes to reselect the controller node when a node fails.

3.3.3 Cloud-Native Architecture

The underlying layer of cloud-native SCP is based on optimized Kubernetes (K8s), which is lighter and safer, compatible with edge environments and heterogeneous deployments. The cloud-native technology enables SCP with health checks, failure self-healing, fast rolling upgrades, advanced application deployment, service discovery, continuous monitoring and deployment, and other automation features. The internal SCP services are deployed through containerized integration and decoupled from the operating systems so that business services can run on different operating systems. SCP makes business services diverse, pluggable, and hot-loadable by extending cloud-native XaaS applications.




Figure 11: Cloud-Native Architecture

3.3.4 Large-Scale Management and Small-Specification Optimization

SCP implements large-scale cloud platform management through caching, distribution, message bus, dynamic-static separation, and other technologies with the following mechanism:




Figure 12: Large-Scale Management Mechanism

This mechanism mainly features the following technologies:

1. The access gateway enables fast page response and high concurrency by querying the configuration cache. The cache query framework aggregates data from cached databases of multiple nodes through distributed technology.
2. When the resource configurations are changed, the message bus will notify the data collector to update the cached data.
3. The SCP cluster pressure is dispersed by scheduling the clusters managed by each SCP node through the distributed technology.

- The configuration data collection is reduced by applying the dynamic-static separation for data configuration and monitoring.

In addition, SCP supports the deployment of resources at low specifications (such as 8 cores, 10 GB memory). Resource usage is reduced by using service start/stop and dynamic business service merging technologies. SCP adopts the following optimization mechanism for low-specification resources.




Figure 13: Small-Specification Optimization Mechanism

The main technical characteristics are as follows:

- Pluggable modules:** For some non-critical functional modules, services can be started or stopped on the advanced service page or by importing a license key.
- Service merging:** The process model is switched to a coroutine model if the system specification is detected to be low, after which module processes are automatically merged.

In addition, SCP supports the deployment of resources at low specifications (such as 8 cores, 10 GB memory). Resource usage is reduced by using service start/stop, dynamic business service merging, and memory compression technologies. SCP adopts the following optimization mechanism for low-specification resources.




Figure 14: SCP Low-Specification Resource Optimization Mechanism

The main technical characteristics are as follows:

- **Pluggable modules:** For some non-critical functional modules, services can be started or stopped on the advanced service page or by importing a license key.
- **Service merging:** The process model is switched to a coroutine model if the system specification is detected to be low, after which module processes are automatically merged.
- **Dynamic scaling:** Services can dynamically scale down the number of processes based on request volume.
- **Memory compression:** Memory can be compressed using ZRAM technology.

3.4 HA Architecture

3.4.1 Principles

SCP in cluster mode has high availability capability. As shown in the following figure, stateless business services can run on all nodes in the cluster, while stateful services can only run on the leader node. The backup node acts as a replica node to provide strong database synchronization, and the follower node is used for load balancing.




Figure 15: Cluster Mode Structure

When cluster management components detect that the leader node is abnormally offline, the backup node will be switched to the leader node, and the follower node will be upgraded to the backup node to continue strong database synchronization, and then the offline leader node will be set to the follower node and wait to come online to join the cluster automatically to ensure consistent data backups between the active and standby nodes.




Figure 16: Active/Standby Node Switchover

After an offline node comes back online and passes checks, it will rejoin the cluster to run.




Figure 17: Recovered Offline Node

3.4.2 Components

The high-reliability architecture of SCP cluster consists of cluster management components, distributed scheduling components, and database management components, and requires cluster IP address and AD. When the original controller node is abnormal, the cluster IP address can float to a new node to prevent the failure from being perceived by the users; and AD can schedule the external requests to the nodes with low loads to mitigate the single-node pressure.

3.4.2.1 Cluster Management Components

The cluster management components are mainly used for adding or deleting cluster nodes and performing heartbeat detection on them to automatically isolate failed nodes to ensure cluster reliability. The architecture and main functional modules are shown below:




Figure 18: Cluster Management Components

- **Etcd:** An open-source consistency component based on the Go programming language that provides heartbeat detection, event watch mechanism, and consistency configuration storage.
- **Rubik:** The cluster management service developed by Sangfor for the entire cluster lifecycle management. Rubik on each node will register a heartbeat with etcd. When a node hangs, its heartbeat stops, and other nodes will identify the offline node and process correspondingly.
- **Communication Proxy Service:** Load software used to forward traffic accessing services on a single node to the leader node.

3.4.2.2 Distributed Management Components

The distributed management components are mainly used for excluding failed nodes to perform interactive scheduling and load balancing scheduling of cluster tasks. The architecture and main functional modules are shown below:




Figure 19: Distributed Management Components

- **Distributed Scheduling:** Used to schedule tasks to available nodes for execution by querying node statuses with the consistent hash algorithm.
- **Distributed Locks:** Used to mutually exclude tasks through distributed Redis and etcd locks that are stored in memory and on disks respectively.

3.4.2.3 Database Management Components

The database management components are mainly used for managing databases in the cluster by forming a cluster of databases from multiple nodes

to ensure strong data consistency and reliability, and automatically rebuild the database cluster when a database is detected as abnormal. The architecture is shown below:




Figure 20: Database Management Components


The main features include database cluster creation, abnormal node database isolation, and consistent database data synchronization management.

1. The database of the controller node acts as the active database and the other available nodes act as standby databases in a database cluster.
2. A standby node that is detected as abnormal will be removed from the database cluster and replaced by another available node.
3. When the controller node is offline, the database of a standby node will be switched to the active database, and the database cluster will be reorganized by other nodes.
4. Data reliability is ensured by the strong consistency mechanism of database synchronization.

3.4.2.4 Trend Data Management Components

The trend data management components are mainly responsible for collecting trend data for various resources in the cluster. They aggregate the trend databases on multiple nodes in the cluster into a cluster and use a multi-write

mechanism to create data replicas to ensure reliability. The architecture of the trend data management components is shown below.



- In the active/standby mode or cluster deployment mode, the monitoring data will not be synced from existing nodes to new standby nodes or standby and slave nodes newly added to the cluster.
- In scenarios such as poor network connectivity or network disconnection among nodes, the monitoring data cannot be written from active nodes to standby nodes. In such cases, the data will not be backfilled to the affected nodes again and will only be retained on the nodes where the write operation succeeded.

- In cluster deployment mode, if an active node fails and leaves the cluster, any monitoring data that failed to be written to standby or slave nodes due to poor network connectivity or network disconnection will be permanently lost unless the active node rejoins the cluster.
- **When SCP is upgraded from SCP 6.11.2 or earlier to SCP 6.11.3 or later, trend data migration will be performed. The migration will take up to 15 minutes based on the trend data size in the historical version. You can view the progress in Tasks.**

3.5 Hierarchical Management Architecture

In hierarchical management scenarios, Sangfor Cloud is deployed in the primary data center room, and SCPs are deployed in each data center room with HCI clusters as needed. Cloud platform administrators, tenants, and second-level operators access Sangfor Cloud to log in to cloud platforms for O&M management and services.

3.5.1 Design

Based on the unified infrastructure Sangfor HCI and connected to Sangfor Cloud, SCP provides comprehensive cloud service directory capabilities for data center cloud platforms and manages heterogeneous resource pools with multi-tenancy capabilities.




Figure 21: Hierarchical Management Design

3.5.2 Architecture

The following figure shows the logical architecture of hierarchical management of Sangfor Cloud and SCP for cloud data centers:




Figure 22: Hierarchical Management Architecture

The cloud management platform consists of three layers:

- Cloud Service Layer:** Provides business service features externally, which can be divided into operations plane, O&M plane, and homepage visualization according to the features.

The operations plane mainly provides global data center account and quota management, resource statistics, and login authentication features, while the O&M plane provides monitoring, alerts, and data center room features.

- **Basic Component Layer:** Provides basic capabilities internally to serve the cloud service layer, mainly including command and data bus, API gateway, message notification, user permission system, and data center management.
- **Cloud Platform Resource Layer:** Deploys a cloud platform agent in each cloud data center to collect their operational data and O&M monitoring information and summarize them to the cloud service layer for operations and O&M management.

3.5.3 Capacities

Sangfor Cloud:

1. The operations plane provides distributed cloud platforms with global data center account and quota management, resource statistics, and login authentication features, while the O&M plane provides monitoring, alerts, and data center room features.
2. Sangfor Cloud uses a command bus to synchronize global operations performed by users on cloud service platforms, such as account creation to each data center, and supports single sign-on to data centers in different regions for resource configuration to improve user experience and utilization efficiency.
3. Sangfor Cloud uses a data bus to deploy the DBLog Agent in both the cloud data centers and the cloud service platforms to collect log and alert data for unified analysis and display of user operations. It also uses big data and other processing engines at the database bus layer for subsequent data analysis, BI reporting, and predictive analytics to make operations smarter.

SCP:

4. In hierarchical management scenarios, administrators or second-level administrators perform O&M operations on managed clusters through SCP in the same way as in the general management architecture.
5. Cloud tenants log in to Sangfor Cloud through a unified portal, and then redirect to the SCP corresponding to an available region to perform operations on the resources in the same way as in the general management architecture.

3.5.4 User System

The cloud management platform has the following user roles in the hierarchical management architecture to form a user system for different scenarios:

Role	Description
Super Admin	The default role with all administrative permissions for the platform system.
Platform Admin	An extended role with a subset of the super administrator's permissions to assist with some operations.
MSP Admin	A role in the managed cloud services also known as operations administrator that proxies resources for selling and can create tenants.
Tenant	The default role with all the tenant's permissions.
Tenant Co-Admin	An extended role with a subset of the tenant's permissions.
Tenant User	The end-user of the resources that cannot create resources but only apply for them from the tenant through a work order.

Table 1: Platform User Roles




Figure 23: User Level Model

As shown in the figure above, user roles are classified into three levels: the super administrator level, the MSP administrator level, and the tenant level. The MSP administrator level can be divided into two sub-levels, whose roles have the same capabilities but have hierarchical relationships; the tenant level consists of the basic level and the extended level, and the extended level roles inherit some of the capabilities from the basic level roles.

4 IT Infrastructure

The IT infrastructure layer is the underlying hardware base for building a cloud platform with resources including physical servers, switches, storage, hardware firewalls, etc.

SCP supports a rich variety of IT infrastructure hardware, which can be viewed through the Sangfor HCI Compatibility Chart.

5 Virtualized Resource Pools

Based on the underlying IT infrastructure, hosts, storage devices, network devices, and other hardware are connected to pool resources through Sangfor HCI to enable high availability, scalability, and stability of cloud platforms with resource pools for compute, network, storage, security, and backup.

HCI is composed of aSAN, aSV, and aNET components. In terms of logical architecture, each component can be grouped into a different data layer and a unified management and control layer, which involves main features such as configuration management, resource management, monitoring services, and advanced features. HCI can be integrated with SCP through APIs for cloud data center management.

The HCI system architecture can be functionally characterized into three layers. The description and figure of the layered architecture are as follows:

1. **Physical Layer:** It consists of infrastructure hardware, including physical resources such as CPU, memory, HBA cards, GPU, NICs, RAID cards, SSDs, and HDDs. These physical resources can be managed and controlled by the HCI cluster management software.
2. **Virtualization Layer:** As the main component of HCI, it comprises aSV, aSAN, aNET, and aSecurity and provides virtualized computing, storage, and network resources for the upper layer and cloud layer.
3. **Management and Control Layer:** It is responsible for cluster management on HCI and functions configuration management, resource management, monitoring service, and advanced features. This layer also supports integration with the cloud platform through northbound APIs to implement the management of cloud data centers. A layered architecture avoids deploying management layers on every node and reduces resource consumption and message interactions, thereby enhancing management efficiency.
 - **Configuration Management:** Used to convert or deliver configurations for aSV, aSAN, aNET, and aSecurity, which

includes topology calculation, error recovery, data storage, framework distribution, and read/write paths.

- **Resource Management:** Used to manage resources for aSV, aSAN, aNET, and aSecurity, which includes nodes, storage capacity, network visualization, VMs, network topology, virtual storage, and physical hosts.
- **Monitoring Service:** Used to monitor the components from the following aspects: usage, read/write speed, alert notification, information collection, physical host, and storage status, VM and virtual storage status, host health, one-click virtual environment status check, and performance and system bottlenecks.
- **Advanced Features:** Used to offer advanced features such as backup and continuous data protection (CDP), snapshots, HA, dynamic resource scheduling (DRS), dynamic resource eXtension (DRX), aSecurity, and network functions virtualization (NFV).




Figure 24: HCI Layered Functional Architecture

Sangfor SCP unifies the management of server clusters with multiple architectures (e.g., x86 and ARM) to meet the business needs of users in digital transformation.




Figure 25: Unified Management of ARM and x86 Clusters

6 Resource and System Management

The resource management and system management are the basic modules for SCP to integrate with HCI resource pools and to provide administrators and cloud tenants with cloud resources and services. This chapter describes their main features.

6.1 Resource Management

The resource management module is responsible for the management of the underlying HCI resource pools, VMware resource pools, and AWS Cloud resources and quotas. It also provides the Application Store to flexibly add custom cloud services.

6.1.1 Clusters

On a cloud platform, a cluster is the physical environment for building resource pools and consists of multiple physical servers. These servers are deployed with hyper-converged software for unified management and scheduling and are connected to each other through high-speed networks for joint provisioning of compute, storage, and network resources, thus forming a highly integrated system. A cluster can be scaled by rapidly adding compute, storage, and network resources based on business needs.

SCP supports adding multiple HCI clusters and VMware clusters for management. Only resources from clusters associated with a resource pool can be used for VMs or applications on SCP.

6.1.2 Resource Pools

A resource pool is a collection of cluster resources on SCP. Administrators can associate added clusters to different resource pools, and manage and allocate resources in them to tenants. Currently, a resource pool can only be associated with only one cluster. A cluster can be added with multiple storage. It is recommended to label storage of different types and performance with

different storage tags.

6.1.3 Dedicated Server Groups

A dedicated server group is a collection of specified physical hosts (which can also be a single host) in a cluster on SCP. Compute resources are used separately and storage resources are shared between dedicated server groups in the same cluster. A super administrator can allocate a dedicated server group to a tenant or MSP administrator for the exclusive use of compute resources in the group.

6.1.4 Heterogeneous Resource Management

SCP supports unified management of resources on heterogeneous multi-cloud platforms such as VMware vCenter, AWS Cloud, and bare metal servers with minimal operations to improve O&M management efficiency.

6.1.4.1 VMware Clusters

SCP supports unified management of VMware vCenter through its APIs without affecting customer use.

After SCP manages VMware vCenter, users can perform bulk startup and shutdown, restart, backup and recovery, and migration on VMware VMs. SCP provides VMware VMs with the same operations and experience as HCI VMs, such as full lifecycle management that includes creating, editing, and deleting, as well as work order approvals, metering and billing, and multi-tenant management.

VMware VMs can also be bidirectionally migrated between VMware vCenter and Sangfor Cloud and recovered from backups to Sangfor Cloud. VMware VMs that are migrated or recovered from backups to Sangfor Cloud are automatically installed with vmTools, which can improve O&M efficiency.

6.1.4.2 Physical Hosts

SCP supports unified monitoring and management of physical hosts and data

center servers, which has the following two parts:

- Managing host hardware, including component alert status, and power-on and power-off.
- Monitoring host system performance, including CPU, memory, and network status.

Physical host hardware management relies heavily on the Intelligent Platform Management Interface (IPMI), an open standard hardware management interface specification that defines a specific method for embedded management subsystem communication. IPMI messages are communicated through the baseboard management controller (BMC, located on a hardware component of the IPMI specification).

Intelligent management of low-level hardware using IPMI has the following advantages over management using the OS:

1. Supports managing out-of-band servers, such as power-on, power-off, and restart.
2. Eliminates the need for the OS to monitor the status of the transmission system and supports monitoring the physical health of servers, such as CPU, memory, power supply, and fan status.
3. Supports obtaining and viewing alert information of hardware configurations of physical servers.

UI Operations and Displays:

SCP supports adding a single host and bulk adding hosts by importing Excel files, powering them on and off, tagging, exporting, allocating and deallocating, and monitoring them, moving them to other groups, accessing their consoles, editing their attributes, and sending their alerts.

Physical Machines > Add Physical Machine

IPMI Address:

Port:

Username:

Password:

Specifications

IP Address:

OS:

CPU: core(s)

Memory: GB

Asset Type:

Server Room:

Server Rack:

Vendor:

Figure 26: Physical Host Specifications

The alert information of hardware configurations of physical servers is obtained through IPMI.

Tenants can power on and off and tag hosts, moving them to other groups, and editing their attributes; tenants users can power on and off hosts and edit their attributes.

6.1.4.3 AWS Cloud

SCP acts as a hybrid cloud management platform that supports unified management of local private cloud platforms and public AWS Cloud environments.

Application Scenarios:

1. Add an AWS Cloud region on SCP to view and manage VMs in that region locally.
2. Create a VM in a region of AWS Cloud on SCP and remotely log in to its console to quickly deploy and put business services online.

UI Operations and Displays:

Environment Creation: Supports creating an AWS Cloud environment with the

access key ID and access key secret of an AWS Cloud account, and configuring the resource update interval.

Resource Management: Supports creating, powering on, powering off, and releasing elastic compute servers (ECSs) of AWS Cloud.




Figure 27: AWS Cloud Resource Management

6.1.5 Quotas

SCP manages the number of clusters and physical CPU cores and the permissions of each module based on the quota license key. The number of host CPU cores licensed by SCP cannot be less than the number of CPU cores in all managed clusters. Additionally, SCP supports licensing and quota management for security resources and services.

6.1.6 Application Center

As business models and new technologies grow, the evolution of enterprise architectures to microservice architectures creates changes in development requirements. As shown in the figure below, normal application development requires eight steps, whereas PaaS-based development in a microservices architecture requires only two steps. Steps that are grayed out are automatically performed by the platform so that developers can focus on the application requirements.




Figure 28: Application Development Steps

Sangfor **Application Center** provides databases, middleware, big data components, and public components of third-party commercial applications and standard development frameworks for customers' development and production environments to meet the needs of developers for quick development and coming online of applications; it also provides monitoring and alerts as well as log traceability for easy fault location and fixing and simplified O&M. Sangfor works with vendors of independent software, applications, and middleware to build a comprehensive, free, and diversified cloud application ecosystem.

The **Application Center** uses a domain-specific language (DSL) to define the configurations of an application package, including the deployment process, resources and scripts, UI, and basic internal information. The application package structure is shown below. This script defines all the information to quickly publish various applications.




Figure 29: Application Package Structure

Type	Name	Version	Deployment Mode
Database	MySQL	5.7	Single instance, active-standby, and cluster
Database	MongoDB	3.4.24	Single instance and cluster
Database	Redis	5.0.12	Single instance, sentinel, and cluster
Middleware	Kafka	2.7.0	Single instance and cluster
Message Queue	RabbitMQ	3.7.28	Single instance and cluster

Table 2: Application Types Supported by Application Center




Figure 30: Application List

Name	Category	Provider	Tag	User	Type	Time Created	Last Updated	Operation
Kafka_2.7.0...	Database and ...	-	kafka_im...	-	Public	2024-02-20 15:...	2024-02-20 15:...	Edit Delete
MongoDB_3...	Database and ...	-	mongodb...	-	Public	2024-02-20 15:...	2024-02-20 15:...	Edit Delete
RabbitMQ_3...	Database and ...	-	rabbitmq...	-	Public	2024-02-20 15:...	2024-02-20 15:...	Edit Delete
Redis_5.0.12...	Database and ...	-	redis_ima...	-	Public	2024-02-20 15:...	2024-02-20 15:...	Edit Delete

Figure 31: Application Package Management

6.1.7 Homogeneous Hybrid Cloud

Previous hybrid cloud solutions only implemented resource pool management through SCP, without data plane and network plane connections or consistency of security capabilities and service directories. Sangfor provides Sangfor O2O Synergized Cloud, which is a safe and reliable lightweight hybrid cloud architecture with a smooth experience, unified resource management, and in-depth capability integration, based on offline private cloud and online managed cloud in completely consistent architectures.

SCP achieves unified management and a one-page view of resources and business services of the local and online resource pools.

6.2 System Management

The system management acts as a basic module of SCP, providing management of policies, users, platform security, system, and certificates.

6.2.1 Users

The user management module is responsible for user access authentication and licensing management of SCP, which employs role-based access control (RBAC) for authentication at the resource-level granularity. In RBAC,

permissions are associated with roles, and users can be granted corresponding permissions through roles and retrieved from roles to retrieve permissions, which greatly simplifies the management of user permissions. A user can have multiple roles and a role may have multiple permissions.

6.2.2 Policies

The policy management module is responsible for the management of VM policies on SCP, including lifecycle policies, scheduling policies, and auto scaling policies. For example, it supports setting the run location of VMs according to a scheduling policy. The platform can also automatically select an optimal host for a VM to run on based on the remaining resources and the resources required by the VM.

6.2.3 Certificates

The certificate management module is responsible for SSL certificate configuration and management. By configuring an SSL certificate to create an SSL tunnel between the client and SCP, the client can be trusted by the browser to prevent HTTPS sites from being marked with a danger warning. SCP supports the manual import of trusted certificates.

6.2.4 Platform Security

SCP provides rich features to secure the platform, including login protection, authentication settings, SSO settings, key pairs, and SM encryption mode.

6.2.4.1 Login Protection

SCP supports login configurations such as password policy, session timeout limit, and authentication for all platform accounts (including administrators, tenants, tenant users, and other role accounts).

SCP supports two-factor authentication login with a password plus SMS authentication or OTP authentication for enhanced account security.

6.2.4.2 Authentication

SCP supports integration with lightweight directory access protocol (LDAP) servers to import LDAP users into the platform and grant them roles to host their password authentication.

SCP based on HCI 6.9.0 and higher supports integration with third-party OTP servers through the standard Remote Authentication Dial-In User Service (RADIUS) protocol. OTP code is required additionally for login for account security. The RADIUS protocol is used for network access control and the RADIUS server receives identity information of users attempting to access the network from the network devices for authentication purposes.

6.2.4.3 SSO

SCP supports single sign-on (SSO), a feature that enables users to gain access to associated systems and applications in the SSO system after logging in to the identity authentication server.

SCP supports integration with customers' CAS systems and granting roles to their users who can then use the account information for SSO to reduce the customers' O&M workload.

6.2.4.4 Key Pairs

VM login using the account password has weak security and may involve the risk of brute-force attacks. SCP provides the VM key pair feature, a secure and convenient login authentication method. A key pair contains a public key and a private key. The public key is sent to the bound VM, to which the user can log in remotely through SSH with the private key.

6.2.4.5 SM Encryption Mode

SCP supports encrypting the platform using an SM encryption algorithm to meet customers' business system needs.

1. Supports enabling SM encryption mode according to the configured HSM.

2. Supports encrypting sensitive data on the platform using the SM encryption algorithm SM4.
3. Supports verifying the integrity of important data using the SM encryption algorithm SM3-HMAC.

6.2.5 System

The system management module contains service management, date and time management, platform configuration, data center settings, SMS and email settings, and component management.

6.2.5.1 Services

SCP supports activating and stopping the following services:

1. API: An external interface service that allows custom-made software development and integration with third-party cloud platforms.
2. Billing Center: A user resource billing and metering service that includes the report of total costs, itemized bills for tenants, and pricing.
3. Application Center: Provides full lifecycle management of conventional and cloud-native application services from quick deployment and configuration to monitoring and backup.
4. aOPS: A unified O&M product system based on cloud business services.
5. KMS Encryption: Encrypts VM data, network and security resources using the AES-256 algorithm to avoid data theft.
6. Advanced Services: Third-party and value-added services such as aSecurity (aSEC), Network Insight (aNI), aDesk VDI, Database Management Platform (DMP), object storage, and Cloud Security Service Platform (CSSP), etc.

6.2.5.2 Date and Time

SCP supports customizing the cluster time and enabling time synchronization to ensure time accuracy.

6.2.5.3 Platform

SCP supports changing the platform IP address and configuring DNS servers and routing as needed to meet the network requirements in different scenarios.

6.2.5.4 Data Center

SCP supports managed cloud service account access by configuring the information including local data center URL, tenant login URL, etc.

6.2.5.5 SMS and Email

SCP supports configuring an SMS server for login through SMS two-factor authentication as well as an email address for alert notifications to keep customers informed of the cluster situation.

6.2.5.6 Components

SCP supports displaying or hiding VMs for the following components: VDC, aSEC, aNI, aOPS, etc.

6.2.6 Licensing

SCP supports unified licensing for all product components on the platform (SCP, HCI clusters, VNF devices, etc.) and licensing for an individual HCI cluster after managing multiple HCI clusters for easy user management and subsequent flexible adjustment.

6.2.6.1 Unified SCP Licensing

Unified licensing requires a USB key and a license key number. The USB key is a physical device that needs to be plugged into a physical host running an HCI cluster for activation so that a unique key ID can be mapped to SCP. Then the customer imports the license key file containing licensing information for all Sangfor Cloud products purchased to SCP to activate the licenses for the

relevant products.



The key ID and license key number must match each other; otherwise, the activation will fail.

The process of SCP unified licensing is shown in the figure below. The customer first plugs a USB key into a physical host running an HCI cluster to map the key ID to SCP and then imports the corresponding license key file to SCP to obtain the licenses of the corresponding products and allocate them.




Figure 32: Unified SCP Licensing

After importing the license key file on SCP, the customer's product licenses, including information such as the number of licenses and their expiration dates, will be displayed on the **Resources > Management > System > Licensing** page.

6.2.6.2 Individual HCI Licensing

SCP supports licensing an HCI cluster individually to flexibly meet customer needs. An individually licensed HCI cluster operates independently but is still functionally under the unified management of SCP.




Figure 33: Individual HCI Licensing

6.2.6.3 Licensing for SCP

SCP will be automatically activated if the license key file imported on it includes its license.

6.2.6.4 Licensing for HCI Clusters

SCP supports unified licensing for HCI clusters so that customers can conveniently split and merge licenses and use idle licenses to flexibly meet business changes. License configuration does not require a new license key number.

The license key file contains the number information of HCI cluster licenses, which will form a license pool. Users can allocate licenses of up to the number of licenses in the pool to any cluster (some clusters may require a minimum number of licenses), or reclaim licenses from a cluster and reallocate them to other clusters. It is very flexible as the license information does not need to be bound to a cluster.

6.2.6.5 Licensing for Security Components and Advanced Components

The license key file contains license information for security components such as Network Secure (vNGAF), Application Delivery Controller (vADC), and Internet Access Gateway (vIAG), and advanced components such as Continuous Data Protection (CDP), and Heterogeneous Virtualization Mgmt

(aHM). The license information includes the number of licenses and their expiration dates. Users can create up to the licensed number of virtual devices for security components and advanced feature components.

6.2.6.6 vKEY Licensing

SCP 6.8.0 and higher support licensing with vKEY software to prevent cluster services from being affected by missing or failures of hardware USB keys. It is more efficient as the USB key does not need to be sent back or reapplied for license updates.

6.2.6.7 Grant License Based on CPU Cores

For versions earlier than SCP 6.11.2, licensing for aSV, aNET, and aSEC was uniformly priced based on the number of physical cores. For SCP 6.11.2 and later versions, licensing for aSV, aNET, and aSEC has been adjusted to a tier-based pricing rule based on the number of CPU cores.

The pricing rules for aSV, aNET, and aSEC licenses based on the number of CPU cores are consistent and are divided into three tiers:

- (1) Specification for Pro CPUs: Applicable to all numbers of physical CPUs.
- (2) Specification for Plus CPUs: Applicable to physical CPUs with fewer than 64 cores.
- (3) Specification for Base CPUs: Applicable to physical CPUs with up to 32 cores.

When licensing an HCI cluster, you can grant one or more licenses with different specifications to aSV and aNET at the same time. However, the specifications and number of licenses must match the actual number of physical CPUs and cores in the HCI cluster; otherwise, the licensing will fail. For example, if an HCI cluster contains 6 Base CPUs, 6 licenses of Base CPUs are required. If an HCI cluster contains 6 Base CPUs and 4 Plus CPUs, 6 licenses of Base CPUs and 4 licenses of Plus CPUs are required. In addition, 10 licenses of Plus CPUs or 10 licenses of Pro CPUs can also be used because higher licenses can activate CPUs with lower licenses, but not vice versa.

Note:

SCP 6.11.2 and later versions support importing license key files in earlier versions. The licenses of CPUs for aSV, aNET, and aSEC in earlier versions will be automatically identified as licenses of Pro CPUs in SCP 6.11.2 and later versions after import.

7 IaaS-Based Services

7.1 Compute

7.1.1 VM

7.1.1.1 Introduction

A VM is an easily accessible, elastically scalable, and on-demand virtual cloud computing server consisting of virtual CPU, memory, and disks. It combines services such as virtual private cloud and cloud backup to create an efficient, reliable, and secure computing environment for continuous and stable service running.

The VM services of SCP are based on the underlying compute virtualization (aSV) technology of HCI, which provides users with a wizard-like deployment for quick business service configuration including image, resource pool, cloud storage, and network. Multiple VMs can be managed in groups. Tenants can label VMs running the same business service in different resource pools with the same tag for grouping across resource pools.

When creating an elastic server, users can set its virtual CPU, memory, image type, and login authentication method as needed. After creation, computing and storage resources are used in the same way as local PCs and physical servers on demand, greatly reducing hardware costs. Meanwhile, the unified management platform simplifies maintenance and management, allowing users to focus on business services.

7.1.1.2 Application Scenarios

The VM services provide the base operating system for the following broad business deployment scenarios:

1. Simple applications and portals such as blogs and official websites, which do not require servers with high computing and storage performance. It is sufficient to use an ordinary VM in this scenario.
2. Databases and other applications that perform fast data exchange and processing, which require servers with high I/O capabilities. It is recommended to create a resource pool using high-performance NVMe SSDs to provide high read/write performance and low latency in this scenario.
3. Multimedia, video production, and image processing, which require servers with good image processing capabilities. It is recommended to use VMs with GPU image computing capabilities in this scenario.
4. AI analysis, machine learning, and deep learning, which require GPU as the accelerator to implement near-real-time and highly concurrent massive computing. It is recommended to use VMs with high GPU computing capability in this scenario.

7.1.1.3 Supported Features

1. Multi-Type Support: Supports multiple types of VM, including x86, ARM, GPU, VMware, and AWS.
2. Full Lifecycle Management: Supports powering on and off, suspending VMs, resetting their passwords, and setting their expiration dates, as well as deallocating expired VMs according to predefined policies for efficient resource management.
3. Group Management: Supports managing VMs in groups across resource pools, facilitating users to manage a large number of VMs.
4. Data Protection: Supports performing operations such as snapshots, backups, and CDP on VMs.

7.1.2 Bare Metal

7.1.2.1 Introduction

SCP supports unified management and monitoring of physical hosts as bare

metal services. The management involves host hardware, including component alert status, and power-on and power-off, and the monitoring involves host system performance, including CPU, memory, and network status.

Physical host hardware management relies heavily on the Intelligent Platform Management Interface (IPMI), an open standard hardware management interface specification that defines a specific method for embedded management subsystem communication. IPMI messages are communicated through the baseboard management controller (BMC, located on a hardware component of the IPMI specification).

7.1.2.2 Application Scenarios

To manage and monitor physical servers on cloud platforms, users can use SCP to manage physical hosts to build bare metal services for O&M.

7.1.2.3 Supported Features




Figure 34: Bare Metal Service

1. Add a single host and bulk add hosts by importing Excel files.
2. Power on and off, restart, tag, and monitor hosts, move them to other groups, and send their alerts.
3. Eliminate the need for the OS to monitor the status of the transmission system and use IPMI to monitor the physical health of servers, such as temperature, voltage, fan status, power status, etc.

4. Monitor host systems through protocols such as SNMP, SSH, and WMI with the help of Sangfor Cloud's Monitor Center (aMC).

7.1.3 GPU-Enabled VM

7.1.3.1 Introduction

SCP supports the management of various enterprise-grade and customer-grade GPUs, such as NVIDIA's Tesla A100, Quadro RTX8000, GeForce RTX3090, etc.

SCP supports splitting enterprise-grade GPU cores into vGPUs through the GRID driver and license controller of NVIDIA. vGPUs are then allocated with video memory resources and core computing capacity of physical GPUs through three schedulers: preemption, equal share, or fixed share to support compute acceleration, graphic computing, and machine learning on VMs. The customer-grade GPUs only support the passthrough mode, in which the graphics card can be directly allocated to VMs.



The models of GPUs running on the same physical host must be the same; otherwise, the GPU features will be unavailable.




Figure 35: GPU Services

7.1.3.2 Supported Features

1. GPU Passthrough and Splitting

GPU passthrough refers to a technology that passes through a single physical GPU core on a graphics card to a VM through DMA and Virtual Function I/O (VFIO).

The DMA technology allocates a segment of address space from the kernel of the HCI system to the hardware device (graphics card). When the applications at the upper layers determine to use this device, they will directly put the data into the corresponding address space for its processing, eliminating the consumption of CPU transit and scheduling.

The VFIO technology is a protection mechanism for I/O scheduling that guarantees the reliability of I/O transfers. The KVM host identifies the physical cores of the graphics card and maps them in a one-to-one mode to VMs as virtual hardware, which is then used by guest OSes.




Figure 36: GPU Splitting and Passthrough

GPU splitting is similar to GPU passthrough, but the graphics card is taken over at the KVM host layer by the NVIDIA vGPU host driver, which can provide management such as status monitoring and resource splitting for data center-grade GPUs (generally the Tesla and Quadro series). The vGPU host driver

splits the GPU into multiple vGPUs, each of which can be allocated to a single VM. For example, a single Tesla P40 GPU core can be virtualized into up to 24 vGPUs to support 24 users at the same time. GPU cores can also be virtualized into different sizes of vGPUs based on scenarios and loads, and GPU-enabled VMs can be migrated between GPU servers, which are all outstanding features of the new-generation graphics card virtualization. Meanwhile, NVIDIA has increased the price of software licenses for the new hardware virtualization technology.

2. GPU-Enabled VM Scheduling

SCP provides two scheduling modes for GPU computing resources:

- **Best Performance:** Scheduling GPU-enabled VMs to physical hosts with idle GPUs, so that they can evenly use different graphics cards to ensure the best performance of each VM in scenarios that require high GPU performance, such as deep learning, video rendering, and scientific computing, etc.
- **GPU Consolidation:** Preferentially scheduling GPU-enabled VMs to physical hosts that are running GPU-enabled VMs to make full use of the resources in scenarios such as AI college training in the education industry and general workstations.

Users can select the GPU-enabled VM scheduling mode based on actual application scenarios and business reliability requirements, and **GPU consolidation** is selected by default.

3. GPU-Enabled VM Creation and Quota Management

SCP supports the creation of GPU-enabled VMs using resource pools with GPUs. Users can configure a GPU core or slice it into vGPUs for those VMs. Different operating systems support different GPU models and require the graphics card drivers of the versions supported by SCP for normal GPU use. For supported GPU models and ways to obtain graphics card drivers, users can refer to the configuration page tips.

For the management and allocation of tenant GPU resources, SCP provides two types of quotas: vGPU and graphics card model. vGPU quota limits the number

of vGPUs and graphics card model quota limits the number of graphics cards of a certain model that a tenant can configure. In GPU splitting mode, the graphics card model quota is consumed according to the percentage of splitting. For example, a GPU-enabled VM configured with a split A100-8C vGPU (8 GB of video memory, 0.2 A100 core computing) consumes one vGPU quota and 0.2 Tesla A100 graphics card model quota.

7.1.4 Image

7.1.4.1 Introduction

An image is an elastic server template that contains the software and necessary configurations. It includes the operating system, preinstalled public applications, and the user's private applications or business service data.

SCP provides simple and convenient image self-management. Users can flexibly apply for VM services by using public images, private images, or network device images, and create private images through VM services or external image files.

- Public Image

A public image contains common standard OSes and preinstalled public applications and offers simple and convenient image self-management, viewable to all users.

- Private Image

Private images involve tenant images that are viewable to the tenants' users and cloud platform images that are only viewable to the platform administrators. A private tenant image is created by the tenant based on a VM and includes the OS, preinstalled public applications, and the tenant's private applications. Creating a cloud server from a private image saves time on reconfiguration.

- Network Device Image

SCP provides a number of virtual network devices such as vNGAF, vIAG, and vADC. Platform administrators can upload VM images containing these

network device services to create and configure them for use in the network topology.

During the SCP O&M management process, users can create and upload images in the following ways:

1. Create a public image from an existing VM

Platform administrators can create an image using a VM that has been configured on demand as a template. This method requires the VM to be powered off with the OS and vmTools installed and creates an image that retains all software and data on the VM system disks. If the image is created in a resource pool higher than HCI 6.8.0, it will retain the data on all disks of the VM.

2. Upload a local image file

Platform administrators can create a new VM by uploading a local image file on SCP, which currently includes ISO files for OSes (not viewable to tenants) and VMA files provided by Sangfor or other vendors.

7.1.4.2 Supported Features

The administrator and tenant can upload, edit, and distribute an image.

The process of image upload and distribution is described in the figure below. Firstly, the administrator uploads images on the SCP platform and then bulk distributes them to resource pools. The SCP platform only stores some image information (such as operating system types and file sizes of images) rather than the image files. The image files are distributed to and stored on the shared datastore of clusters within the selected resource pools, with each cluster maintaining an independent copy. The distribution process allows the administrator to manually specify the distribution source and the scope to be distributed. After that, the image will be distributed from the source to the selected resource pools. The distribution method is implemented using cross-cluster VM migration capability on HCI. Breakpoint resumption is not supported. If an image upload process is interrupted, the next upload will restart from the beginning.

The administrator can edit the images at any time after they are uploaded, for example, change the image name, image description, and operating system name, and add or delete resource pools. The administrator can also add or delete resource pools. If a resource pool is added, SCP distributes images to this resource pool from other resource pools that have the images, which requires connectivity between the resource pools in the HCI management network. **By default, distribution speed is limited to 90 MB/s for 1GE interfaces and 250 MB/s for 10GE interfaces, which can be configured in the migration task.** The image will become unavailable in deleted resource pools.

The super administrator and tenants can upload, edit, and distribute images.

The process of uploading and distributing images is shown in the following figure. First, the administrator uploads images on SCP and bulk distributes them to resource pools. SCP stores only some information of the images (e.g., OS and file size) and does not save the image files. The image files are then distributed to the shared storage of the corresponding clusters based on the resource pools selected by users and are copied individually by each cluster.

The administrator can edit the images at any time with information such as their name, description, OS name, and resource pools to which they belong, and add or delete them. When adding an image to a resource pool, SCP distributes the image from one of the other resource pools that already have it to that resource pool (the HCI management network between the two resource pools must be connected). Images deleted from a resource pool will be unavailable in it.




Figure 37: Image Management Service

The administrator can clone existing images. During the cloning process, the configuration for the cloned images (such as CPU, memory, disk capacity, and advanced options) can be modified. The resource pool to which a cloned image will be distributed can also be configured, following the same distribution logic as uploading image.

The administrator can update existing images. During the update process, the image configuration (such as CPU, memory, disk capacity, and advanced options) can be modified, and the system file can be updated by accessing the console of the image operating system.

7.1.5 Auto Scaling

7.1.5.1 Introduction

Auto scaling (AS) refers to the automatic adjustment of business resources through policies based on needs. Users can customize AS configurations and policies to reduce manual resource adjustments to respond to business changes and peak workloads in order to save resources and labor costs.

7.1.5.2 Supported Features




Figure 38: Auto Scaling

SCP supports the auto scaling actions within a VPC network with the vADC component.

1. Create an AS group for the tenant using a template and associate it with an AS policy.
 - **AS Group:** A collection of instances running stateless business services with the same application scenario, which is the basic unit for enabling and disabling AS policies and performing AS actions.
 - **AS Policy:** The time (scheduled tasks) and resource (alert tasks) trigger thresholds configured for AS groups, such as scaling, removal, and alert thresholds.
2. The system obtains the AS group configuration and continuously monitors the business processes by metrics.
3. When the scaling threshold is triggered, the system process executes the corresponding AS action based on the AS rules.
4. For example, scale up nodes by starting a new VM through the configured VM template from the AS group.
5. After the new node is started, its information is added to the node pool of the load balancer and it becomes a business node.
6. A large number of subsequent business connections will be scheduled to the new node by the load balancer with a built-in cooldown period (during which the AS action will not be triggered again to avoid excessive resource consumption).

7.1.6 aDesk VDI

7.1.6.1 Introduction

Sangfor aDesk VDI is a desktop as a service (DaaS) from Sangfor Cloud that delivers complete Windows virtual desktops to end users through the internet.

This turnkey service can be easily scaled up or down on demand and simplifies desktop IT management tasks.

With virtualization technology, aDesk VDI hosts and manages all desktops, applications, and data in data centers, which are originally running and stored on local traditional PCs. With hyper-convergence technology, Sangfor VDS creates VMs based on CPU, memory, and storage resources according to users' needs and delivers the VMs to endpoints through the desktop delivery protocol, providing a virtual desktop environment similar to the local desktop.

7.1.6.2 Supported Features

1. A Sangfor Virtual Desktop Controller (VDC) is used as a workspace on SCP and manages an aDesk VDI cluster by administrators.
2. Create workspaces by connecting to the OpenAPIs of the VDCs to integrate with their service ports and connection keys.
3. Grant tenants aDesk VDI permissions by associating their accounts with workspaces.
4. Create virtual desktops in persistent or nonpersistent mode in floating pools for tenants in corresponding workspaces, and perform operations such as password recovery, configuration modification, user association, and IP address specification on them. Tenants and tenant users can use the created desktop resources on their endpoints through the network.

7.1.7 VM Customization Specifications

7.1.7.1 Introduction

VM customization specifications are predefined templates for standardized deployment of system parameters. You can apply a customization specification when creating or cloning VMs to automatically execute predefined procedures, such as initial configuration and environment verification, during the VM's first startup, enhancing deployment efficiency. The implementation principle is shown below:

During the VM's first startup, vmTools establishes a communication channel with the node where the VM runs. After the auto update and network configuration are complete, SCP converts the customization specification into executable commands through the channel for injection and execution.




Figure 7.1-1 Custom Specification Implementation Process

7.1.7.2 Application Scenarios

Used for large-scale deployment of VMs, enabling bulk initialization of guest OSes, such as resetting the guest OS SID, configuring networks, and installing required software packages, which significantly simplifies the business environment deployment.

7.1.7.3 Supported Features

- Supports resetting SID and adding users to working groups and domains on Windows.
- Supports executing custom Sysprep answer files on Windows.
- Supports executing custom BAT scripts on Windows.
- Supports executing custom scripts on Linux.
- The customization specifications support import and export to facilitate the specification migration across multiple SCPs.
- The customization specifications can be executed in the following scenarios: 1. Creating VMs from built-in images. 2. Cloning VMs. 3. Creating VM through auto scaling.

7.2 Network

7.2.1 VPC

7.2.1.1 Introduction

Virtual private clouds (VPCs) are collections of private resources for tenants that are logically isolated from each other at layer 2. The VPC features provided by SCP enable customers to locally separate VPC networks and use network services such as elastic IP addresses, direct connections, object storage, security services, and backups according to their business needs.

Private networks can be categorized into the following types based on the

traffic type:

1. Internal Network: Networks through which east-west traffic passes between VMs in the same or different subnets within a VPC. VMs communicate with each other through allocated internal IP addresses. East-west traffic does not go through the firewall by default.
2. External Network: Networks through which north-south traffic passes when VMs within a VPC access the external internet or provide external network services. VMs communicate with the external internet through allocated elastic IP addresses. When the cloud firewall service is enabled, all north-south traffic in the internal and external networks will go through the firewall for cleaning.
3. Direct Connection: Used for layer-2 network communication between VMs in a VPC and local data centers of a tenant or devices managed in a server room. Through the direct connection, the VPC can reside on the same network as on-premises devices.
4. Shared Service Network: Provides connectivity for tenant access to shared services such as Windows authorization, databases, and object storage.




Figure 39: VPC Network Topology

All four networks support individual gateway and interface multiplexing. When there is a multiplexed physical interface, all traffic is isolated through a virtual local area network (VLAN).

7.2.1.2 Supported Features

1. Multi-Line and Multi-Gateway

The administrator can centrally plan multiple VPC gateways for each independent resource pool and communication domain. Users can select lines and gateways for traffic forwarding as needed when creating an elastic IP address.




Figure 40: Multi-Line and Multi-Gateway

After SCP is deployed, the administrator needs to create an edge in each resource pool or communication domain, and it must be connected to the physical interface from which the traffic is expected to egress so that the traffic of the created virtual private network goes out from the selected gateway. When different traffic needs to multiplex an edge, VPC gateway traffic can be isolated from others by configuring a VLAN ID.

2. Multi-VPC for Tenants

Tenants can create multiple VPC networks within a resource pool. VPC networks are virtual private networks that are logically isolated from each other and can be divided into multiple subnets and reuse the same network segments.




Figure 41: Multi-VPC for Tenants

3. DVS

Each subnet corresponds to a distributed virtual switch (DVS) that isolates traffic using the VXLAN protocol at the backend. Each DVS is allocated with a VXLAN ID unique to the data center. VPC networks are isolated from each other, in which VMs need to be configured with an elastic IP address and a VPC gateway to communicate with the external network.

4. Router

A VPC network consists of two routers and multiple DVSs, where the internal router is used to forward east-west traffic within the VPC, and the external router, where the elastic IP addresses and routing rules are in effect, acts as a VPC gateway to translate and forward north-south traffic.




Figure 42: Virtual Router

Routers in a VPC provide basic routing features, including configuring ACL rules, static routes, DNS servers, DHCP services, and DNAT. They can also function as gateways, enabling the VPC to access the external network and the shared service network, and enabling direct connections. A router can be associated with multiple elastic IP addresses, one of which is the primary.

5. DNS Server and DHCP

Routers in a VPC create a pool of DHCP servers for each subnet, which contains the default gateway, domain name system (DNS) server address (same as the default gateway), and IP-MAC address binding entries to be sent to VMs. When a VM starts up, it initiates a Dynamic Host Configuration Protocol (DHCP) request, which is received and processed by a router. The DHCP server then queries the MAC address in the request to check whether there is a corresponding IP-MAC address binding entry, and if there is, it will return the IP address, the default gateway, and the DNS server address to the VM.

6. DNAT

Tenants can allocate elastic IP addresses to routers to configure destination network address translation (DNAT) rules on them for external services for

elastic IP address multiplexing and security control. A single service on a VM can be configured with multiple DNAT rules.

7.2.2 EIP

7.2.2.1 Introduction

An elastic IP address (EIP) is used for the external network gateway of a cloud resource and can be associated with a VM, a router in a VPC, and a VNF device. A VM associated with an EIP can access the external network and provide external services. A vADC associated with an EIP can receive access requests from the public network and distribute them to VMs.

A router can be associated with multiple EIPs belonging to the same or different lines (e.g., China Unicom and China Telecom). When a router is associated with EIPs of different lines, the traffic goes out of different gateways. Associating an EIP to a router actually distributes an IP address to the outbound interface of the router (which has no IP address configured by default) to associate it with a VNF device or VM. The primary EIP is the outbound interface IP address by default, and the others are standby, all of which can be used to configure DNAT rules.

To associate an EIP with a VM, the EIP is first configured as the standby IP address for the router's outbound interface. Then, a secure network address translation (SNAT) targeting all ports from the internal network IP address of the VM to the EIP and a DNAT targeting all ports from the EIP to the internal network IP address of the VM are sent. The EIP information cannot be viewed within the VM.




Figure 43: Elastic IP Address

7.2.2.2 Elastic IP Pool

An elastic IP pool is used to assemble and manage EIPs for specific lines on SCP. Super administrators plan elastic IP pools from which EIPs are allocated, giving great flexibility to tenants and reducing security risks for platform administrators.




Figure 44: Elastic IP Pool

As shown in the figure above, an elastic IP pool is directly associated with a network gateway, while a user is directly associated with a VPC network. Administrators need to create an edge and then a gateway for a created resource pool, and specify a gateway when creating a VPC network or an elastic IP pool for a tenant. When creating an elastic IP pool, it is necessary to specify:

1. whether it is globally available, and if not, specify its associated gateways.
2. Associated tenants.
3. Line type.

For flexibility, administrators can specify whether an elastic IP pool is globally available or available to specific gateways at the time it is used.

To access the external network, tenants need to apply for an elastic IP address and associate it with the VM. Through the VPC network specified by a tenant to use the EIP, administrators can locate available elastic IP pools corresponding

to the gateways associated with the VPC network to allocate an EIP.



Elastic IP pools should be selected by priority, with priority given to elastic IP pools associated with the specified gateways, then to partially available elastic IP pools, and lastly to globally available elastic IP pools.

To avoid the trouble of a failure of a single elastic IP pool, platform administrators can purchase elastic IP pools from multiple internet service providers (ISPs) and associate them with HCI resource pools on SCP. An elastic IP pool can be associated with multiple resource pools and vice versa. Elastic IP pools associated with a resource pool can be viewed in all VPC networks associated with that resource pool.

7.2.3 Shared Service Network

7.2.3.1 Introduction

A shared service network is a network channel for tenant VMs in a VPC to access shared services on SCP such as file storage, object storage, FTP, and Windows authorization as needed.

The basic principle of the shared service network is to allocate an EIP for the outbound interface of the router to access the shared service network only, and then send a port translation rule to translate the source addresses of all traffic accessing this network to the outbound interface IP address, and at the same time, send a policy-based route to forward the shared service network traffic to the shared service gateway, enabling the access from a VPC network to physical shared services.




Figure 45: Shared Service Network Topology

7.2.3.2 Network Flow Direction




Figure 46: Network Flow Direction

1. Access to shared services through IP address: The VM first initiates a request to access a shared service by sending a packet to the internal router, which

forwards the packet to the shared service router based on the built-in policy-based route. The shared service router changes the source IP address of the packet to the associated EIP (allocated from the elastic IP pool of the shared service line type) through NAT, and then forwards the packet to the corresponding shared service gateway to access the shared service through the planned edge.

2. Access to shared services through domain name: If there is no domain name in the local DNS cache, a VM in a tenant VPC subnet will send a domain name resolution request to the DNS server when accessing a shared service through a domain name. The internal router then responds to this request through the DNS proxy and returns the corresponding IP address. The VM fills the packet with that IP address as the destination address and sends it out through the same process as access through an IP address.

7.2.4 Network Orchestration

7.2.4.1 Classic Network

SCP provides virtualized network orchestration. Platform administrators can drag and drop all basic network devices (routers, switches), VMs, and VNF devices of tenants for management in the global view.




Figure 47: Classic Network Topology

7.2.4.2 Tenant Network

Tenants can use existing classic networks or VPC networks, which greatly simplify the deployment of business services with a fixed orchestration schema.



Figure 48: Tenant Network Configuration

The VPC network feature on SCP allows users to quickly deploy and use private networks without customizing their internal configurations. To customize tenant network configurations, SCP provides administrators with the feature of quickly selecting the original networks of tenants for viewing and configuration.

Administrators can view the selected original networks of tenants on the Topology page and modify the connections and configurations in the topology. This feature is mainly used for the configuration of tenants' direct connections, but also for troubleshooting in case of network disconnections (configuration changes may cause disconnections within tenants' VPCs. Please operate with caution).




Figure 49: Tenant Network Topology

7.2.5 Communication Domain

A communication domain is a network domain of multiple resource pools that are connected together through the peer connection. In a communication domain, the virtual devices (eg., VMs) in different resource pools can communicate with each other, which facilitates effective management and delivery of network resources in large-scale data centers.

Compute resources (VMs, VNF devices) are isolated between resource pools within a communication domain, and communication domains and resource pools correspond to network resources and compute resources respectively. In order to implement the capability of large clusters, the concept of region is introduced. Users can perceive the concept of three layers: region, communication domain, and resource pool.




Figure 50: Communication Domain

7.2.5.1 Introduction

1. Resource Pool Roles

Resource pools in a communication domain play three roles: primary controller, standby controller, and standby resource pool. The primary controller and standby controller are specified by users when creating the communication domain. The primary controller runs the network management plane services of VN to manage the network resources and configurations within the communication domain. The standby controller backs up the virtual network configuration database of VN and switches to the primary role when the primary controller is abnormal to ensure the HA of services. All the rest of the resource pools are standby resource pools.

2. Controller IP Address

The resource pool where the primary controller resides is required to be configured with four controller IP addresses on the same network segment as the management interface, three of which are allocated to the controller node where the primary controller resides for load balancing of services within the communication domain, and one for network resource management to fail over between controller nodes. The primary controller and standby controller guarantee high reliability of network resource data. When the primary controller fails, the standby controller can be manually switched to a new primary controller to continue providing services. The controller IP addresses of the standby controller can be configured as needed and only used when the standby controller is set as the primary controller.




Figure 51: Resource Pool Roles in Communication Domain

Prerequisites and precautions for adding resource pools to a communication domain:

1. Only resource pools connected to the network can be added.
2. Only resource pools of HCI 6.7.0 and higher can be added.
3. Only resource pools that meet the limits for communication domain size and data merge conflicts can be added.
4. Up to 128 resource pools can be added to a communication domain, and the number of online physical hosts clustered in each resource pool must be between 3 and 32.
5. The network configurations associated with resource pools must automatically convert to associate with the communication domain after resource pools are added to it.
6. It is prohibited to perform operations on the IP pool for VXLAN on HCI after resource pools are added to a communication domain.

Prerequisites and precautions for removing resource pools from a communication domain:

1. Only resource pools that have no VMs, nor connected VMs or VNF devices can be removed.
2. The network configurations of resource pools will be cleared by the communication domain after resource pools are removed from it.

Prerequisites and precautions for deleting a communication domain:

1. Only communication domains with one primary controller can be deleted.

2. Deleting a communication domain will delete its network configuration, but not that of the primary controller.

7.2.5.2 Application Scenarios

A newly created communication domain requires to be configured with an IP address for the VXLAN interface to interconnect the VXLANs of all physical hosts in it, so that the virtual devices in different resource pools can communicate with each other. The main application scenarios include:




Figure 52: VXLAN Interconnection in Communication Domain

1. VPC Network Interconnection Across Resource Pools
 - VPC networks in a communication domain can connect to each other across resource pools. For example, if VM 1 of resource pool A and VM 2 of resource pool B are in the same VPC subnet, then they can ping each other. VMs can be scheduled within a resource pool. VPC subnets in a communication domain are connected to each other.
 - VPC subnets are connected to the internal router, and east-west traffic between them passes through the VXLAN interface.

- North-south outbound traffic from VPCs goes through the external virtual router, whose reliability is ensured by the Virtual Router Redundancy Protocol (VRRP). The primary instance is scheduled to a resource pool configured with an edge.
- The EIP applied to a VPC can be associated with a VM in any VPC subnet in a communication domain, or the external router for communication with the external network.

2. VPC Gateway Migration Across Resource Pools Through HA Mechanism

When a resource pool in a communication domain is abnormal, its routers of the VPC gateway and shared service gateway can be automatically migrated to other available resource pools in the communication domain through the HA mechanism. This scenario has three deployment requirements:

- The edge configured for the VPC network and shared service network of the abnormal resource pool must be configured with a physical NIC of another resource pool, meaning that there is an instance of the edge in another resource pool.
- The connection between the resource pools that can be migrated to and the elastic IP pool associated with the abnormal resource pool must be normal because EIPs are migrated with the routers. Please make sure that the EIPs can continue to provide services after the migration.
- BGP routes to the EIPs from the data center or parent network must be updated after the migration to make sure that the external access can reach the correct location.

3. Direct Connection Gateway Migration Across Resource Pools Through HA Mechanism

When a resource pool in a communication domain is abnormal, the direct connection gateway of its VPC can be automatically migrated to other available resource pools in the communication domain through the HA mechanism. This scenario has two deployment requirements:

- The edge configured for the direct connection gateway of the abnormal resource pool must be configured with a physical NIC of another resource pool, meaning that there is an instance of the edge in another resource pool.
- Resource pools that can be migrated to are accessible to each other through direct connections.

7.2.6 Traffic Mirroring

Traffic mirroring refers to copying the messages of traffic from specified ports (VM interfaces or edge-connected interfaces) to other ports (VM interfaces or edge-connected interfaces) of any hosts in the same cluster or communication domain. The target ports are connected to a data monitoring device that users can use to analyze the mirrored messages for traffic audit, network monitoring, and troubleshooting.

7.2.6.1 Application Scenarios

1. Network Intrusion Detection for Traffic Protection

Self-developed or third-party security software can fully check traffic to capture all potential vulnerabilities and intrusion threats for rapid response to attacks.

2. Traffic Audit in Finance or Government

Financial or government business traffic with high security and compliance requirements can be transparently mirrored to unified audit platforms for analysis.

3. Issue Location for Network O&M

O&M engineers can troubleshoot network issues by directly checking transmission content (such as analyzing TCP retransmit) through traffic mirroring, eliminating the need to access the inside of VMs to capture messages.

7.2.6.2 Supported Features

1. Mirror the traffic from a VM interface or network device interface to another one on the same host.
2. Mirror the traffic from a VM interface or network device interface to an edge-connected interface on the same host.
3. Mirror the traffic from an edge-connected interface to another one on the same host.
4. Mirror the traffic from a VM interface or network device interface to another one on another host.
5. Mirror the traffic from a VM interface or network device interface to an edge-connected interface on another host.
6. Mirror the traffic from an edge-connected interface to another one on another host.
7. Automatically migrate traffic mirroring policies as VMs migrate, avoiding manual configuration.

Mirroring the traffic from a physical interface requires specifying a VLAN ID to ensure that the mirrored traffic is isolated from other traffic. Mirroring the traffic between different hosts uses the Encapsulated Remote Switch Port Analyzer (ERSPAN) to encapsulate the traffic from the source port and send it to the target port on the target host. During encapsulation, TCP segmentation offload (TSO) is performed on the memory messages to prevent outer-layer segmentation of the encapsulated messages. Users can configure the mirror percent to adjust the number of mirrored packets and set the direction of mirrored traffic.

7.2.6.3 Policy Configuration

Users only need to configure one traffic mirroring policy to mirror the traffic from multiple VM interfaces or edge-connected interfaces to other interfaces on the same host or different hosts.

Create Traffic Mirroring Policy

Support traffic mirroring within a virtual network, from a virtual network to physical network, and between physical interfaces. [Configuration Guide](#)

Name:

Mirror Source:

Select

Mirror Target:

Select

VLAN ID:

Mirror Percent ⓘ :

100

%

Traffic Direction ⓘ :

☒ All

☐ Inbound

☐ Outbound

Status:

☒ Enabled

OK

Cancel

Figure 53: Configuration of Mirroring to VM Interfaces

Create Traffic Mirroring Policy

Support traffic mirroring within a virtual network, from a virtual network to physical network, and between physical interfaces. [Configuration Guide](#)

Name:

Mirror Source:

3 Edge-Connected Interface selected

Mirror Target:

Select

VLAN ID:

Mirror Percent ⓘ :

100

%

Traffic Direction ⓘ :

☒ All

☐ Inbound

☐ Outbound

Status:

☒ Enabled

OK

Cancel

Figure 54: Configuration of Mirroring to Edge-Connected Interfaces

Steps:

1. Log in to HCI or SCP.
2. Add a traffic mirroring policy.
3. Specify the **Name**.
4. Select the interface type and interfaces for **Mirror Source**.
5. Select the interface type and interfaces for **Mirror Target**.
6. Specify the **VLAN ID** (configured only when **Mirror Target** is edge-connected interfaces).
7. Specify the **Mirror Percent**.
8. Select a **Traffic Direction**.
9. Click **Enable** for **Status**.
10. Click **OK**.

7.2.7 QoS

7.2.7.1 Introduction

1. Global QoS for Communication Domain or Resource Pool

System administrators can configure global Quality of Service (QoS) for a communication domain or resource pool by setting its outbound and inbound bandwidth to up to 20,000 Kbps.





Figure 55: Resource Limit

2. QoS for VM NIC

System administrators and tenants can configure QoS for a VM NIC by setting

its outbound and inbound bandwidth to up to 20,000 Kbps. When a VM NIC is configured with QoS and the resource pool in which it resides is configured with global QoS, the QoS for the VM NIC prevails.



The screenshot shows the configuration window for NIC1. The 'NIC' section has a checked 'Enable' checkbox. The 'Connected To' dropdown is set to 'ManageSW', with 'Specify IP' and 'Disconnect' buttons. The 'Advanced' section is expanded, showing 'MAC Address' as 'fe:fc:fe:d7:fe:7e' with a refresh icon, and 'Adapter Model' as 'Intel E1000'. The 'Traffic Limit' section has a checked 'Enable' checkbox. Below it, 'Outbound' and 'Inbound' settings are both set to 'Optional' with a unit dropdown menu set to 'Kbps'.

Figure 56: Traffic Limit

7.2.7.2 Application Scenarios

When a VM on a virtualization platform is poisoned due to, for example, ransomware or ARP attacks, it will send a lot of packets to other VMs on the platform, occupying a large amount of bandwidth and affecting the normal communication of business services. In addition, the live migration, backup, clone, and deployment through templates of VMs all consume significant amounts of bandwidth. Therefore, VMs running critical business services need to have a certain amount of bandwidth guaranteed.




Figure 57: QoS Application Scenarios

7.3 Storage

7.3.1 Block Storage

7.3.1.1 Introduction

The block storage service mainly provides block storage space in the form of cloud disks for VMs for data running, storage, and backup. Users can add cloud disks to a VM when creating it. A cloud disk is used in the same way as a disk of a physical server, but features higher data reliability and I/O throughput, as well as being easier to use. It provides high storage performance and reliability through technologies such as read/write cache tiering, striping, I/O location, and multi-replica.

7.3.1.2 Application Scenarios

1. Hosting of key business services, such as enterprise resource planning (ERP) system, manufacturing execution system (MES), and other systems.
2. Hosting of enterprise self-built databases, such as Oracle, MySQL, and SQL Server.
3. Use for system disks, data disks, and shared disks.

7.3.1.3 Supported Features

1. Disk Addition

Users can add disks to VMs. A VM can be mounted with up to 16 disks.

2. Disk Initialization

Newly added disks need to be initialized (partitioned and mounted) by users for use.

3. Disk Encryption

Users can encrypt disks to improve data security and meet compliance needs.

4. Disk Capacity Expansion

Users can expand the capacity for disks with insufficient capacity. Before expansion, the VM on which the disks are mounted must be shut down. After expansion, users can allocate the expanded capacity to existing partitions or format it into a separate new partition, based on the OS type, disk partitions, and file system of the VM.

5. Disk Deletion

Users can delete disks that are no longer needed by detaching or detaching and deleting them. The VM on which the disks to be deleted are mounted must have no snapshot; otherwise, the disks cannot be deleted.

6. Recycle Bin

Users can contact the administrator to retrieve detached disks.

7. Storage Volume Splitting

SCP supports host volume splitting and disk volume splitting. Host volume splitting groups hosts in an HCI cluster to different storage resource pools. Disk volume splitting splits volumes at disk granularity into clusters of three or more hosts.

8. Storage Policies

SCP provides redundancy policies based on multi-replica and erasure coding (EC) mechanisms to meet business requirements. The multi-replica mechanism saves multiple replicas of data to prevent business service interruption due to

data corruption, while the EC mechanism encodes the original data to obtain redundancy through the EC algorithm, splits the data into segments, expands and encodes the data blocks, and stores the data and redundancy together in different locations for fault tolerance.

SCP supports three disk allocation modes, including thin provisioning, pre-allocating, and dynamic provisioning.

9. Advanced Storage Features

Sangfor Managed Cloud Services' block storage service supports iSCSI protocol and virtual shared disks.

10. Fault Prevention and Handling

SCP supports arbitration mechanisms, silent error detection, data rebuilding, and disk health checks.

7.3.2 File Storage

7.3.2.1 Introduction

The network-attached storage (NAS) is a high-performance distributed file system that can be shared and accessed massively and scaled elastically.

Sangfor file storage service has the following features:

1. Built-in self-developed distributed database PhxKV, which processes massive unstructured metadata to meet performance requirements.
2. Automatic tiering and caching, which stores hot data in SSDs, thereby reducing data transfer latency and improving storage IOPS.
3. High reliability guarantee. Customers can flexibly use the multi-replica mechanism of 2 or 3 replicas and the EC mechanism of 4+2 or 8+2 according to their needs.
4. Visualized intelligent management, which monitors and displays all hardware on the platform, including CPU, disks, NICs, cables, and interfaces, as well as the performance metrics, running trends, and resource usage of each resource pool. Customers can easily view the running situation of resource pools and deal with their problems.

7.3.2.2 Application Scenarios

File storage is widely used for enterprise application data sharing, containers, AI learning, web service and content management, application development and testing, media and entertainment workflows, and database backup.

Examples of typical scenarios:

1. Shared Storage Containers

A container cluster stores configuration files or initial load data of container images on shared file storage and reads them in real time when bulk loading containers. Persistent data is shared between multiple pods through the NAS. A pod can fail over in case of failure.

2. Websites and Web Applications

High-performance websites compose services into clusters, and store code files, configuration files, or images (most of which are as small as tens of kilobytes) and other business service data on the NAS for shared access, which requires a high level of system response speed.

3. Real-Time Log Saving

As business services rapidly grow, the logs generated by a system continue to increase with a large amount of valuable information. Storing data locally may crash disks and cause system failures. It is necessary to save these logs conveniently, centrally, and cost-effectively on the NAS.

4. Development and Testing Environments

A development and testing environment shares code programs through the NAS to speed up data transfer between different platforms, and continuously integrates the code data of the shared access of the system for rapid engineering releases.

7.3.2.3 Supported Features

1. Quota Management

Supports creating quotas to limit the capacity (1-100 PB) and number of files (1,000-10 billion) of shared directories.

2. Multi-Protocol Sharing

- Supports sharing a NAS directory to multiple business systems (Windows, Linux, and Mac) to implement multi-protocol file collaboration like writing a file on a Linux client while reading it on a Windows client.
- Supports shared access to a shared directory through NFS, CIFS, and FTP protocols. By configuring mapping rules for AD domain users and local authenticated users, they can access file directories created through NFS and FTP protocols across protocols.

3. AD Domain Authentication

Supports configuring AD domain names and user account information on Sangfor Enterprise Distributed Storage (EDS) to add EDS to the customers' AD domains to synchronize the user information on the domain controllers. By selecting user or user group information in an AD domain for the shared directory of the EDS configuration file, clients in the domain can access the shared directory without authentication.

4. Multi-User Through SMB

Supports up to 500 users to access the CIFS shared directory of EDS through the Server Message Block (SMB) protocol at the same time.

5. File Audit Logs

Supports administrators to enable file audit logs to easily filter the logs related to file creation, deletion, and metadata modification to trace abnormal files.

6. File Metadata Write

Supports storing file business metadata on cache disks to improve file access performance and massive access stability.

7.3.3 Object Storage

7.3.3.1 Introduction

The object storage service (OSS) is an object-based massive storage service that offers secure, highly reliable, and cost-effective cloud data storage and processing. SCP provides customers with the object storage service through the integration with Sangfor EDS.

7.3.3.2 Application Scenarios

With features such as bucket and object versioning management and QoS, object storage is widely used for business services that require high concurrency, low latency, high reliability, and cost-effectiveness. It can effectively meet the backup and archiving needs of enterprise applications, databases, and unstructured data, making it the best choice for massive data storage. Examples of typical scenarios:

1. Massive Storage of Files such as Images, Audios, and Videos

Object storage provides highly scalable, reliable, and cost-effective storage space with high performance and low latency for massive files such as images, audio, videos, and logs. Various endpoints, web programs, and mobile applications can write to and read from the object storage directly.

2. Enterprise Cloud Disk Applications

Object storage works with the enterprise cloud disk feature on SCP to create a storage system with high concurrency, low latency, high reliability, and cost-effectiveness, and whose capacity can be automatically expanded as users' data volume increases.

Endpoints first send dynamic data requests to the business system of the enterprise cloud disk on SCP, and the latter processes and returns them directly to the endpoints. The static data is saved in the object storage service and processed by the business system through the internal network. The endpoints request and retrieve static data directly from the object storage

service. Meanwhile, the object storage service has a lifecycle feature, which enables automatic conversion between various object storage and effectively saves the cost of storage resources.

7.3.3.3 Supported Features

1. Account Management

- Supports automatically synchronizing every 10 minutes or manually synchronizing object storage accounts on EDS.
- Supports creating object storage accounts, selecting the resource pools and users (administrators or tenants) to which they belong, and setting capacity quotas for them. After the account is created, a key file will be downloaded automatically. It cannot be downloaded repeatedly and must be saved safely. Each account can generate up to five key files, and users must delete a discarded key file to regenerate another when there are five.
- Supports the super administrator or MSP administrators to allocate object storage accounts to or reclaim them from tenants (individually or in bulk).

2. Bucket Management

Supports full lifecycle management of buckets in object storage pools, and display of content such as **Outbound Traffic** and **API Calls** in fixed or custom periods. The daily data is synchronized between 00:00 and 06:00 the next day.

Supports creating buckets and selecting the object storage accounts to associate them with. An object storage account can be associated with multiple buckets, which have the following settings:

- Permissions

Supports setting permissions for all object storage accounts and all visitors to storage pools of buckets.

- QoS

Bandwidth: Unlimited by default and can be customized.

Requests per Second: Unlimited by default and can be customized.

Hide QoS for Specified Accounts: Supports individually setting bandwidth and requests per second limits for specific accounts, which are only subject to these settings.

- Advanced

Versioning: Disabled by default. If enabled, all versions of objects with the same name in the buckets will be retained.

- Lifecycle Management

Supports clicking **Create Rule** to create lifecycle management rules for objects in the buckets.

3. Summary

Supports collecting and querying the usage of object storage, and exporting them as reports by selecting object storage pools, users, and time ranges on the **Summary** page of the object storage module.

7.4 Cloud Security

7.4.1 Distributed Firewall

7.4.1.1 Introduction

Attackers can easily target weakly protected servers when they manage to bypass the edge defenses using various methods, enabling them to attack other servers within the data center with ease. Manual SP installation on a large number of servers in the data center is too time-consuming to prevent extensive paralysis of the data center caused by the rapid spread of viruses. Therefore, SCP offers a distributed firewall (DFW) to extend security from the data center edge to its core. This enables micro-segmentation between VMs and ensures L3 and L4 security protection for internal traffic, significantly reducing the impact of attacks on the data center.




Figure 58: DFW Principle

7.4.1.2 Application Scenarios

In 2016, the Shadow Brokers disclosed multiple remote Windows vulnerability exploit tools, shocking the entire IT community. Windows servers were at severe risk of attacks if any of the following ports were open: 135, 137, 139, 445, and 3389. Bitcoin ransomware had already emerged by the time the vulnerability was discovered.

However, SP installation on all servers for a data center with hundreds of Windows servers is time-consuming. If a hacker penetrates the firewall and gains access to the data center during this process, any Windows server with the corresponding ports open becomes an easy target.

7.4.1.3 Supported Features

1. Comprehensive Protection for All VMs

DFW can be configured on any distributed switch port to protect the inbound and outbound traffic in a region of SCP and ensure micro-segmentation.

Corresponding policies are synchronized when a VM is migrated to another

node.

2. Rapid Enforcement of Protection Policies

DFW can be directly and swiftly deployed on SCP without installing any plugins due to the simple architecture. The bypass and interception log display features ensure quick verification of rule effectiveness.

3. Simplified Network Structure for Easy Troubleshooting

SCP DFW has a simple network structure featuring easy O&M through "What-You-Draw-Is-What-You-Get." The connectivity test enables quick troubleshooting.

4. Collaboration with Sangfor vNGAF for Enhanced Security

DFW is often used to control internal traffic in virtual networks because its distributed architecture ensures platform-wide policy effectiveness, while vNGAF is usually deployed at the edge to defend against external traffic attacks. DFW is available once aNET is licensed. DFW, in combination with vNGAF, can provide all-around security protection for data centers, covering the north-south, east-west, external, and internal traffic.

7.4.1.4 Policy Configuration

7.4.1.4.1 Create Policy

Sangfor DFW offers east-west network isolation to protect all VMs within a data center using a single rule.

Create Policy

Name:

Scope:

Priority:

Priority	Name	Source	Destination	Service	Action	Status	Operation
----------	------	--------	-------------	---------	--------	--------	-----------


No data available

Figure 59: DFW Policy Configuration

- **Policy Name:** Deny inter-VM communication on ports 135, 445, and 3389 within the data center.
- **Source:** All VMs in the data center.
- **Destination:** All VMs in the data center.
- **Service:** Select ports 135, 137, 139, 445, and 3389 for both TCP and UDP.
- **Action:** Deny.

7.4.1.4.2 Import Policy

Sangfor distributed firewall supports importing distributed firewall policies from VMware vCenter before the VM live migration, preventing network disconnection during the migration.



Policies, rules, and services remain consistent with those on VMware vCenter, and users can edit these policies and rules.

■Technical Principles

The migration tool enables the migration of distributed firewall policies on VMware vCenter. Solutions: The tool authenticates with VMware vCenter and calls OpenAPI provided by VMware vCenter to retrieve the complete distributed firewall configuration, including policies, rules, VMs, VM groups, and IP address groups. Perform the corresponding data transfer to generate a standard CSV file that can be identified by SCP. Log in to SCP, import the file, and create distributed firewall.

■Description

The distributed firewall migration feature consists of two modules: Distributed firewall policy export and distributed firewall policy import. After distributed firewall policy migration is complete, VM migration starts. The policy will automatically identify referenced VMs after VM migration and apply the firewall feature to these VMs. Distributed firewall policy export: You can connect to the VMware NSX environment using the migration tool provided by SANGFOR Checks (6.1.1), obtain all VMware vCenter distributed firewall policies, and export them as a predefined standard CSV file for subsequent import. Distributed firewall policy import: After the CSV file is exported using the tool, you can import the file into SCP to create corresponding distributed firewall policies.

Note: The imported policies associated with VMs will take effect only after the VMs are migrated.

7.4.2 VPC Firewall

SCP supports configuring access control policies for the edge routers of tenant VPCs to create secure network boundaries. vNGAF is deployed in Layer 3 mode at VPC gateways to provide a service-oriented security protection capability that is ready to use by tenants.




Figure 60: VPC Firewall

When a tenant enables vNGAF protection, SCP automatically configures routing rules to direct all external network traffic to vNGAF through vNGAF's built-in protection policies. vNGAF then cleans the traffic according to its rules and sends the secure traffic to the router. Non-external traffic (internal traffic, direct connection traffic, and shared service traffic) does not need to go through vNGAF.

Tenants can troubleshoot by stopping the protection. If stopped, external network traffic will no longer pass through the firewall.

Tenants can also limit the maximum traffic of applications through the traffic control of vNGAF to prevent ordinary business services from preempting the bandwidth of critical business services, thus safeguarding the latter's access experience.



The VPC firewall is built based on the security component vNGAF.

7.4.3 aSecurity

Built-in security is a concept that allows a cloud platform to directly provide security capabilities for cloud-based services, eliminating the need for add-on security components. On SCP, this concept is implemented through the construction of aSecurity. Compared to add-on security components, built-in security can better meet business requirements, ensuring security and reliability while reducing security O&M costs. This brings significant benefits to enterprise digital transformation.

aSecurity offers capabilities such as built-in VM security and cloud firewall, enabling features like antivirus, anti-ransomware, intrusion defense, and application security protection.

7.4.3.1 Technical Architecture

The technical architecture of aSecurity consists of three layers:

1. Basic Platform Layer

The core technology is Hypervisor, which virtualizes physical resources to maximize resource utilization. HCI provides basic computing, storage, and network features, while SCP offers cluster, tenant, and permission management capabilities.

2. Core Engine Layer

Based on the capabilities at the basic platform layer, the core engine layer encapsulates interfaces of different platforms (such as Endpoint Secure, CWPP, Neural-X, HCI, and SCP). It integrates the capabilities of various management platforms and security platforms to develop basic component features (such as multi-version management and AK/SK authentication) for aSecurity.

3. Function Presentation Layer

Based on the basic component features provided by the core engine layer, actual business scenarios, and feature characteristics, this layer is divided into multiple functional modules for management, such as MGR management, threat management, and vulnerability management. The business interfaces are called on the UIs that are embedded on HCI and SCP pages. Interface requests are routed from the platforms to Sangfor-WAF of aSecurity and then to the aSecurity services. Then, the UIs receive service responses.




Figure 61: aSecurity Architecture

7.4.3.2 Dashboard

The **Dashboard** page of aSecurity displays an overview of the platform's security, including **Pending Tasks**, **Threat Overview**, **Security Trend**, and **Protection and Response**. The **Pending Tasks** at the top of the page cover security events, vulnerabilities, and unprotected assets, and are of the following types: New assets requiring vmTools installation or security protection plan configuration, ransomware risks, general security events, hot events, and general vulnerabilities. Users can click **Fix Now** to quickly redirect to the corresponding fixing page.




Figure 62: Dashboard

7.4.3.3 VM Security

aSecurity provides the **VM Security** module for unified asset management, where all VMs on SCP are displayed along with their basic and security information, eliminating the need for complex inventory. This module includes the following features:

1. **Asset Management:** Supports automatically obtaining the basic and security information of VMs, including severity, running status, name, IP address, OS, application, protection status, and security events.
2. **Virus and Vulnerability Scan:** Through AI-based Engine Zero, Gene Analytics Engine, Behavioral Analytics Engine, and Cloud Detection and Killing Engine, supports issuing virus and vulnerability scan tasks to specific VMs, and detecting and fixing viruses and vulnerabilities with the following operations: **Fix, Trust, Restore, Untrust, and Ignore.**
3. **Quarantine:** Supports quarantining encrypted assets in seconds after a virus infection to prevent the virus from spreading.
4. **One-Click Security Event Fixing:** Provides full-stack wizard-like fixing for security events from network quarantine, snapshot creation, to virus killing.

5. Ransomware Recovery: Supports full-stack wizard-like fast and secure business recovery from ransomware quarantine to network recovery without expert intervention.
6. Network Protection: Provides cyber attack protection capabilities for cloud-based business systems across seven layers of the communication process.

7.4.3.4 Security Capabilities

The **Security Capabilities** module includes features such as security event management, vulnerability management, cyber attack protection, and data protection policy.

1. Security Event Management: Provides management and fixing of security events such as virus infections, brute-force attacks, and cyber attacks.
 - Provides full-stack wizard-like fixing for virus infections from network quarantine, snapshot creation, to virus killing.
 - Provides blocking, trusting, and ignoring operations for brute-force attacks.
 - Provides cyber attack blocking by adding their source IP addresses to the blacklist to prevent system paralysis, data leakage, and information security threats.
2. Vulnerability Management: Supports vulnerability monitoring views for all VMs to fully understand the vulnerability risks of business systems and one-click Windows vulnerability bulk repair to strengthen business security protection.
3. Cyber Attack Protection: Supports configuring cyber attack protection policies for VMs in the virtual network topology, providing protection for cloud-based business systems across seven layers of the communication process.
 - **Intrusion Prevention System (IPS):** Detects and defends against attacks exploiting OS vulnerabilities, application vulnerabilities, and protocol vulnerabilities.
 - **Web Vulnerability Protection:** Anti-attack policies for web servers that detect and defend OS command injection, SQL

injection, XSS attacks, and other attacks against web applications.

4. **Data Protection Policy:** Supports enabling the data protection policy feature to protect all assets by taking daily snapshots, to trigger snapshots for data protection immediately when a security software uninstallation or suspected ransomware is detected, and to persist the original and periodically protected files to prevent them from being overwritten by poisonous data when an anomaly is detected.

7.4.3.5 Logs

aSecurity provides the capability to query antivirus logs, brute-force attack logs, cyber attack logs, and vulnerability scan logs to help administrators analyze and fix security events.

7.4.3.6 Security Settings

The **Security Settings** module includes features such as alert options, signature database update, blacklist and whitelist, troubleshooting, security agent installation, security capability management, aSecurity upgrade, and remote maintenance.

1. **Alert Options:** Supports detecting in real time security events that threaten VMs, such as brute-force attacks, WebShell backdoors, ransomware, and Trojans, and sending notifications to the users' email addresses when they are identified; supports setting the entities, limits, and frequency of notifications, as well as the severity of the vulnerabilities.
2. **Signature Database Update:** Supports updating the WAF rule base, and manually and automatically updating the IPS rule base to identify and block new threats in time, improving the security of the system.
3. **Blacklist and Whitelist:** Supports blacklisting cyber attack events and brute-force attack events, and whitelisting specific IP addresses, VMs, security policies, and security events.

4. **Troubleshooting:** Supports troubleshooting when the network is disconnected or a service or application cannot be accessed with two methods: precise traffic analysis and global passthrough analysis. The precise traffic analysis obtains the data flow whose packets have been dropped due to firewall policies to analyze intercepted access requests. The global passthrough analysis disables all firewall policies within a specified range for use when the firewall policy that intercepts business accesses cannot be identified. To use it, disable global firewall policies and make sure that the business network connection is normal.
5. **Security Agent Installation:** Supports viewing VMs without security agent installed and selecting VMs to install the security agent on.
6. **Security Capability Management:** Supports configuring the manager component for VMs, providing comprehensive capabilities including threat detection, defense, and response; supports enabling VM security protection and cyber attack protection for resource pools.
7. **aSecurity Upgrade:** Used to instruct administrators to upgrade the aSecurity version.
8. **Remote Maintenance:** Supports contacting a technical support representative for remote diagnostics, troubleshooting, and technical problem resolution to improve system performance or recover business services.

7.4.4 Security Components

aSecurity-NFV provides a diverse range of security services through the security resource pool, including vSSL VPN, vNGAF, aTrust, vIAG, vOSM, vDAS, Sangfor Host Security, Cyber Command-Logger, and vCyber Command & vSTA.

7.4.4.1 vNGAF

Traditional firewalls can no longer defend against more and more attacks, which requires a more effective next-generation application firewall. Network Secure (NGAF) is equipped with comprehensive business protection and visualization capabilities, identifies applications, users, and content for access

control, and aggregates edge security devices for multi-module processing as well as integrates with other data, systems, and devices for management and analysis all while safeguarding performance.

Network security edges refer to the edge between the cloud platform and the internet or other external areas (cloud edge) and the edge between internal application systems within the cloud platform (intra-cloud). Deployment of vNGAF at the cloud edge enables cyber attack protection, intrusion defense, and application security protection.




Figure 63: vNGAF Modules

vNGAF consists of functional modules of Intrusion Prevention System (IPS) and Web Application Firewall (WAF). It takes application assets as its core, builds a comprehensive security protection system applicable to multiple scenarios based on edge defense, providing five core capabilities: comprehensive application asset identification and protection, precise application access control, professional application security protection, comprehensive unknown threat detection and blocking, and continuous threat information sharing.

7.4.4.2 aTrust

aTrust is dedicated to providing government and enterprise clients with essential product components and comprehensive solutions for establishing a zero-trust security system. It caters to the security access requirements of clients in diverse business scenarios, including intranet and internet access, multi-branch connection, cloud migration, and mobile office, facilitating the

swift migration of government and enterprise businesses to a zero-trust security architecture.

Sangfor aTrust follows the NIST ZTA zero trust architecture standard. It combines the software-defined perimeter (SDP) and unified endpoint management (UEM) technologies to provide enterprise clients with a comprehensive hybrid office solution that includes secure access and data leakage prevention.

The following figure presents the overall architecture of aTrust.




Figure 64: Architecture of Sangfor aTrust

The aTrust client currently supports all mainstream OSes, including Windows, macOS, UOS, Kylin, Ubuntu, Android, iOS, and HarmonyOS. It offers features such as SPA knocking, SSL tunnel access, endpoint environment detection, trusted application identification, and data security sandbox.

Following the design principle of separating the control plane and the data plane on the server side, aTrust consists of two core components: aTrust Controller and aTrust Proxy Gateway.

aTrust Controller serves as the control plane of the entire aTrust solution. It functions as the policy decision point (PDP) in the NIST ZTA and the controller in the SDP, responsible for making final access authorization decisions. It provides management and control features such as user management, endpoint management, adaptive authentication, dynamic access control, and

policy management. Additionally, it offers open capabilities for integration and synchronization with existing enterprise identity infrastructure (such as 4A and PKI).

aTrust Proxy Gateway serves as the data plane of the aTrust solution. It functions as the policy enforcement point (PEP) in the NIST ZTA and the accept host (AH) in the SDP, responsible for executing the authorization policies from the controller, forwarding authorized access requests, and intercepting unauthorized access requests. It is typically deployed in a logical cascading manner at the uplink of protected resources, reducing resource exposure. External users and endpoints cannot directly access the protected resources but need to go through strict identity authentication and authorization decisions by the Controller. The access is then proxied by the Proxy Gateway. aTrust Proxy Gateway supports various access protocol proxies, including Layer 7 web proxy, Layer 4 TCP proxy, and Layer 3 IP proxy.

For detailed technical information about aTrust, refer to the [Sangfor aTrust Technical White Paper](#).

7.4.4.3 vIAG

With the development of internet technologies, organizations' business models and employees' work patterns and habits are constantly changing, introducing new risks to enterprises:

1. Increased intranet security risks due to unidentified identities and devices accessing the network.
 - **Unauthorized Access:** External individuals can easily connect to the intranet using an Ethernet cable, posing risks to various business systems.
 - **Non-Compliant Devices Accessing Network:** Non-compliant devices, such as those without antivirus software, using non-compliant OSes, or with weak passwords, pose significant threats to the intranet environment, potentially leading to the rapid spread of viruses within the network.

2. Hidden risks behind seemingly normal internet access activities.

- **Bandwidth Occupancy:** Employees using applications (such as P2P applications and streaming media) during work hours occupy great bandwidths, which severely affect email sending, data downloading, video conferences, and other activities critical to the company's core business.
- **External Leakage of Confidential Information:** Sensitive data and files, including personal privacy information, organizational confidential information, government documents, and official files, may be sent to external objects unintentionally or casually without monitoring, alerts, or compliance auditing. The operator is unknown.
- **Cybercrime:** The organization's network may be used to spread rumors, launch personal attacks, and indiscriminately distribute reactionary, gambling, and pornographic content, resulting in significant harm to the company and organization and potential legal consequences.
- **Challenges in Internet Access Supervision:** The workplace becomes an unregulated internet café, leading to decreased work efficiency. Advanced encryption and proxy technologies complicate the control over unauthorized content. Additionally, an application has both "sides", making it challenging to determine whether to implement control measures.

3. Abnormal business access is hard to detect, and it is difficult to manage account misuse, resulting in tracing issues.

- **Invisible Business Behavior:** When employees access the business system, comprehensive auditing operations make it difficult to determine the responsibility in cases of operational errors or data leaks.
- **Difficulties in Detecting Abnormal Behavior:** When employees access the business system, some malicious or unintentional actions may harm the system, such as excessive

data download, crawling of all data, and execution of sensitive operations like deletion and clearing.

Sangfor Internet Access Gateway (IAG) effectively addresses these risks and provides control over endpoint access, user internet access, and endpoint behavior.

1. Endpoint Access Control

Sangfor IAG accurately identifies connected endpoint assets, offers various identity authentication methods (including 802.1x and Portal), conducts security baseline checks, and controls peripheral device permissions. This helps users establish secure endpoint access standards to reduce security risks and data leakage.

2. Internet Access Control

For endpoints connected to the network, Sangfor IAG utilizes databases containing vast URLs and application rules to achieve precise application control. It ensures comprehensive compliance auditing and supports dynamic flow control and application traffic forwarding to allocate bandwidth properly. This helps users maintain internet access compliance, enhance user experience, and prevent unauthorized activities.

3. Endpoint Behavior Control

Lightweight desktop control features include port control, application network connection control, software installation management, screen watermarking, external connection control, and peripheral device control. They enable administrators to effectively manage and control endpoint behavior, thereby improving system security.

Sangfor IAG integrates network and endpoint control to promote more standardized, efficient, and secure office operations.

For detailed technical information about vIAG, refer to the [Sangfor IAG Technical White Paper](#).

7.4.4.4 vSSL VPN

The vSSL VPN component aims to create value for customers by leveraging

innovative solutions tailored to their specific business needs. Based on the comprehensive understanding of the customer's business situation, it effectively resolves challenges encountered during the process of internet-based transformation. In addition to conventional SSL VPN solutions like mobile office solutions and multi-party access permission allocation schemes, vSSL VPN also addresses security issues of customers' current critical systems by leveraging SSL security features. It ensures the security and confidentiality of core customer data through SSL encryption and logical isolation.

7.4.4.5 vADC

Sangfor Application Delivery Controller (ADC) is a professional device that offers a comprehensive solution for application delivery. It provides features such as multi-data center load balancing, multi-link load balancing, and server load balancing. With performance optimization, single-edge acceleration, and multiple intelligent management schemes, it monitors the status of data centers, links, and servers in real time and delivers user access requests to appropriate data centers, links, and servers based on predefined rules. This optimizes data flow allocation and maximizes utilization of data centers, links, and servers. It expands the processing capacity of the application system, improves stability, enhances user experience, and reduces IT investment costs.

7.4.4.6 vOSM

vOSM is a crucial component for MLPS compliance. It provides users with integrated O&M security audit services, including account management, identity authentication, single sign-on, resource authorization, access control, and operation audit. It effectively audits the O&M processes of servers, network devices, security devices, and databases, transforming O&M auditing from event-based to content-based. Using preventive, control, and audit measures implemented through the internal control management platform, it comprehensively addresses O&M security issues.

7.4.4.7 vDAS

Sangfor Database Audit System (DAS) is a compliance management system that performs detailed auditing of database operations in a business environment. It analyzes, records, and reports on the behavior of business personnel accessing the system, assisting users with pre-planning, real-time monitoring, response to violations, compliance reporting, and incident tracking to ensure the normal operation of core assets. vDAS can monitor the real-time operation flow of the database server, intelligently analyze various operations, and provide log reports for post-analysis and forensics.

7.4.4.8 Sangfor Host Security

Sangfor Host Security helps users centrally collect, analyze, and manage security configurations for cloud assets. It is a crucial support for daily cloud information security work. As a unified security configuration check and management system, it accurately and quickly identifies and summarizes security configuration issues and vulnerabilities across different hosts, databases, middleware, networks, and security products. Its main features include:

1. Task Formulation: Flexibly formulates security baseline check tasks of various types and cycles, allowing easy and quick setup of check objects and policies.
2. Collection and Analysis: Comprehensively and centrally checks and analyzes local security configuration issues, vulnerabilities, weak passwords, and other issues, reducing redundant work caused by scattered device management.
3. Violation Reporting: Provides comprehensive, detailed, and clear management and comparative analysis of scan reports.
4. System Reinforcement: Provides detailed and practical system reinforcement solutions to guide users in resolving security issues.

7.4.4.9 vCyber Command-Logger

vCyber Command-Logger continuously collects and aggregates log information

from cloud tenant hosts, OS, and business systems in real time. It assists users in security analysis, compliance audit, and timely detection of security events and violations.

7.4.4.10 aSI

Sangfor HCI supports security integration with Sangfor Cyber Command, enabling intelligent fixing of security events.

Cyber Command is a security detection product that utilizes AI and big data security detection technology. It enhances detection accuracy, identifies unknown threats, and takes action by integrating with various security devices. Through Stealth Threat Analytics (STA) for hardware, it mirrors the traffic of core switches in data centers. In scenarios where Cyber Command functions as the local security center, it integrates with different security devices to form a comprehensive security protection system.

It relies primarily on STA for detection:

1. STA is an independent hardware device connected to core switches to forward the traffic to Cyber Command.
2. In addition to traffic forwarding, STA can also analyze traffic and provide the result to Cyber Command.




Figure 65: Cyber Command & STA Workflow

Currently, Cyber Command has integrated with various security devices,

including Endpoint Secure, Network Secure, and IAG. It collects security logs from these devices for analysis, statistics, and display and takes correlated fixing measures.

Cyber Command has strong capabilities for detecting and identifying security issues, while HCI has full operational control over business VMs. By combining the detection capabilities of security products with the fixing capabilities of HCI, a more effective and comprehensive security solution can be achieved.




Figure 66: Cyber Command & STA & HCI Workflow

HCI is managed by Cyber Command as a cluster with the account and password provided. It provides interfaces for synchronizing the VM list, shutting down and suspending VMs, configuring the distributed firewall, and taking storage snapshots. Therefore, the list and on/off status of VMs in the HCI cluster can be viewed on Cyber Command. The HCI cluster can automatically sync the VM list with Cyber Command at an interval of 30 seconds to 1 minute.

When a security event is detected on Cyber Command, it supports operations such as shutting down and suspending VMs, taking storage snapshots, and configuring distributed firewall security isolation strategies for VMs. The processing mechanism based on the integration of Cyber Command and HCI is shown in the following figure.




Figure 67: Processing Mechanism Based on Integration of Cyber Command and HCI

The integration between HCI and Cyber Command is based on north-south traffic. East-west traffic of VMs cannot be mirrored externally, so no effective protection is provided for east-west traffic.




Figure 68: Traffic Acquisition Path of Cyber Command & STA & HCI

The following are application scenarios for Security Integration (aSI):

1. When Cyber Command detects ransomware, it calls the storage snapshot API of HCI to provide fallback protection.
2. When Cyber Command detects VMs infected with viruses such as crypto mining viruses, it works together with HCI to shut down or suspend the VMs to prevent them from consuming the entire platform's performance.
3. When Cyber Command detects VMs in a botnet, it automatically isolates infected VMs using the distributed firewall of HCI.

7.4.5 Data Encryption

When users need to encrypt and protect data stored in the cloud for security or regulatory compliance, they can utilize the data encryption feature of Sangfor HCI. By integrating with the key management service (KMS) of SCP, users do not need to build, maintain, and protect their own key management infrastructure and encryption implementation mechanism. The data encryption feature of Sangfor HCI ensures the privacy and autonomy of cloud data. It supports encryption, decryption, and integrity verification of sensitive data on both HCI and SCP. Additionally, it enables tenant hardware-level key management and the encryption of VM and NFV images. SCP supports integration with third-party password service platforms through OpenAPI. Users can use an encryption card to provide high-performance data encryption and decryption services.

The encryption card is a hardware security module (HSM) specifically designed to store and manage keys, certificates, and other sensitive data. HSM employs physical features and encryption algorithms, such as the anti-counterfeiting identifier, random number generator, symmetric encryption algorithm, and asymmetric encryption algorithm, to ensure the security of keys and data. VMs can use HSM to safeguard sensitive data, such as encryption keys and digital certificates. Virtualization management software (such as VMware vSphere, KVM, and Hyper-V) typically provides interfaces and tools to map HSM to VMs,

thereby achieving VM security protection. Encryption cards offer two modes of service on HCI:

- Passthrough Mode

In this mode, HSM is directly connected to the PCIe bus of the host, allowing the host to directly access the hardware resources of HSM, such as the encryption engine and random number generator. Passthrough mode ensures optimal performance and security by bypassing the virtualization layer, resulting in reduced latency and security risks.

- VF Mode

In VF mode, HSM is virtualized into multiple virtual devices, each of which can be assigned to a VM. SR-IOV-VF enhances the performance and efficiency of VMs by directly mapping the data traffic of VMs to the virtual devices of HSM, thereby eliminating the need for virtualization layer involvement.

7.5 Disaster Recovery

7.5.1 Data Protection Overview

SCP provides the **Data Protection Overview** page for administrators to view the overall statuses of scheduled backup, CDP, and remote disaster recovery services.




Figure 69: Data Protection Overview

7.5.2 Scheduled Backup

7.5.2.1 Introduction

SCP supports tenants to customize scheduled backup policies for VMs based on the importance of running business systems to protect their business data.

7.5.2.2 Application Scenarios

Only the reliability measures of business servers (such as active-active storage, multiple replicas, HA, and automatic system restart in case of failures) are not enough to address business data failures such as deletion of data by mistake or ransomware attacks within a VM, because such failures directly damage the data source and thereby damage the replicated data as well. For example, the recent surge in ransomware attacks has resulted in internal data corruption within business systems, rendering multiple replicas ineffective as each replica is corrupted.

7.5.2.3 Supported Features

1. Quick Backup

SCP provides a quick backup mechanism, which uses an initial full backup, followed by incremental backups, along with the bitmap technology (which is used to mark the changed locations in the QCOW2 file, also known as dirty data locations) for quick backup.

The backup operation is very simple as users can add a backup policy compliant with business characteristics in four steps as prompted in the wizard.

The automatic backup period can be flexibly configured by week, day, or hour based on the actual business requirements. Automatic cleanup is available to save backup repository space to the maximum. In addition, no plugins are required, making it simple to use.

2. Quick Recovery

In the event of a business system failure, the VM can be recovered to the most recent point of normal business operation by querying the backup files if data backup has been enabled. Generally, perform the following operations for backup and recovery: migrate the backup data to the running location and start the VM at the running location to resume business services after the data migration is completed. However, substantial business data volume and long migration duration may result in a large recovery time objective (RTO). As a result, SCP further optimizes the backup and recovery features to address this issue.

3. Periodic Full Backup, Compression, and Archiving

SCP provides the initial full backup, followed by incremental backups, by default. However, to meet certain compliance requirements for periodic full backup and secondary data archiving in addition to incremental backups, SCP also provides solutions for periodic full backups and compression and archiving.

7.5.3 CDP

7.5.3.1 Introduction

SCP provides tenants with the continuous data protection (CDP) feature for their business VMs. CDP is an advanced backup technology for I/O in seconds. Data protection is implemented by recording every data modification made by customers and restoring data with zero data loss.

7.5.3.2 Application Scenarios

Generally, common hour-level backup is performed, which indicates that at least 1-hour data loss is inevitable in the event of a business failure. Such data loss is unacceptable for some critical businesses. Some critical business systems require a more granular backup approach, with the recovery point objective (RPO) measured in minutes or seconds. CDP is therefore implemented to meet this requirement.

7.5.3.3 Supported Features

To optimize CDP technology for enhanced data protection, SCP incorporates the CDP module at the QCOW2 file layer of the VM instead of using the traditional approach of embedding CDP software as an agent at the guest OS layer. SCP provides a cost-effective, easy-to-deploy, and high-value CDP solution that is more suitable for VM environments. Its main features include:

1. CDP Backup

Supports minute-level or even second-level CDP for VMs on which business applications run; supports quick browsing backup files and log files, making it easier to quickly obtain data files from the CDP backups.

2. Backup Consistency

Supports performing CDP backups for VM disks and consistency checks on the data stored on them to ensure that each recovery point (RP) is valid.

3. Quick Recovery

In the scenario where a business failure occurs (such as a virus attack and deletion by mistake) and CDP backup files are required for VM recovery, users can find backup files according to the timestamps. Through incremental backups and innovative design of I/O repositories, quick browsing of backup files is supported for file-level and business-level data recovery.

7.5.4 HotAdd-based Agentless Backup

7.5.4.1 Introduction

HotAdd is a capability provided by the cloud platform that supports adding devices, including CPUs, memory, NICs, and disks, while the VM is running.

HotAdd-based agentless backup is a transport method for agentless backup. A third-party backup proxy app is deployed on a DR/backup proxy VM and accesses the disk data of application VMs by using the HotAdd capability, thereby enabling backup and recovery of VMs on the cloud platform.

7.5.4.2 Supported Features

1. DR/backup proxy VM management: Supports creating, configuring, and managing DR/backup proxy VMs for HotAdd-based agentless backup operations.
2. Disk hot add: Supports hot adding linked clones of application VM disk snapshots to DR/backup proxy VMs.
3. Disk data access: In the guest OS of the DR/backup proxy VM, the backup proxy app can directly read data from the hot-added disk to complete the backup.
4. Disk hot removal: Supports removing the hot-added disk from the DR/backup proxy VM after backup is complete.
5. Backup and recovery: Supports full and incremental backups for application VMs and recovery to new VMs and source VMs.

7.5.5 Remote Disaster Recovery

7.5.5.1 Introduction

SCP provides a VM disaster recovery (DR) solution of local backup and remote recovery with RPOs from one second to one week. The local data center (primary site) provides second-level CDP and transmits backup data to the remote data center (secondary site) completely, securely, and quickly through encrypted transmission, compressed transmission, and interrupted transmission. When a VM at the primary site fails, users can quickly recover it from local backups or to the secondary site, and then migrate its business services back to the primary site.

7.5.5.2 Application Scenarios

When a running host in a cluster fails, the VMs on it can be recovered on a suitable host in the cluster through the HA mechanism to minimize business

service interruptions. When a business system encounters a logical failure (such as a virus attack and deletion by mistake), it can be recovered through CDP and backups to reduce data loss as much as possible. However, when a data center fails (due to earthquake, fire, etc.), neither HA, multi-replica, or local backups can be used for business recovery, and remote backups retain only the data and exclude the working environment. In this case, recovering the entire data center is time-consuming and can cause huge losses to customers.

To address the failure of an entire data center, SCP provides a remote disaster recovery solution that can quickly recover business services, reduce their interruptions, and protect them with multiple replicas, HA, and backup systems for comprehensive security and reliability.

In response to the explicit regulations and needs of some industries (such as banking, education, and healthcare) for disaster recovery building, SCP is continuously innovating to provide better services for customers.

7.5.5.3 Supported Features

1. Provides second-level local CDP to recover failed VMs from local backup data and remote VM recovery with different RPOs (1 second, 10 seconds, 10 minutes, 30 minutes, 1 hour, 2 hours, 4 hours, 8 hours, 12 hours, 1 day, 2 days, and 1 week). Both primary and secondary sites must be Sangfor's enterprise cloud platforms.
2. Supports bandwidth settings based on RPO needs, and provides interrupted transmission and compressed transmission to help users save bandwidth resources and costs.
3. Supports one-click recovery to the secondary site with an RTO of 10 minutes and early configuration of the network for recovery.

A disaster recovery process is divided into two main phases: real-time data backup and business failure recovery. The backup phase includes local backup and remote synchronization (synchronizing data to the secondary site), and the recovery phase includes local recovery (to the primary site), remote recovery (to the secondary site), and business migration back to the primary site. The

following figure illustrates them except for disaster recovery initialization and test.




Figure 70: Remote Disaster Recovery Process

The seven steps of disaster recovery are explained in detail below, in order of disaster recovery initialization, local backup, remote data synchronization, local recovery, remote recovery, business migration, and disaster recovery test.

Disaster Recovery Initialization:

1. Create a disaster recovery task and synchronize the configuration at the primary and secondary sites.
2. Automatically create placeholder VMs and synchronize VM configurations at the secondary site.
3. Manually pre-configure (recommended) the network topology at the secondary site.
4. Configure a link between the primary and secondary sites. The inter-site communication interface should preferably be independent and not multiplexed with the business or cluster interface to prevent the disaster recovery process from preempting business traffic. Both interfaces between clusters and within clusters must be configured.
5. Configure a protection group and set the startup priority upon recovery for VMs in the group.

6. Create a recovery plan.

Local Backup and Remote Data Synchronization:

A local backup policy is configured according to RPO needs and supports normal backup or CDP, corresponding to the following two asynchronous data synchronization methods:

- Local scheduled backup with remote backup file synchronization for RPOs in hours, days, or weeks
 1. The primary site takes VM backups periodically (such as every hour) based on the specified RPO.
 2. When a backup is completed, a task is triggered to transfer the backup file to the secondary site.
 3. Supports interrupted transmission, compressed transmission, encrypted transmission, compared transmission, load balancing, and dynamic traffic control.
 4. Supports setting RPO warning and alert thresholds when configuring a remote protection policy to prevent frequent RPO alerts within a short time due to network jitter or insufficient bandwidth. The warning information is displayed on the DR status and data overview page, and the alert information is shown on the alert page, which can be synchronized to users in real time through email and SMS.
 5. Supports creating replica seeds from backup files, encrypting and packing them, and transferring them through external disks.
 6. Supports archiving backup files by month.
 - Local CDP backup with remote I/O log synchronization for RPOs in seconds or minutes
 1. The primary site periodically generates RPs (i.e., I/O logs) based on the specified RPO and backup files based on the specified schedule.
 2. The I/O log data generated by the CDP at the primary site is synchronized in real time to the secondary site by replay, as well as backup files.

3. Supports interrupted transmission, compressed transmission, encrypted transmission, compared transmission, load balancing, dynamic traffic control, and degrading treatment.
4. Supports setting RPO warning and alert thresholds when configuring a remote protection policy to prevent frequent RPO alerts within a short time due to network jitter or insufficient bandwidth. The warning information is displayed on the DR status and data overview page, and the alert information is shown on the alert page, which can be synchronized to users in real time through email and SMS.

Local Recovery:

When business services on some VMs at the primary site are unavailable due to a failure, these VMs can be recovered directly from local backups to quickly get business services running again.

Remote Recovery:

Remote recovery is used for disaster recovery when the entire data center at the primary site fails and business services cannot be recovered locally and for scheduled recovery without failures.

- **Disaster Recovery:** After a disaster, unsynchronized data may be lost. Business services can be recovered to the latest RP at the secondary site by default to prioritize them to come online as soon as possible.
- **Scheduled Recovery:** Cross-site business recovery planned by customers requires the primary site to be online for scenarios such as real DR tests and upgrades during VM shutdown.

Steps:

1. Manually shut down all VMs at the primary site.
2. Synchronize unsynchronized data to the secondary site.
3. Recover the VMs at the secondary site without loss of business data.

Business Migration:

The primary site is the production environment used by customers to run

business services and is often better equipped than the secondary site. Therefore, business data including full image files, VM configuration files, incremental backup files generated periodically, and real-time I/O logs need to be sequentially migrated back to the primary site after recovery.

The business migration back to the primary site requires all VMs at the secondary site to be shut down; otherwise, only scheduled backups and synchronization of unsynchronized data between the primary and secondary sites will be performed automatically. Before migration, a data backup and synchronization of unsynchronized data between the two sites is performed. Depending on the primary site failure and secondary site data, there are two types of migrations as follows:

- **Full Migration:** Used when the data center at the primary site has lost backup files in the event of a failure, and the secondary site has full image files and incremental backup files synchronized from the primary site. In this scenario, the backup chains of the two sites are inconsistent.
- **Incremental Migration:** Used when the data center at the primary site has no backup files lost in the event of a failure. In this scenario, the backup chains of the two sites are consistent.

Disaster Recovery Test:

There are real and simulated DR tests. A real DR test is performed at the final stage before business services come online to guarantee that the entire recovery and migration process is available. A simulated DR test is a disaster recovery simulation that includes the five of the steps introduced above: local backup and remote data synchronization, local recovery, remote recovery (disaster recovery and scheduled recovery), and business migration.

The requirements that must be met in a DR test are as follows:

1. Newly generated data during the backup phase can be synchronized to the secondary site based on the configured RPO.
2. VMs during the recovery phase can be individually and quickly recovered from local backups.

3. VM data is not lost during disaster recovery or scheduled recovery.
4. VM data is not lost during business migration back to the primary site.
5. After recovery to the secondary site, VMs can be quickly connected to the network based on the topology configured in advance to recover business services.

The simulated DR test is mainly used for the daily maintenance of customers to verify the reliability and availability of backup files at any time. It consists of six steps: creating a DR policy and executing it, creating a protection group and a recovery plan, performing the DR test and finishing it.

Test VMs will be recovered at the secondary site through a simulated DR test. Users can only verify but not create business services on them. The test VMs automatically create test networks based on the placeholder VMs at the secondary site (requires manual configuration if creation fails). The test networks cannot be connected to the network used for recovery at the secondary site; otherwise, the recovery network may be affected. Users can finish the test after verifying the correctness of the business services, and the system will automatically remove the automatically created test VMs and network devices. The simulated DR test does not perform business migration back to the primary site.

7.5.6 Cloud Disaster Recovery

7.5.6.1 Introduction

With Managed Cloud Services, Sangfor provides the cloud disaster recovery service with second-level RPO and minute-level RTO, which does not require building a disaster recovery center. On-premises applications can be backed up to the cloud to support offsite disaster recovery, guaranteeing data security and business continuity. Additionally, SCP enables easy on-premises deployment, automatic cloud resource management, and centralized console control.

7.5.6.2 Application Scenarios

Disaster recovery construction is required to meet on-premises key business continuity and data security compliance requirements. However, due to the limitations of money, technology, and O&M, it is difficult and time-consuming to build a disaster recovery center. Sangfor Cloud provides the disaster recovery service of homogeneous hybrid cloud solutions, which can reuse the resources on Managed Cloud Services to quickly build a secondary disaster recovery center.

7.5.6.3 Supported Features

1. Unified Management: Offers a unified asset view, resource search, and value display of on-premises and cloud applications.
2. Cloud Interconnect: Offers a cloud disaster recovery center that is connected to on-premises applications, interconnecting layer 2 and 3 networks to lower cost, latency, and difficulty.
3. Rapid Capability Acquisition: Offers rapid access to trial resources of Managed Cloud Services and value guides to understand the corresponding business effects to develop on-premises capabilities.
4. High Reliability: Offers up-to-second RPO and replicates data in real time, enabling VM recovery to the cloud in minutes.
5. Fast Verification: Offers multiple recovery points and a disaster recovery test that can be done within 10 minutes.

7.6 Memory Tiering Service

7.6.1 Background

The increase in memory prices leads to the stagnation of existing projects and difficulties in promoting new projects. This has become a core challenge in frontline operations and will continue to impact order revenue and profit margins.

Industry research shows that the actual memory usage in data centers has remained low for a long time, with a large amount of cold data occupying high-performance memory space. According to VMware statistics, the active memory ratio is only about 50%.

A low-cost delivery solution can be implemented using the memory tiering module to reduce hardware procurement while preserving software revenue and profit margins.

7.6.2 Technical Principles

Based on memory tiering technology, memory overcommitment is achieved by dividing memory into a DRAM high-performance tier (Tier 0) and an NVMe SSD expansion tier (Tier 1). Through hot and cold data identification and migration, memory can be expanded with minimal performance loss.


This method can break through memory expansion bottlenecks and significantly reduce its cost.

Tier 0: DRAM, which is expensive and has a small capacity. It is used to store hot data.


Tier 1: NVMe. Its latency differs by three orders of magnitude from DRAM, but it is inexpensive and has a large capacity. It is used to store cold data.

Memory tiering is transparent to VMs, for which it is a unified memory address space.


The memory tiering system continuously monitors the VM's memory accesses, identifying hot and cold data. When memory usage is high, the cold data is migrated from the DRAM tier to the NVMe tier (demotion). When the VM requires accessing the migrated cold data, the tiering system migrates the cold data from the NVMe tier back to the DRAM tier (promotion).




Technical Architecture:




7.6.2.1 Configuring Node Memory Tiering




Step	Details	Remarks
Check memory	The total committed-RAM contains only physical committed-RAM, excluding tiered committed-RAM. a. Total committed-RAM - Kernel group memory - 5% * Total committed-RAM > Total memory allocated for running VMs with memory tiering disabled. b. Total committed-RAM - Kernel group memory + min(Physical memory * Tiering ratio, Tiered disk size) ≥ Total memory allocated for running VMs.	a. Required when configuring memory tiering for the first time. Not required when disabling memory tiering, increasing or decreasing the memory tiering ratio, or adding or removing disks. b. Required when configuring memory tiering for the first time or decreasing the memory tiering ratio. Not required when adding disks or increasing the memory tiering ratio.
Mark the disk as in-use	/cfs/disk.info diskused = 5 This prevents the disk from being used for other purposes (such as datastore creation) during the addition process.	Supported disk types: NVMe SSDs (recommended) and SATA SSDs (supported at the code level but not publicly advertised). Number of disks: A node supports multiple disks, and both types of SSDs can be used at the same time.
Format	Format all partitions as swap space.	
Run the swapon command	Supports configuring multiple disks at the same time, including NVMe SSDs and SATA SSDs. Swap priority: NVMe SSD (10) > SATA SSD (0) > System disk (-2)	If multiple disks are selected as tiered disks at the same time, some disks may succeed while others fail.
Update memory tiering configuration	<pre>{ "Node" : { "disks" : { "Disk ID" : { "Disk capacity" : "960197124096", "Protocol type" : "NVMe" } } }, </pre>	<pre>{ "host-b0a3f202a0d4" : { "disks" : { "SNVMe_SAMSUNG_MZQL2960HCJR-00B7C_S63WNE0T102530_b0a3f202a0d4" : { "disksize" : "960197124096", </pre>

	<pre>"Tiering ratio" : 100 } }</pre>	<pre>"interface_type" : "NVME" } }, "tier_percent" : 100 } }</pre>
Update memory overcommitment settings	Set the node's memory overcommitment ratio to 100%.	If the operation fails, the task description will indicate the failure, but the mounted disk operations and the tiering configuration will not be rolled back.
Disable transparent huge page	Disable the transparent huge page, including specifying a flag bit to prevent the transparent huge page from being enabled after reboot, and disabling the currently active transparent huge page.	Check the mount parameters for qemu_vhost_user. If the huge value is never, the transparent huge page is disabled. The figure below shows it as enabled.  <pre>Sangfor:aSV/host-b4055dba394a(asv-controller) tmpfs on /mnt/shared/qemu_vhost_user type tmpfs Sangfor:aSV/host-b4055dba394a(asv-controller)</pre>

7.6.2.2 Clearing Node Memory Tiering Configuration




Step	Details	Remarks
Check memory	Total committed-RAM - Kernel group memory $\geq$ Memory allocated for running VMs.	If the check fails, an error message will be displayed prompting the user to shut down or migrate the VM.
Clear memory tiering configuration	Clear the memory tiering configuration first.	Running the swapoff command may take a long time. If the configuration is cleared first, the upper layer is unaware of the existence

Step	Details	Remarks
		of tiered memory. This prevents VMs from occupying tiered memory during the power-on.
Run the swapoff command	1. During the swapoff process, the task displays the remaining swap memory that has not been swapped to physical memory in real time. 2. Timeout = Used swap space/10 GB/min * 4 (quadrupled) 3. If multiple disks are used as tiered memory, some disks may succeed while others fail.	The task description will indicate the failed disks. The tiering configuration corresponding to the failed disks will be written back.
Clear disk usage mark	Clear the usage mark for disks on which the swapoff command is run successfully.	If the update fails, the task description will indicate the failure.
Enable transparent huge page	Enable the transparent huge page, including removing the mark that prevents transparent huge page from being enabled after reboot, and enabling the current transparent huge page.	Check the mount parameters for qemu_vhost_user. If the huge value is advise, the transparent huge page is enabled, as shown in the figure below. 

7.6.2.3 Memory Calculation Algorithm Description

Version	HCI 6.11.3
---------	------------




1.1.1. Enabling/Disabling Memory Tiering for VM Creation/Cloning

The business process remains unchanged, but the API input parameters and parameter handling must be adapted to support the field.

Default: Disabled

7.6.2.4 Enabling/Disabling Memory Tiering for VMs




Step	Details	Remarks
Check before configuration	VMs with encryption cards or with huge-page memory, NIC passthrough, or GPU passthrough enabled do not support memory tiering.	Memory is locked after such VM start, so memory tiering is not supported.

Step	Details	Remarks
Check the VM status	If the VM is shut down, update the VM configuration. If the VM is running, go to the next step.	
Check the status of the memory tiering feature for the VM	If the feature is disabled, update the VM configuration to enable it. If the feature is enabled, go to the next step to disable it.	
Check whether the node where the VM is running has memory tiering enabled	If not, update cgroup (Set memory.swap.max from 0 to -1). If yes, go to the next step.	If the memory tiering feature is not configured for the node, configuring the memory tiering for the VM will not take effect, and no error will be reported.
Check node memory	Total committed-RAM - Kernel group memory - (Total memory allocated for running VMs with memory tiering disabled + 5% * Total committed-RAM) > Memory allocated for VMs on which memory tiering is being disabled. If the memory is insufficient, error will be reported. If the memory is sufficient, go to the next step. Note: Bulk editing scenario 1. Group the VMs by the nodes on which they are running. 2. For multiple nodes, partial success or failure is allowed.	A 5% buffer memory is reserved for services, and a portion of physical memory will also be locked for network forwarding plane.
Configure cgroup	memory.swap.max 0: Swapping is disabled. -1: Swapping is enabled. Scenario 1: When memory tiering is enabled on nodes and VMs, the value of -1 indicates that	

Step	Details	Remarks
	swapping is enabled. Scenario 2: When memory tiering is enabled on nodes and disabled on VMs, the value of 0 indicates that swapping is disabled. Scenario 3: When memory tiering is disabled on nodes and enabled on VMs, the value of -1 indicates that swapping is enabled, which means that the swap memory mounted on the native system is used. Scenario 4: When memory tiering is disabled on nodes and VMs, the value of -1 indicates that swapping is enabled, which means that the swap memory mounted on the native system is used.	

7.6.2.5 Powering on VM



Step	Details				Remarks
		Node	VM	Algorithm	
Schedule	Filter				

Step	Details				Remarks
		Disabled	Disabled	Memory tiering does not take effect. Assessment is provided based on the previous memory overcommitment algorithm.	
		Disabled	Enabled	Memory tiering does not take effect. Assessment is provided based on the previous memory overcommitment algorithm.	
		Enabled	Enabled	 Total committed-RAM + Tiered memory - Memory allocated for VMs with memory tiering enabled or disabled > Memory allocated for VMs to be started	
		Enabled	Disabled	 Total committed-RAM - Memory allocated for	

Step	Details				Remarks
				VMs with memory tiering disabled - Total committed-RAM * 5% > Memory allocated for VMs to be started Total committed-RAM + Tiered memory - Memory allocated for VMs with memory tiering enabled or disabled > Memory allocated for VMs to be started	
	Rate	Switch matching rating plugin (New): 10 Node remaining physical memory plugin (New): 90 Node load plugin (Existing, unchanged): 100 Other plugins: Remain unchanged. Effect: 1. VMs with memory tiering enabled are preferentially scheduled to nodes with memory tiering enabled. 2. VMs with memory tiering disabled are preferentially scheduled to nodes with memory tiering disabled. 3. Physical memory of the nodes in the cluster is preferentially used.			
Power on VM		The original logic remains unchanged. The memory tiering configuration is not associated with the startup parameters of the VM.			
Create cgroup		Create a VM ID group under normal_compute_group or important_compute_group based on the VM's importance level. Based on the status of the memory tiering feature for VMs and nodes on which they are running, configure memory.swap.max to control whether swapping is enabled. Scenario 1: When memory tiering is enabled on nodes and VMs, the value of -1 indicates that swapping is enabled. Scenario 2: When memory tiering is enabled on nodes and disabled			

Step	Details	Remarks
	on VMs, the value of 0 indicates that swapping is disabled. Scenario 3: When memory tiering is disabled on nodes and enabled on VMs, the value of -1 indicates that swapping is enabled, which means that the swap memory mounted on the native system is used. Scenario 4: When memory tiering is disabled on nodes and VMs, the value of -1 indicates that swapping is enabled, which means that the swap memory mounted on the native system is used.	
Notify the kernel to collect active memory statistics	The kernel needs to be notified to collect active memory statistics.	

7.6.3 Identifying Memory Hot and Cold Pages


The feature is optimized based on the kernel's Data Access Monitoring (DAMON) technology. The address spaces of VMs are dynamically divided into regions, and regular sampling is performed to record the access frequency.

The sampling result represents the final result of the corresponding region.

Memory with the highest access frequency over a period of time is considered hot memory (active memory), while the rest are considered cold memory.

Region sampling enables a relatively accurate assessment of memory activity and keeps CPU overhead low.

DAMON technology uses random and sparse sampling to infer the overall access frequency of memory on each region. The process can be divided into the following steps:



Each sampling segment undergoes iterative steps, including random and sparse sampling, merging of similar and neighboring regions, and splitting of existing regions. In most cases, the DAMON technology can converge to a relatively accurate and low-overhead steady status.

7.6.4 Migrating Hot and Cold Pages

When memory usage increases, cold pages can be proactively identified and migrated to Tier 1 in advance, preventing VMs from directly triggering memory reclaiming. This reduces performance decrease caused by memory migration and ensures stable system operation even under high load.

Technical Implementation:

VM memory is managed through the memory cgroup system. By configuring the `memory.high_async_ratio` parameter, asynchronous memory reclaiming can be triggered when the cgroup memory usage reaches the configured threshold.

When a VM accesses a cold page on Tier 1, a page fault mechanism is triggered, migrating the cold page back to the DRAM tier.

7.6.5 Monitoring Hot and Cold Data in VM Memory

The system monitors the hot pages and the frequency of hot and cold data

swapping in VM memory in real time to assess the health status of the tiering system. If the data swapping is too frequent with a massive volume over a period of time, alerts will be triggered, and optimization policies (such as live migration) will be provided.

7.6.6 Disabling Memory Tiering for VMs

After memory tiering is disabled for a VM, the VM's memory will be reserved in the DRAM tier. However, memory that has already been moved to Tier 1 will not be immediately migrated back to the DRAM tier. Instead, the memory will be migrated as required when accessed.

Technical Implementation:

- Memory Reservation

The memory is not truly reserved, but rather swapping is disabled. When memory is required, memory preemption will be performed on another VM with low priority. The preempted memory will be locked.

- Migration as Required

This is triggered by page faults.

Provide an interface for immediate migration in the aCLI tool, allowing memory to be migrated back to the DRAM tier from Tier 1 immediately when necessary.

Two methods are available for technical implementation:

Method 1:

The VM memory is mapped from files in the tmpfs system. When immediate migration is required, read through the entire file using the dd command.

When using the dd command, care must be taken not to turn a sparse file into a pre-allocated one.

However, immediate migration is not supported for VMs without vmTools installed because such VMs use anonymous memory, which cannot trigger migration via the dd command.

(There are other methods available, but considering that VMs without vmTools installed are rare, support is not provided at this time.)

Method 2:

The Euler kernel adds a `force_swapin` interface to `cgroup`, which is used to forcibly swap VM memory from Tier 1 back to the DRAM tier.

However, the test indicates that it does not work well with transparent huge pages (it swaps in only the first 4K subpage of a transparent huge page). This prevents memory from being swapped as expected because transparent huge pages are enabled on VMs by default.

Using the `force_swapin` interface is much faster than running the `dd` command. However, due to its limitations in the main scenarios, the `dd` command is recommended.

Application of active memory in VM management

- Startup
- DRS
- HA

In the VM scheduling scenarios described above, when selecting an optimal node to run a VM, the node's active memory is taken into consideration. The node with more available space after deducting active memory is preferred to minimize the interference with the performance of already running VMs.

However, if active memory is excessive and physical memory is insufficient, memory preemption among VMs will inevitably lead to performance loss.

7.6.7 Replacing NVMe SSDs

When NVMe SSDs reach the end of their lifetimes or fail, replacement in time is required.

To replace an NVMe SSD, ensure that services on the node are fully evacuated, or that the memory required for services does not exceed the capacity of the DRAM tier.

After services are evacuated, disable memory tiering on the node, replace the NVMe SSD, and enable memory tiering again.

7.6.8 Deployment Assessment

7.6.8.1 Hybrid Deployment Mode

Nodes with memory tiering enabled and disabled can be deployed in the same cluster.

Different types of VMs can be scheduled on all nodes without restriction.

Users can achieve their special scheduling requirements by configuring the preferred node where VMs run, affinity, and other capabilities.

7.6.8.2 Workload Assessment and Tiering Ratio

Active memory is mainly used to assess whether the workload is suitable for memory tiering.

The default capacity ratio between the DRAM tier and the NVMe tier is 1:1. In this case, active memory cannot exceed 50% of the total physical memory before enabling memory tiering.

To change the capacity ratio between the DRAM tier and the NVMe tier, the active memory ratio must meet the corresponding requirements. For example, with a 1:4 ratio, the active memory ratio cannot exceed 20% before memory tiering is enabled.

Simply put, regardless of the tiering ratio, the total active memory for all services cannot exceed the capacity of the DRAM tier (the overhead of the HCI platform must also be deducted).

If the active memory ratio is generally below 50% but occasionally increases above 50% due to service bursts during certain periods, enabling memory tiering may result in performance degradation during service bursts.

7.6.8.3 Server Hardware Configuration

Users can choose the server and hardware configuration based on their own requirements. However, the following recommendations are provided for optimal performance and reliability:

- Memory modules are evenly distributed across nodes.

Based on the test results, 32 GB memory modules are recommended. Enabling more memory channels can improve performance.

- Supports PCIe Gen4 and later.
- Supports NVMe SSD hot-swapping with sufficient interfaces.

7.6.8.4 Selecting NVMe SSD

As an extension of the DRAM tier, the NVMe SSD requires high performance and reliability.

Performance: 4K random read/write IOPS must be at least 100,000.

Endurance: At least 3 DWPD and 7300 TBW

Enterprise read/write hybrid disks are recommended because the NVMe SSD is used as a memory extension, which may lead to frequent read and write operations.

7.6.8.5 NVMe SSD Failure Scenarios

The failure rate of NVMe SSDs is relatively low. If a failure occurs, the corresponding impact scope depends on whether the VM actually uses the Tier 1 space.

The Tier 1 space is primarily used to store data that is infrequently accessed by services. When the NVMe SSD fails, the data stored in the Tier 1 space will become inaccessible:

If the Tier 1 space is not used, the service will not be affected.

If some cold data of the VM is stored in the Tier 1 space, accessing the corresponding data will trigger an abnormal crash. The VM will then automatically restart and resume operation on another available node via the HA mechanism.

Note that whether a VM immediately encounters errors depends on whether it accesses data on the faulty NVMe. If the service never accesses the relevant cold data, the VM may continue to run normally.

8 PaaS-Based Services

8.1 Database

8.1.1 Introduction

As a data storage and management software, a database is essential to an enterprise's IT architecture. With the development and popularization of cloud computing, more and more enterprises migrate their data to the cloud.

Relational Database Service (RDS) is a ready-to-use, stable, reliable, elastic, and scalable online database service with multiple security measures and a perfect performance monitoring system. It provides professional database backup, recovery, and optimization solutions, enabling customers to focus on application and business development.

Sangfor, with its technology and experience in the field of cloud computing and conforming to the trend of Information Technology Application Innovation, has launched the Sangfor Database Management Platform (DMP), which focuses on solutions for various database scenarios such as deployment, usage, O&M in the era of cloud computing.

Based on SCP, Sangfor DMP provides database O&M capabilities and various database services, including RDS, migration, backup, automated O&M, and operation services, to enhance database reliability and simplify O&M to meet customers' needs for servicing and centralized management of databases.




Figure 71: DMP Features

8.1.2 Application Scenarios

DMP acts as a PaaS capability on SCP to provide supporting capabilities such as full lifecycle management, business reliability guarantee, data backup and protection, monitoring and operation, and maintenance of mainstream relational databases in the following typical scenarios:

1. Database Hosting Platform Creation and Expansion




Figure 72: Database Hosting Platform

As a critical application of business hosting, the database requires its hosting platform to meet a high level of reliability, standardized delivery, easy expansion, and high-performance of databases, all of which can be provided by DMP.

Key Values:

- **Quick Delivery:** Supports racking an all-in-one machine within 3 hours and deploying and putting a database online within 1 hour.
- **Active High Reliability:** Provides integrated high-reliability guarantee and significantly reduces the O&M burden through various types of active monitoring, checking, and backup.
- **Excellent Performance:** Supports the smooth expansion of the performance and capacity of the underlying storage and adjustment and optimization of the database configuration.

2. Database Hosting Platform Transformation and Replacement




Figure 73: Database Hosting Platform Transformation and Replacement

The traditional database construction solution is to purchase physical machines or minicomputers for hosting databases, which requires customers to configure the hardware, operating system, database, and network one by one and invest human resources to continuously maintain them in daily operation. However, due to the scalability difficulty, poor reliability, and high operation and maintenance cost, physical machines and minicomputers can no longer support customers' business development needs. DMP provides a database resource pool construction solution that enables database administrators to perform unified O&M management of databases to reduce the O&M burden significantly and supports standardized database delivery and data transmission service.

Key Values:

- **Elastic Expansion:** Deployed on SCP and supports elastic expansion of underlying hardware without impacting business services.
- **Smooth Transformation:** Supports smooth transformation to new platforms and database migration to other platforms through the data transmission service.
- **Simple O&M:** Supports unified management of the full lifecycle of the created databases.

3. Database Security Protection




Figure 74: Database Security Protection

Database administrators' primary duty is to perform backups for databases against unexpected events that may threaten the data security of databases, such as hacker virus intrusion, employee resignation after deleting databases, and natural disasters. Common database backup methods generally require third-party software or scripts with different maintenance methods for full and incremental backups, and it is difficult, even impossible, to verify the validity of the backup data in real time. DMP provides database transaction backup capability. Based on the specific backup policy configured by the administrator for a database, DMP can automatically perform scheduled full and incremental backups, automatically clean up expired backups, and create a brand new database directly from any backup file to verify its validity at any time.

Key Values:

- **Backup Visualization:** Supports backup result viewing, backup task status management, information statistics, and task failure alerts for visualization and management of backup validity.
- **Verifiable Backups:** Supports one-click recovery of backup files to new single instance databases or RAC databases so that the validity of backups can be quickly verified in daily operation.

- **Enhanced Data Security:** Strengthened continuously for data security.

8.1.3 Supported Features

8.1.3.1 Automated Deployment

- MySQL Database

DMP supports the automated deployment of database clusters, enabling the database deployment with simple configurations to facilitate business services to come online. At present, all the deployments and subsequent failovers of databases on DMP are implemented using the task flow mechanism by splitting the tasks.

- Oracle Database

Users can perform a one-click deployment by only specifying the necessary parameters, such as description, specifications, Oracle database configuration, and network to be connected.

Oracle database lifecycle uses the capabilities provided by SCP for image-based deployment, resource pool disk management, and IP address management to quickly create a VM that runs the Oracle database and then connects to the VM through the Agent interface to deploy a database automatically. The deployment processes are shown as follows:




Figure 75: Single Instance and RAC Database Deployment Processes

An SCP interface is first called to create a VM configured based on the parameters set by the user. After the shared disks applied in the above figure are mounted to the VM, DMP connects to it through the Agent and executes an automated Oracle database deployment through a script.

1. Set system parameters.
2. Configure SSH to communicate between nodes (only required for RAC architectures).
3. Install GRID (only required for RAC architectures).
4. Create an ASM disk group (only required for RAC architectures).
5. Install the database software.
6. Create a database.
7. Configure the monitoring feature.
8. Configure the listening port.
9. Optimize performance parameters.
 - SQL Server Database

SQL Server is a Relational Database Service (RDS) released by Microsoft.

During the deployment without a domain controller, DMP performs a series of configurations and parameter initialization for the operating system and SQL Server program to deliver a single instance database (a relational database running the Windows Server 2016 operating system) or a database cluster (including at least two databases).

In addition to the backup and monitoring features necessary for RDS, DMP implements the synchronization of SQL Server database objects (such as login names and orphaned databases) between nodes in the cluster to ensure compatibility with SQL Server features and high availability.

8.1.3.2 Lifecycle Management

DMP supports the lifecycle management of database clusters, including

enabling, disabling, restarting, deleting, and modifying configurations. Different operations can be performed on clusters in corresponding statuses.

8.1.3.3 Self-Healing

DMP supports the self-healing of single instance and primary-replica MySQL databases as well as single instance Oracle databases and Oracle RAC RDS clusters with the following processes:




Figure 76: MySQL Database Self-Healing Process




Figure 77: Oracle Database Self-Healing Process

The self-healing process can be divided into detection, analysis, and self-healing phases. The detection (or node monitoring) phase collects service

status and data synchronization information of the monitored node, the analysis phase defines the type of failure according to the node information and the database architecture and type, and the self-healing phase adopts a proper solution to recover the cluster based on the failure type and practices (DBA and O&M experience).

8.1.3.4 Data Protection

DMP provides transaction consistency backup of databases, which significantly improves database security. Its backup management modules run independently, manage database backups, and maintain database tables, enabling features including manual and scheduled backups, auto backup cleanup, backup deletion, and recovery.

DMP supports data backup and recovery, as well as CDP, for MySQL, Oracle, and SQL Server databases.

8.1.3.5 Oracle Database Disaster Recovery

DMP provides disaster recovery protection groups (DR protection groups) based on Oracle Data Guard, including its creation and deletion, adding and removing standby databases to and from it, switchover and failover, its list, and listening services.

DMP's Data Guard protection groups (DR protection groups on DMP) can only contain active Oracle databases of version 11.2.0.4.0. DR protection groups include physical standby databases, which copy data online through RMAN. These groups are managed with the Data Guard broker, which supports manual failover.

8.1.3.6 DTS

Data Transmission Service (DTS) is a sub-module of DMP, providing stable and reliable database services of data migration and data upload to the managed cloud with automated and standardized processes to efficiently meet users' needs.

DMP supports performing full and schema transmissions on Oracle databases from outside to within a VPC network and a classic network and uploading their data to the managed cloud service. The data of Oracle databases can also be transmitted across platforms or architectures and between databases across minor versions. MySQL databases support full and full plus incremental transmissions, and their data can also be transmitted between databases across minor versions (5.6 to 5.7) or platforms.

DTS consists of features such as database connectivity tests, transmission configuration item settings, pre-transmission checks, data transmission, and data validation after transmission, enabling data migration with simple operations. Before transmission, DTS will perform pre-checks on relevant databases, systems, objects, and data, and create tablespaces, temporary tablespaces, and DB links to ensure the data is successfully transmitted to the destination database.

8.1.3.7 Account Management

DMP provides the **Accounts** module for MySQL and Oracle database account management, which includes the following query and modification features of accounts and account permissions for database clusters:

Add, edit, and delete accounts, modify account system permissions, and add remarks to accounts.

Add and delete account object (table, view, and stored procedure) permissions.

Synchronize account information.

With the convenient account management feature, users can query, add, and delete accounts and account permissions without logging in to the VM on which the database is running. By clicking the **Sync Account** button, changes made to database accounts or permissions can be synchronized to the **Accounts** page, preventing inconsistencies between the two sides from causing operation failures. They can also quickly search for accounts and permissions by keywords such as username, IP address, account type, and permission name.

8.1.3.8 DMC

There were previously inconvenient for users to log in to a database instance through the console on HCI or SCP. To improve the user experience, DMP provides the Database Management Consoles (DMCs) of the database instances, allowing the login in to a console by simply entering the username and password and specifying a database instance.

Currently, only MySQL and Oracle databases support DMC, and SQL Server databases only support the remote desktop feature.

This module comes with the following benefits:

1. Easy to Use: Visualizes the interface to match the operating habits of database administrators.
2. Highly Scalable: Compatible with various types of databases and underlying database drivers with a unified interface for external calls.
3. Easy to Manage: Automatically and uniformly manages session connections, including session overflow, timeout, etc.
4. High Security: It can only be logged in from DMP and controlled by DMP.

8.1.3.9 Operations Management

DMP is derived from SCP with native multi-tenant operations management capabilities, which can be directly applied to multi-subsidiary and multi-departmental scenarios. Instead of recreating tenants, DMP automatically inherits from SCP the information of tenants that meet the following conditions:

1. With RDS quota for DMP.
2. With the DMP service started.

8.1.3.9.1 Multi-Tenant Isolation

On SCP, resources are allocated to tenants for hierarchical management. For example, a company with multiple departments (tenants) naturally requires

that each department's resources be isolated, which is an essential capability of SCP. DMP also supports multi-tenant isolation, where different tenants' resources and virtual private networks (VPCs) are completely independent.




Figure 78: Multi-Tenant Isolation

A tenant can use an SCP account to directly access DMP and use the database resources within its quota without problems such as resource preemption or network IP address conflicts with other tenants.




Figure 79: DMP Home Page

8.1.3.9.2 Quota Management

While enterprise administrators pay management costs for multi-tenant management, the quota management feature empowers tenants to implement self-management and self-service within the quota to reduce the

administrators' burden.

An administrator can pre-allocate resource quotas for resource pools of tenants on SCP, and tenants can use resources as needed within those quotas.

Using a database requires the following resource quotas:

1. Compute
 - VM: Cannot be less than MySQL RDSs.
 - CPU: Based on the number and specifications of the corresponding VMs.
 - Memory: Based on the number and specifications of the corresponding VMs.
2. Storage (Select one of the following three based on the needed storage capacity)
 - High performance storage.
 - Performance sensitive storage.
 - Capacity sensitive storage.
3. Database Management Platform
 - Cloud Database-Oracle: Allocated in different specifications based on the number of tenant databases, where a RAC database requires two quotas.
 - Cloud Database-MySQL: Allocated in different specifications based on the number of tenant databases, where a primary-replica database requires two quotas.
 - Load Balancer: Required for active-standby databases.

8.1.3.9.3 Work Orders

The quota management feature eases the administrative burden and enhances the self-service capability of tenants, while the work order feature further widens the self-service scenarios. Tenants can submit work orders to apply for resource quotas on demand to the corresponding administrators, who only need to follow the process for approval, providing a highly flexible

self-service delivery and use mode to support the agile transformation of business services.

Tenants need to specify the applied resource entries and fill in the reason before submitting a work order. After approval by administrators, the corresponding resource quota will be allocated to tenants.




Figure 80: Work Order Template

Administrators can view, approve, and reject pending work orders, as well as modify resource quotas applied by tenants on the unified work order management page, significantly improving operation and management efficiency through the management process.




Figure 81: Work Order Processing by Administrator

8.2 Middleware

8.2.1 Introduction

SCP provides the middleware service capability through the **Application**

Center at the PaaS layer, supporting users to customize the business hosting middleware, such as message queues Kafka and RabbitMQ, Redis and Memcache caches, and MySQL, PostgreSQL, and MongoDB databases.

8.2.2 Application Scenarios

As an important concept in computer systems, middleware refers to the software layer of a system between the application programs and the operating system. It provides services and support for interaction and collaboration between applications and is widely used in various fields.

A cloud environment is a virtualized environment in which applications and services can be deployed. Middleware is used for communication and management between applications to ensure the reliability and security of the cloud environment. For example, a load balancer can distribute requests in a cloud environment to different VMs or instances for system scalability and availability.

8.2.3 Supported Features

SCP provides all middleware services that can be added based on business needs through the **Application Center**. The following only introduces common message queues RabbitMQ and Kafka.

1. RabbitMQ

SCP provides an easy-to-deploy, stable, and reliable message queue RabbitMQ of 3.7.28 standard version in standalone mode and cluster mode. Its typical application scenarios are as follows:

- **Business Decoupling:** RabbitMQ can reduce the business coupling of each process of applications, solve the problem of multi-business correlation in complex systems, effectively improve processing efficiency, and provide a transactional guarantee.
- **Traffic Peak Cutting:** RabbitMQ can pile up the messages of business services that require traffic control to effectively cut

down traffic peaks and protect the backend business services from being impacted.

2. Kafka

SCP provides highly concurrent and available message middleware Kafka 2.00 in standalone mode and cluster mode, which applies to traffic data processing, business decoupling, traffic peak cutting, and other scenarios. Its typical application scenarios are as follows:

- **Log Analysis:** Kafka can be integrated with a big data development kit to build a complete log analysis system.
- **Business Decoupling:** Kafka can reduce the business coupling of each process of applications, improve processing efficiency, and lower the risk of change.

9 O&M Management

9.1 Monitoring Center

9.1.1 Monitoring Dashboard

SCP can monitor VMs, physical hosts, elastic IP addresses, and shared bandwidth, and display the monitoring information in panels. The monitoring metrics are listed below:

1. VM: CPU Usage, CPU Used, CPU Used (Hz), CPU Wait, CPU Limit Exceeded, Memory Usage, IOPS, IO Speed, IO Latency, Throughput, Throughput (pps), and Disk Usage.
2. Physical Host: CPU Usage (Optimized), CPU Usage (Hz), Memory Usage, IOPS, IO Speed, and Throughput.
3. EIP and Shared Bandwidth: **WAN Bandwidth**.

The **Monitoring Dashboard** feature supports customizing panels, setting monitoring time ranges, and exporting results.

- Metric Collection Method

1. CPU Used

Collect the CPU used by each vCPU per VM through the corresponding thread on the host.

Read the user CPU time and system CPU time of each thread from the `/proc/<PID>/task/<TID>/stat` file (<PID>: process ID, <TID>: thread ID) and sum the two values (corresponding to columns 14 and 15 in the file). By collecting the sum value twice and calculating the difference, the CPU time during the interval is obtained. Dividing this value by the sampling interval yields the CPU used.

2. CPU Used (Hz)

Multiplying the CPU Used metric by the CPU clock speed can convert the value into CPU used in Hz.

3. CPU Wait

Collect the CPU wait time of each vCPU per VM through the corresponding thread on the host. Obtain the guest time (column 43) of each thread from the `/proc/<PID>/task/<TID>/stat` file. By collecting the value twice and calculating the difference, the CPU time in the guest during the interval is obtained.

Obtain the CPU ready time from the `/proc/<PID>/task/<TID>/schedstat` file (indicated by the increment of the value in column 2).

Calculate the CPU wait time using the following formula:

CPU wait time = Total sampling interval - CPU ready time - Guest CPU time Differences

Method on HCI: Sample and calculate based on native Linux kernel interfaces. Guest time indicates the time that a task is being executed on the VM, and CPU ready time indicates the time that a task is ready to be executed on the VM but is blocked due to host scheduling or other causes. Therefore, by subtracting the two components above from the total time, the remaining time indicates pure idle time, during which the core is neither executing tasks nor waiting to be scheduled.

Method on VMware vCenter: Sample and calculate through modifying the kernel. The results contain idle time and time such as swap wait and CPU IO wait, providing a broader statistical dimension.

4. CPU Limit Exceeded

Read the `cpu.stat` file in the cgroup where the VM resides. By collecting specific values twice and calculating the difference, CPU limit exceeded time during the interval can be obtained.

5. IO Latency

To obtain the read and write latency, use the QMP interface to run the human-monitor-command and "command-line":"info iostat" parameter commands (full QMP input: `{"execute": "qmp_capabilities"}\n{"execute": "human-monitor-command", "arguments": {"command-line": "info iostat"}}`) and check the `r_await` and `w_await` fields.

6. Disk Partition Usage

To view the usage of mounted partitions, run the corresponding commands in the guest OS. Run the `df` command on Linux, and the `wmic logicaldisk get deviceid, size, freespace` command on Windows.

7. Node CPU Usage

With turbo boost disabled, the actual CPU usage is consistent with CPU usage displayed. With turbo boost enabled, a CPU core with two threads is considered busy if either thread is busy. In this case, the actual CPU usage for the core is defined based on this definition of busy. During the interval between two samples, the ratio of time the CPU core is busy multiplied by the frequency ratio (the ratio of the current core frequency to the base frequency) is the actual CPU usage.

- Metric Function

CPU Used: Monitor the utilization of each CPU core of VMs. Previously, CPU usage for VMs was provided as an overall “CPU usage” metric. This metric provides a more granular, per-core view of utilization. For example, a single-threaded app only consumes one vCPU of the VM. Therefore, the following scenario may occur: A specific vCPU used for an app is used up, but the overall CPU usage remains low.

CPU Used (Hz): Used to quantify the actual CPU computing frequency consumed by each CPU core of a VM (Unit: Hz).

CPU Wait: Used to measure the idle time of each CPU core of a VM.

CPU Limit Exceeded: Used to monitor the cumulative time during which the used CPU of a VM exceeds the configured maximum frequency.

IO Latency: Used to check whether a VM has IO performance bottlenecks.

Disk Partition Usage: Used to check whether the partition space usage on the VM is excessively high, preventing business issues caused by insufficient disk space.

Node CPU Usage: Used to reflect the actual computing capacity used. The metric is related to %UTIL, hyper threading, and frequency scaling. Two hyper-

threads on a physical core share the cache and execution units. Therefore, the performance when both hyper-threads are running at the same time is not twice (about 1.3 times) the performance of a single hyper-thread running alone.

- Metric Reference Value

Metric Name	Unit	Range	Attention Level
CPU Used	%	[0, 100]	Normal
CPU Wait	ms	[0, 3500)	High latency
		[3500, 20000]	Normal
IO Latency	ms	[0, 55)	Normal
		[55, +∞)	High latency
Disk Partition Usage	%	[0, 90]	Normal
		(90, 100)	Disk partition space is nearly full, which may impact business.
Node CPU Usage	%	[0, 120)	Normal
		[120, +∞)	High latency

9.1.2 Reports

SCP supports statistics on the real-time status and monitoring charts of platform resources, summarizing resource configuration information, and previewing and exporting monitoring reports. The report content to be displayed can be selected as needed, including **Resource overview**, **Resource pools**, **Clusters**, and **VMs**. Users can configure report templates for subsequent export.

9.1.3 Alerts

SCP supports detecting and alerting on VMs, physical hardware, networks, and license key information, as well as setting alert trigger conditions and emails to receive alerts. All alerts are logged and can be viewed by administrators on demand.

1. Supports configuring alert policies for VMs, physical resources, networks, platform, and SCP cluster, which include alert thresholds and durations.

2. Supports viewing the alert logs of the above resources by filtering by multiple conditions such as duration and severity.
3. Supports configuring notification policies, which include the email for receiving alerts and the unified notification interval.

9.2 SkyOPS

9.2.1 Introduction

Sangfor SkyOPS is an intelligent O&M management component for SCP and is quickly deployed on HCI from an image template, so that failure alerts and potential risks in business environments can be encrypted and transmitted to it. Based on the Sangfor Tiangong AIOps Engine, SkyOPS is designed to help users quickly fix failures, predict risks, and provide recommendations. It offers remote O&M management of cloud platforms deployed in user data centers.




Figure 82: SkyOPS Business Process

9.2.2 Supported Features

SkyOPS is equipped with Sangfor Tiangong AIOps Engine, which is the first intelligent O&M analysis engine in O2O scenarios in China and protects cloud computing business services from failures in the entire lifecycle. It is applicable to more than 150 failure scenarios and can predict failures in 7 to 30 days, identify them in 1 minute, and locate them in 3 minutes. SkyOPS has the following features:

1. Detailed Monitoring and Alerts: Provides unified full-stack monitoring and alerts from the underlying hardware to the business applications of cloud platform infrastructures.
 - Hardware layer monitoring covers 81 failure scenarios with nearly 200 alert metrics, effectively classifying the severity of failures and improving the efficiency of failure fixing.
 - Business layer monitoring covers VM health status, VM backup and disaster recovery policies, database performance, application health status, and aSecurity.
2. Leading Intelligent Prediction: Provides pre-event prediction, in-event detection, and post-event fixing and recommendations based on AI-first thinking in O&M scenarios.
 - Disk failures are predicted 7 to 30 days in advance with an accuracy rate of no less than 90%.
 - Memory failures are predicted with a recall rate of 40% and more than 90% of them are prevented in advance based on memory CE isolation.
 - Idle VMs and VMs with low resource usage are identified and the corresponding operation recommendations are provided. Resource usage trends are predicted based on their historical changes and recommendations for planning available resources are given to users.




Figure 83: Tiangong AIOps Engine

9.2.2.1 Disk Failure Prediction

SMART attributes that provide disk status information and have a strong correlation with disk failures are obtained through feature engineering. Then the failed disks are sampled based on these key attributes. The analysis in the actual environment reveals that the failed disks that are close to each other have a strong correlation with their failures due to similar business loads and environmental factors such as server voltage, air temperature and humidity. From there, the Attention mechanism of deep learning is introduced for feature extraction from disk locations and integrated with SMART attributes to build a failure prediction model.

The overall model design employs the convolutional neural network (CNN) and transformer. Limited by the continuous change of disk status and lifetime, the traditional machine learning lacks the dynamic feature digging of the time window, and has poor model robustness and generalization. Therefore, temporal and causal relations are extracted based on a causal CNN. Then the transformer is used to dig the neighborhood spatio-temporal features of disks. Disks of different vendors and models have different collection rates, while business requirements impose disk data drift and distribution differences, therefore the transfer learning algorithm is introduced. By adding weights to each training set sample and using the weights to mitigate the test set data with different distributions, the prediction accuracy rate of the model for

disks of different vendors and models is further improved to above 90%.




Figure 84: Disk Failure Prediction

9.2.2.2 Memory Disk Prediction

Memory failures are common in cloud computing infrastructures. As a result, the stability and reliability of servers and even the entire IT infrastructure are compromised, causing an adverse impact on the SLA.

Memory failures are predicted using error characterization:

1. The following figure shows the memory microstructure. All error data on memory is recorded in logs, which are then parsed for error creation time, location, and type. Traditional machine learning struggles to process temporal information effectively and has relatively modest performance on real datasets.
2. The deep learning algorithm is applied to make full use of the temporal information to help the model learn the pattern of memory error occurrences.
3. Advanced semantic signatures of memory errors in temporal context are integrated through deep learning using positional attention to improve failure prediction.




Figure 85: Memory Failure Prediction

9.2.2.3 Resource Prediction

The resource prediction capability based on business characteristics can predict host and VM resource demand trends with an accuracy rate of 96.8%, prevent performance bottlenecks and reliability problems consequently caused due to insufficient resources, and provide accurate resource allocation and scheduling recommendations for business services with consideration of resource profiling to ensure excellent user experience.

There is a strong association between business resources and performance metrics. Therefore, multi-dimensional performance metrics are used to predict resource requirements, while the associated metrics are used for modeling to improve prediction accuracy.

Sangfor's AI model is simple and efficient, utilizing time series decomposition and deep linear fitting to overcome the key pain points of high resource consumption and low prediction efficiency.




Figure 86: Resource Prediction

9.2.2.4 Alert Convergence

SkyOPS provides the alert convergence features in response to the common alert storms in O&M:

1. Deduplicate the alerts and templatize them by NLP algorithm.
2. Use segmentation methods and entropy calculations based on deep learning to sort the alerts according to their entropy values.
3. Use a clustering algorithm to cluster alerts with similar patterns and characteristics into a type of critical event for alert convergence, thus resolving the alert storms during system failures and enabling users to focus on business services.
4. Use a periodic convergence algorithm to converge the alerts (such as every 15 minutes). If the period is too long, the timeliness of the alerts will be affected; if the period is too short, the alerts with a long time cannot be converged.
5. Use a dynamic time window event mechanism with multiple algorithms such as information entropy, NLP, and semantic similarity matching to form the first alert into an event and inform the public about it immediately, and then add new alerts to the event continuously, enabling the convergence of repeated alerts and the analysis of related alerts and event integration.

9.2.2.5 Root Cause Analysis

To perfect the root cause analysis and visualize the anomaly propagation

across metrics, the following essential works were completed:

1. Causal Graph Creation: Business software stacks vary greatly and change frequently in different scenarios. The causal graph of metrics is constructed using the causal inference model together with the metrics time-series data distribution to simulate the propagation path and intervention effect of critical reliability issues.
2. Abnormal Event Digging: Business anomalies may be caused by multiple factors. Through the anomaly detection algorithm with multiple metrics, abnormal behaviors on system performance and reliability can be dug out while maintaining the stability of the model with limited data.
3. Root Cause Analysis: An anomaly occurring on the root cause node can lead to fluctuations in the service metrics of related nodes, spreading anomalies to multiple nodes or multiple metrics of multiple nodes. By analyzing the anomalies between cause and effect, simulating their propagation paths, and integrating root cause location, an explainable anomaly event tracing path can be generated to help O&M personnel quickly locate failures.




Figure 87: Root Cause Analysis

9.2.2.6 Full Link Failure Detection and Location

Network issue troubleshooting is generally very challenging due to the complexity of cloud networks. By thoroughly analyzing the troubleshooting process and major pain points of O&M personnel and using knowledge graphs

to convey the structure and knowledge of complex heterogeneous cloud networks, the full link diagnostic capability of Sangfor Cloud in the vertical fields is greatly empowered.

1. Key network metrics, basic configuration, alerts, logs, and other data are combined with Sangfor's existing troubleshooting cases to build a full link knowledge graph.
2. Based on network architecture and the specialized experience of Sangfor Cloud, the graph is divided into three parts: static assets, dynamic events, and knowledge base.
3. The graph is further divided according to the network layer (IDC network, physical network, host network, and virtual network) and the system layer (application, platform, system, and hardware). Eventually, it covers not only various forwarding network components but also HCI software systems and hardware assets.
4. The dynamic event graph focuses on possible events (including but not limited to anomaly alerts and change events) to associate static assets.
5. The deep learning model extracts asset objects contained in unstructured event data, automates the association between dynamic events and static assets, and then based on their relationship graphs, quickly locates network outages and analyzes them for the key causes of failures.




Figure 88: Full Link Failure Detection

9.2.3 Benefits

1. Proactive O&M

SkyOPS turns traditional passive O&M into proactive O&M, which effectively reduces system failures, and enhances its availability and performance, while mitigating the workload of O&M personnel and improving O&M efficiency and quality.



Figure 89: O&M Mode Comparison

2. Online Guard to Save Labor

Traditional offline O&M requires on-duty personnel to monitor the running status of a data center and identify and address possible failures and problems as soon as possible, often with a 24-hour uninterrupted guard to ensure that the data center is running stably.

SkyOPS implements the online guard, which can quickly find and solve problems, prevent potential failures, and respond to customer needs through remote real-time monitoring and O&M management system. In addition, periodic online expert inspections and optimizations are conducted on it to keep the system healthy.

3. Intelligent Recommendations

System failures can be caused by a variety of reasons such as hardware failures, software errors, and network issues, which are often addressed based on experience. SkyOPS analyzes data such as failure features and logs through

AI, and generates fixing recommendations to help users quickly locate failures and recover their business services, thus improving the efficiency of failure fixing.

4. Real-Time IT System Status Control

SkyOPS provides real-time IT system status through the website and WeChat mini program and monitors full-stack system alert information, including underlying servers, cloud platforms, VMs, databases, and applications.




Figure 90: Mobile Client Display

9.3 Reliability

9.3.1 Holistic View

SCP provides a holistic view that displays reliability services at the business, platform, and hardware layers. In this holistic view, administrators can view the configuration, usage, and status of each reliability service in the entire cloud data center, fully adjust configurations, and quickly locate specific problems.




Figure 91: Holistic View of Reliability Services

9.3.2 HA

9.3.2.1 Introduction

SCP is integrated with the underlying HCI to provide the VM HA feature. When any of the hardware, storage, or network of a host fails, the HA mechanism will automatically recover the VMs on it (including virtual network devices) to run on other hosts to ensure business continuity.

9.3.2.2 Application Scenarios

For business service interruptions caused by external environment failures (such as host network cable cuts or storage inaccessibility) and VM guest system failures, both SCP and HCI provide a mature and reliable HA mechanism to ensure that business services are not interrupted or temporarily interrupted.

9.3.2.3 Supported Features

VM HA requires that the cluster where the VM is running consists of two or more host nodes. When the node on which a VM with HA enabled fails (due to power supply cut, network disconnection, etc.) and is detected by the cluster

heartbeat mechanism, the VM will be automatically restarted on another node with sufficient resources.

The cluster heartbeat mechanism uses round robin, which detects the VM status every five seconds. When a failure is detected and lasts as long as the monitoring sensitivity set by the user (e.g., a maximum sensitivity of 10 seconds in 5 minutes), the HA-enabled VMs will be switched to run on other hosts.

As shown in the figure below, users can configure the HA mechanism and select the **Monitoring Sensitivity** for various failure scenarios on the **High Availability (HA)** page according to business needs.

HA

Resource Pool: Malaysia-EC...

Hardware Platform Application Advanced Settings HA History

You can specify response methods for scenarios with hardware failures or hardware health issues.

Hardware Failure

Check Entity	Impact	Response Method
CPU Failure	Nodes will be restarted or go offline, and the corresponding VMs will be shut down.	
Memory Failure	Nodes will be restarted or go offline, and the corresponding VMs will be shut down.	
System Disk Failure	Nodes will go offline, and the corresponding VMs will be shut down.	
Motherboard Failure	Nodes will go offline, and the corresponding VMs will be shut down.	
Encryption Card Failure	The corresponding VMs cannot run properly, causing business service interruption.	
GPU Failure	The corresponding VMs cannot run properly, causing business service interruption.	Ignore

Hardware Health Issues Host Health Monitoring

Check Entity	Trigger	Impact	Response Method
Low CPU	Enablert	Nodes will be exposed to security vulnerabilities and performance	Ignore

Save **Restore Defaults**

Figure 92: HA Configuration

9.3.3 Auto Scaling

9.3.3.1 Introduction

SCP supports auto scaling (AS) for business services, which enables the VM

creation and deletion based on load monitoring. When the VM load varies greatly and exceeds the specified threshold, new business host nodes will be automatically deployed to meet the resource requirements of applications and to reduce manual intervention.

9.3.3.2 Application Scenarios

To ensure the stable performance of business systems (e.g., e-commerce platforms, website applications) with distinctive peaks and valleys, it is necessary to increase the number of hosting VMs according to the changes in real-time traffic to prevent the systems from being inaccessible due to surges in traffic, as well as to reduce the number of VMs as soon as possible after the peaks in order to increase the resource usage. On SCP, the size of VM groups is adjusted by auto scaling.

9.3.3.3 Supported Features

SCP supports creating AS groups and AS configurations and configuring flexible AS policies to respond to business needs during peaks and valleys.

1. Create AS groups to deploy VMs in them based on business requirements.
2. Create AS configurations to set the specifications of VMs in AS policies.
3. Create AS policies to set the thresholds and actions for scaling.

9.3.4 DRS

9.3.4.1 Introduction

SCP provides the dynamic resource scheduling (DRS) feature, which is implemented through the integration with the underlying HCI. DRS dynamically assesses the resource situation of the entire cluster by monitoring its resource pool usage, and live migrates VMs from overloaded hosts to other hosts with adequate resources to ensure business health and balanced loads in the cluster.

9.3.4.2 Application Scenarios

In scenarios where a physical host is overloaded due to surges in access traffic or processing data from business systems deployed on its VMs, the VMs will be intelligently scheduled to other hosts based on the resource load of each host to ensure business continuity and stability.

9.3.4.3 Supported Features

DRS works as follows: A cross-cluster heartbeat mechanism is implemented to monitor the utilization of the computing resources (including CPU and memory) of hosts in the cluster regularly. When the resource usage of a host reaches a custom threshold and there are hosts with adequate resources, DRS will migrate VMs from the original host to others. Users can customize overload thresholds and duration for host resources to avoid frequent business switches caused by DRS. Additionally, they can select manual or automatic resource scheduling as needed.

Users can customize a dynamic resource scheduling policy by specifying the scheduling method, resource type, and sensitivity.

- **Automated:** The system performs automatic scheduling based on resource load rules to reduce the workload of the administrators.
- **Manual:** The system provides scheduling recommendations based on resource load rules for the administrators to perform manual scheduling.
- **Resource:** The options **CPU** and **Memory** are available, for which a threshold needs to be set.
- **Sensitivity:** The options **Conservative** and **Aggressive** are available, which correspond to 10 minutes and 3 minutes respectively.

Once the usage of the selected resource type exceeds the threshold for the duration specified by the **Sensitivity**, an alert saying "**Status: Unbalanced**"

appears on the page and the system proceeds to perform automatic scheduling or provide scheduling recommendations based on the specified scheduling method.

9.3.5 Resource Reservation

9.3.5.1 Introduction

Resource reservation refers to reserving certain resource space on some hosts, and these resources will not be used by the system or other VMs. When a host fails or a disaster recovery switchover occurs, the reserved resources can be used to recover VMs immediately, reducing the risk of business recovery failure due to insufficient resources.

9.3.5.2 Application Scenarios

When a physical host fails, the VMs on it can be migrated to other host nodes through the HA mechanism and be quickly recovered using the reserved resources on new nodes to avoid business impact. Reserved resources can also be used to recover VMs from a disaster.

9.3.5.3 Supported Features

SCP provides the resource reservation mechanism to ensure sufficient resources for executing the HA feature in the event of host failures. A certain number of resources are reserved on physical hosts. They can neither be allocated nor used until certain host failure occurs and the HA mechanism takes effect. This mechanism prevents HA failure across the entire HCI due to excessive resource utilization.

9.3.6 VM Resource Limits

9.3.6.1 Introduction

SCP supports configuring the upper limits for CPU clock speed, disk I/O, and

NIC traffic of VMs through QoS settings. Administrators can set upper limits for the resources used by all VMs in the selected resource pools based on the criticality of business services.

9.3.6.2 Application Scenarios

Administrators can set upper limits for the resources used by some VMs based on the criticality of business services. This ensures that critical VMs have sufficient resources and abnormal VMs are restricted from consuming excessive resources, thereby maximizing resource availability for normal VMs.

9.3.6.3 Supported Features

Supports setting the upper limits for CPU clock speed, disk I/O, and NIC traffic of VMs through QoS settings within the following ranges:

1. CPU Clock Speed Limit: 100 MHz-1,000 GHz. Note that if the limit is too low, for example, 100 MHz, VM report data may not be updated in time and the VM may restart.
2. NIC Traffic Limit: 1000 Kbps-20000 Mbps.
3. Disk I/O Speed Limit: 128 KB/s-102,400 MB/s. Note that if the limit is too low, for example, 128 KB/s, the VM may restart because the too-low system disk I/O speed may not be enough to support VM running.
4. Disk IOPS Limit: 16-2,147,483,647.

9.3.7 VM Scheduling

9.3.7.1 Introduction

SCP supports setting the host on which a VM runs based on a VM scheduling policy. When the running location of a VM is set to be automatically selected, SCP assesses the remaining resources of the hosts and the resources required by the VM and selects an optimal host for running the VM.

9.3.7.2 Application Scenarios

To guarantee the high reliability of business systems or access between them with low latency, administrators can configure scheduling policies for business VMs to schedule them to different or the same physical hosts.

9.3.7.3 Supported Features

The following anti-affinity and affinity scheduling policies can be configured for VMs and clusters. These policies can be used to accurately control the running locations of the VMs to optimize business services running on VMs.

1. VM-VM Affinity: Schedules VMs to the same physical host. For example, VMs that communicate frequently with each other or in the same business group are scheduled to the same server to speed up access to business services.
2. VM-VM Anti-Affinity: Schedules VMs to different physical hosts. For example, Oracle RAC cluster systems are distributed on different servers to ensure business reliability.
3. VM Group-VM Group Anti-Affinity: Schedules VM groups to different physical hosts. For example, three VMs running business service A and those running business service B are mutually exclusive, preventing the failures of the physical host running business service A from affecting business service B and improving business availability.
4. VM-Host Affinity: Schedules VM groups to different physical host groups. For example, the active-active business service groups A and B are respectively run in the primary and secondary fault domains of a stretched cluster to ensure both active-active data and services.

9.3.8 SCP Status Check

9.3.8.1 Introduction

SCP supports viewing the platform status on the **Holistic View** page, including **Overload Check**, **Network Check**, **Datastore Status**, **Data Rebuilding**, and

SCP Failover, which helps administrators and O&M personnel to understand the running status and hardware status of clusters, identify problems in time, and quickly perform O&M management. There is also the **SCP Status Check** page for them to perform manual checks and operations.

9.3.8.2 Application Scenarios

Administrators and O&M personnel can go to the **SCP Status Check** page to manually check the SCP status and view alert details.

9.3.8.3 Supported Features

1. **Overload Check:** Checks whether CPU, memory, storage, and network throughput are overloaded.
2. **Network Check:** Checks redundancy and connectivity of management interfaces, overlay network interfaces, and storage network interfaces of hosts.
3. **Datastore Status:** Multiple replicas on different hosts ensure accessibility to data even when a host fails unexpectedly. It checks the state of those replicas to ensure data reliability.
4. **Data Rebuilding:** When a host or disk fails and times out, data replicas will be automatically rebuilt on another host or disk in the cluster. If the storage space is sufficient with no need for spare disks or manual replacement of the host or disk.
5. **SCP Failover:** Supports viewing the status of active and standby SCP nodes, which can be deployed to guarantee high availability. When the active node is unavailable, users can log in to the standby node and switch it to the active node.

9.3.9 Hardware Health Check

9.3.9.1 Introduction

SCP supports viewing the status of the underlying hardware resources of the platform on the **Holistic View** page, including **CPU, Memory, RAID Card, SSD, HDD, NIC**, and **External Storage**, which helps administrators and O&M personnel to identify potential availability issues in advance. There is also the **Hardware Health Check** page for them to perform manual checks.

9.3.9.2 Application Scenarios

Administrators and O&M personnel can go to the **Hardware Health Check** page to centrally view the status of the underlying hardware resources of the platform, and quickly perform O&M management to safeguard hardware health.

9.3.9.3 Supported Features

SCP supports the one-click health checks on the hardware resources of the underlying physical servers of the platform, involving CPU, memory, SSD, HDD, NIC, RAID card, and external storage. All the abnormal check items and those with alerts will be displayed on the page with corresponding suggestions.



It is recommended to perform one-click health checks in an SCP cluster during off-peak hours; otherwise, the performance of running business systems will be affected.

9.4 Security

9.4.1 Login Policies

SCP supports configuring login policies for accounts, including **Account and Password, Session Timeout Options**, and **Login Protection**, as well as two-

factor authentication login with a password plus SMS authentication or OTP authentication for enhanced account security.

9.4.2 Disk Encryption

SCP provides the disk encryption feature, which can encrypt storage data of VMs and network and security resources to guarantee data security and avoid information theft. This feature can only be enabled and disabled for shutdown VMs. Users can encrypt VMs in either of the following ways:

1. Use a built-in key

Instructions: Go to **Resources > Compute > Virtual Machines**, click **More** of the VM to be encrypted, click **Enable Disk Encryption**, and then select an encryption key.

2. Use a KMS key

Instructions: Go to **Resources > Management > Security > Key Management Service** to add a KMS server and encryption key, and then follow the instructions described above for the first method.

9.4.3 Data Erasure

When VMs and backups in the **Recycle Bin** are permanently deleted, their data will be erased completely by overwriting to guarantee data security.

9.4.4 Key Pairs

VM login using the account password has weak security and may involve the risk of brute-force attacks. SCP provides the VM key pair feature, a secure and convenient login authentication method. A key pair contains a public key and a private key. The public key is sent to the bound VM, to which the user can log in remotely through SSH with the private key.

9.4.5 Key Management Service

SCP provides a key management service (KMS) that supports adding external

KMS servers, and Sansec and JIT HSMs to the platform, and creating built-in and external KMS server keys.

10 Operations Management

10.1 Multi-Level Roles

There are three roles on SCP, namely platform administrators, tenants, and tenant users.

10.1.1 Platform Administrator

A platform administrator has all the operations and O&M management permissions. It is the role with the most permissions on SCP. The super administrator is the default platform administrator who manages the entire platform and customizes the roles and permissions of platform administrators.

Permission	Content
Resource Pool Management	Clusters
	Storage Tags
	Operation Status Monitoring
	Resource Pooling
VM Management	VM Lifecycle
	VM Groups
	VM Migration and Cloning
	Images
Network Management	VPC Gateway Planning
	Elastic IP Pools
	Network Topology
	Organization Subnets
	Distributed Firewall
Business Reliability Management	Local Backup Policies
	CDP
	Disaster Recovery Tasks
	HA

	DRS
	Resource Reservation
Platform Reliability Management	Overload Check
	Network Check
	Platform Administrators
	Data Replication
	Data Rebuilding
	SCP Failover
	Hardware Health Check
User Management	Organization
	Users
Quota Management	Basic Resource Quota Allocation
	Security Resource Quota Allocation
	Pending Work Order Approval
	Work Orders
	Approval History Viewing
	Workflow Templates
Billing Management	Bill Viewing
	Billing Rule Setting
Monitoring Management	Monitoring Overview (VMs, EIPs)
	aMC Monitoring
	Alert Log Viewing
	Alert Rule Viewing
System Management	Cluster Licensing
	System Configuration, System Maintenance, and Recycle Bin Operations
	COC Connection
Image Management	Public Images
	Private Images
	Network Device Images

	Conversion Between Public and Private Images
Heterogeneous Management	VMware Clusters
	Physical Machines
	ASW Cloud
	Cloud Environments
Application Management	Automated Application Deployment
	Application Template Utilization and Management
	Application Template Import

Table 3: Platform Administrator Permissions

10.1.2 Tenant

A tenant has management permissions for all managed resources and tenant users and is the second-level management role in the Sangfor multi-tenant model.

Permission	Content
VM Management	VM Lifecycle
	Private Images
	GPU-Enabled VMs
	VM Groups
	Auto Scaling
Network Management	VPC Networks
	Network Topology
	Elastic IP Addresses
	Shared Bandwidth
Security Services	Direct Connections
	vNGAF
	Load Balancers
	vSSL VPN
	MLPS Compliance

Business Reliability	Cloud Backups
	Local Backups
	Remote Disaster Recovery
Dedicated Cloud	Dedicated Resource Pool Overview
Backup	Local Backups
Monitoring Management	Monitoring Overview (VMs, EIPs)
	Alert Log Viewing
User Management	Users
Quota Management	Basic and Security Resource Quota Viewing
Work Order Management	Pending Work Order Approval
	Quota Application to Platform Administrator
	Work Order Approval History Viewing
Billing Management	Total Cost Viewing
System Management	Recycle Bin
	Tasks Viewing
Heterogeneous Management	VMware Clusters
	Physical Machines
	ASW Cloud
Application Store	Application Management
	Application Templates
	Application Development

Table 4: Tenant Permissions

10.1.3 Tenant User

A tenant user has management permissions for the associated VMs and is the third-level role in the Sangfor multi-tenant model.

Permission	Content
VM Management	VM Application
	VM Management (Start, shut down, suspend, etc.)

Work Order Management	Work Order Workflow and Result Viewing
Monitoring Management	Monitoring Log Viewing
System Management	Personal Settings
	Tasks Viewing

Table 5: Tenant User Permissions

10.2 Multi-Tenant Management

A tenant is a unit for resource allocation on SCP and is applicable to hierarchical operation scenarios. The Sangfor multi-tenant model has three levels: platform, tenants, and tenant users. For example, if a company (platform) has multiple departments (tenants) each with a number of employees (tenant users), then each department can be allocated with a tenant on SCP, and each tenant can be allocated with the corresponding number of tenant users. The company needs to designate a platform administrator.

If hierarchical management is not required, only a tenant including all members needs to be created on SCP. The multi-tenant model is shown below:




Figure 93: Multi-Tenant Model

A platform administrator can create multiple tenants, which can be a collection of a company department or a specific project. A tenant has a certain resource quota and can request more from the platform administrator. Tenants can be allocated tenant users and allocate resource quotas to them.

A tenant user is an end user in the multi-tenant system and can manage its resource quota and apply for more from the tenant to which it belongs.

The local login portal is different for different roles. Platform administrators can only access SCP in internal networks, while tenants and tenant users can access SCP in both internal networks and the Internet. When SCP is connected to Sangfor Cloud, tenants, and tenant users can log in with authentication through the Sangfor Cloud portal and be redirected to the corresponding SCP according to their user roles.

10.3 Quota Management

The quota management of SCP meets the self-service needs of tenants in service providers, enterprise operations, and education scenarios.

Administrators can allocate basic and security resource quotas based on tenants' needs to improve operational efficiency.

On SCP, quotas can be set for tenants, who can use the allocated quotas in different resource pools as needed.




Figure 94: Tenant Quotas

whether it can be created; if yes, the resource will be reserved. If the creation fails, a rollback will be performed; if successful, the reserved quota will be committed.




Figure 95: Tenant Quota Management

10.4 Work Orders

Self-service work orders provide users with a multi-tenant cloud service directory, a new mode of infrastructure delivery and usage. SCP's resource virtualization and pooling capabilities allow IT departments to abstract physical resources into elastic virtual resource pools that can be provisioned as needed. These pools include VMs, storage, network, and security resources, and provide external services as consumption units (tenants or virtual data centers). IT departments can offer these consumption units, as well as other infrastructure and application service templates including VMs and OS images through fully automated and self-service access. Such self-service truly enables business agility and efficiency and greatly accelerates business innovation.

Below is an illustration of a self-service work order with the process for tenant users, tenants, and platform administrators.

1. The tenant user applies for VM resources

In addition to managing its VMs, a tenant user can initiate a work order to

apply for more VMs or change their configuration. The work order needs to be filled with information such as the number of CPU cores, memory size, and storage capacity to be increased and will be sent to the tenant and the platform administrator. After the work order is approved, the system will automatically issue the corresponding resources to the tenant user.




Figure 96: Apply for VM Resources

2. The tenant applies for resource quota

A tenant can initiate a work order to apply for more resource quota. The work order needs to be filled with information such as the number of CPU cores, memory size, and storage capacity to be increased and will be sent to the platform administrator. After the work order is approved, the system will automatically issue the corresponding resources to the tenant.




Figure 97: Apply for Resource Quota

3. The platform administrator approves the work order

A platform administrator can approve or reject a work order, modify the applied quota, or directly allocate the corresponding resources to a tenant or tenant user without any work order.

10.5 Billing

Every reduction in service costs and improvement in IT resource allocation results in capital expenditures in general IT departments being turned into service and user operations expenditures. Business resource requests compose the cost structure of each department or business service, eventually generating a monthly per-user cost.

SCP provides perfect resource billing services, supporting statistics on VM, storage, and network resources used by tenants based on a time range and exporting resource usage and billing reports for costing and analyzing for each business unit. It contains the following features:

1. Unit Price Setting

Platform administrators can set unit prices for all CPU, memory, and storage resources. Among them, CPU and memory resources can be individually priced according to the on/off status, and storage resources can be priced for high performance storage, performance sensitive storage, and capacity sensitive storage.

2. Cost Overview

Platform administrators can view the bill of each tenant, including the total cost and the costs of CPU, memory, and storage, as well as the costs incurred by each VM. Tenants can view their bills and the resources consumed by each VM.

3. Statistics Based on Time Range

Platform administrators can view the resources and costs consumed by tenants for a specified date range.

4. Data Export as Excel

Platform administrators can export billing data as an Excel file and save it

locally.




Figure 98: Billing Module Principle

The technical principle of the billing module is shown in the figure above. The Ceilometer module keeps real-time statistics on VM usage and updates the data table. The billing module queries the Ceilometer module every 10 minutes for VM usage and calculates the cost based on usage and unit price. It then writes the cost information to a unified data table, which is updated hourly. The UI performs a query in the cost data table and then displays the cost information. The total cost of tenants is updated at each user query.

10.6 Tasks

SCP logs all tasks and synchronizes the results, which can be audited by audit administrators and platform administrators. For example, administrators can go to the **Tasks** page to flexibly filter by time range and view the historical tasks to locate failures.

10.7 Recycle Bin

SCP provides a recycle bin feature to secure platform data. VMs, network devices, NFV devices, and backup files deleted by administrators are removed to the **Recycle Bin** and retained for 30 days, during which they can be

recovered or manually deleted permanently. Newly added tenants and tenant users can quickly identify information of deleted devices.

10.8 Federal Identity Authentication

SCP supports single sign-on (SSO) based on SAML 2.0, the federal identity authentication. Administrators can manage enterprise identity providers, enabling enterprise users to single sign on to SCP. Administrators and tenants can configure third-party service providers to which the users and user systems of the platform can be extended to broaden the service directory.

10.9 Custom Attributes

In customer scenarios with many business VMs and departments, custom attributes can help users categorize and label their assets, as well as perform unified O&M management of assets by resources. SCP supports users to create custom attributes according to the resources provided by the platform, including VMs, physical hosts, storage, and network components, and associate the created tags to the resources.




Figure 99: Custom Tags for VMs

10.10 Portal Customization

SCP supports the customization of the login page, including the platform name and profile, logo, login URL, and other information, which meets the

customization needs of customers in different industries.

11 Access Layer

To secure the platform and provide a better cloud service experience for different user roles, such as tenants and administrators, SCP provides a series of access control policies at the access layer, including WAF gateway, traffic gateway, business gateway, and OpenAPI gateway.

- **WAF Gateway:** Used for HTTP traffic detection and filtering to protect SCP from various attacks, such as SQL injection, cross-site scripting attacks, and file inclusion attacks.
- **Traffic Gateway:** Used for traffic distribution and restriction, which balances user access traffic through flexible traffic control mechanisms and implements granular traffic distribution and management based on demand.
- **Business Gateway:** Used for business service routing and forwarding, which works with the traffic gateway for large-scale access to SCP and business operations.
- **OpenAPI Gateway:** Used for the management of external OpenAPI to unify the external access portal and balancing internal loads, which greatly simplifies the development and maintenance of SCP's APIs; used for simple and fast integration of SCP with external systems (such as Sangfor Cloud, customers' O&M management platforms).

12 SCP and Sangfor Cloud Portals

SCP provides a unified access portal for various user roles such as tenants and administrators, as well as account registration, login, and features related to resource utilization and O&M.

In general management scenarios, tenants and administrators log in to SCP through the SCP portal to use and manage platform resources.



Figure 100: SCP Portal

In hierarchical management scenarios, tenants and administrators log in to Sangfor Cloud through the Sangfor Cloud portal to use and manage platform resources.




Figure 101: Sangfor Cloud Portal

13 Open APIs

SCP provides rich representational state transfer (REST) APIs externally that can be called by HTTPS requests for standardized management and O&M by third-party users.

For API documentation, please contact a Sangfor technical support representative.

14 Glossary

Name	Term	Description
Continuous Data Protection	CDP	Continuous data protection (CDP) supports recording every disk I/O of VM business systems, recovering VMs to any point in time, and viewing and downloading files of any point with an RPO near to zero, which helps a lot in such cases as file deletion by mistake, virus, system crash, data damage, etc.
Super Administrator	N/A	The super administrator (also known as the admin user) has all the SCP permissions. It is the role with the most permissions in HCI resource pools.
Dynamic Resource Scheduling	DRS	Dynamic resource scheduling migrates VMs from hosts with high CPU or memory usage to those with low usage, reducing the usage of overloaded hosts to below thresholds.
Dynamic Resource eXtension	DRX	Dynamic resource eXtension automatically adds CPU and memory resources to VMs during business peaks to guarantee business system availability.
Sangfor Compute Virtualization	aSV	Sangfor aSV virtualizes CPUs, memory, and I/O devices through the Virtual Machine Monitor (VMM).
Sangfor Storage Virtualization	aSAN	aSAN is an exclusive distributed storage system that integrates distributed caching, SSD caching acceleration, global application delivery, and multiple-level data protection technologies. It ensures efficient and stable operation of critical businesses.
Sangfor Network Virtualization	aNET	aNET is an exclusive network virtualization layer that provides a visual representation (What-You-Draw-Is-What-You-Get) of business logic. It supports the deployment of virtual routers, switches, and network security devices.
Sangfor Security Virtualization	aSecurity	aSecurity includes virtual security components such as vNGAF, vIAG, vADC, and SSL VPN.
Sangfor Cloud Platform	SCP	SCP is capable of managing HCI, VMware, and other cloud resource pools.
Sangfor aDeploy	aDeploy	Sangfor aDeploy is an essential tool for the delivery of Sangfor Cloud projects with features such as HCI and SCP health check, interface disorder adjustment, license key collection, and compatibility list query.

Sangfor aServer	aServer	aServer is an exclusive all-in-one machine for HCI and can be quickly deployed with high compatibility.
High Availability	HA	VM HA requires that the cluster where the VM is running consists of two or more host nodes. When the node on which a VM with HA enabled fails (due to power supply cut, network disconnection, etc.) and is detected by the cluster heartbeat mechanism, the VM will be automatically restarted on another node with sufficient resources, ensuring that business services are not interrupted or temporarily interrupted. When a failure is detected by the HA service of SCP and lasts as long as the monitoring sensitivity set by the user (such as 5 minutes), the HA-enabled VMs will be switched to run on other hosts.
Recycle Bin	N/A	It temporarily stores the deleted VM, network device, and shared disk resources. In addition, the deleted resources can be restored.
Network Function Virtualization Device	NFV Device	It refers to virtual network devices such as vSSL, vIAG, vNGAF, vADC, etc.
Clone	N/A	A clone can generate a VM identical to the original one to improve the efficiency of business deployment.
Snapshot	N/A	A snapshot is a record of the state of the data storage at a certain moment. Sangfor HCI records the state (data location) of a QCOW2 file by taking a snapshot of it and recovers VM data from it.
Deployment from Template	N/A	A VM (shutdown and locked) is first converted to a template, from which VMs with the same status as the template are deployed using the template image file and their incremental image files. The deployed VMs can also be configured with data disks.
Heterogeneous Virtualization Management	aHM	HCI supports adding and managing VMware vCenter for centralized management of both platforms. VMs can be migrated between VMware vCenter and HCI.

15 Abbreviations

Abbreviations	Full Name
ACL	Access Control Lists
aNET	Sangfor Network Virtualization
aSAN	Sangfor Storage Virtualization
aSEC	Sangfor aSecurity
aServer	Sangfor aServer
aSV	Sangfor Compute Virtualization
ATS	Atomic Test and Set
BIOS	Basic Input/Output System
BMC	Baseboard Manager Controller
CDP	Continuous Data Protection
COW	Copy On Write
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name Server
DRS	Dynamic Resource Scheduling
DXR	Dynamic Resource eXtension
DWPD	Drive Writes per Day
FC	Fibre Channel
Google OTP	Google One-Time Password
GPT	GUID Partition Table
GUI	Graphical User Interface
HA	High Availability
HBA	Host Bus Adapter
HCI	Hyper-Converged Infrastructure
iDRAC	Integrated Dell Remote Access Controller
IOMMU	Input/Output Memory Management Unit
IOPS	Input/Output per Second
IPMI	Intelligent Platform Storage Area Network

IPSAN	Internet Protocol Storage Area Network
IQN	iSCSI Qualified Name
iSCSI	Internet Small Computer System Interface
JBOD	Just a Bunch of Disks
LVM	Logical Volume Manager
LUN	Logical Unit Number
MBR	Master Boot Record
M-LAG	Multichassis Link Aggregation Group
MTU	Maximum Transmission Unit
NAT	Network Address Translation
NFS	Network File System
NFV	Network Functions Virtualization
NM	N/A
NTFS	New Technology File System
NTP	Network Time Protocol
NUMA	Non-Uniform Memory Access
OTP	One-Time Password
PCIe	Peripheral Component Interconnect Express
PVID	Port VLAN ID
RAID	Redundant Arrays of Independent Disks
ROW	Redirect On Write
RPO	Recovery Point Objective
RTP	Recovery Time Objective
SCP	Sangfor Cloud Platform
SR-IOV	Single Root I/O Virtualization
STP	Spanning Tree Protocol
TPM	Trust Platform Module
UEFI	Unified Extensible Firmware Interface
UPS	Uninterruptible Power System

URL	Uniform Resource Locator
UUID	Universally Unique Identifier
VAAI	vStorage API for Array Integration
viAG	Virtual Internet Access Gateway
vADC	Virtual Application Delivery Controller
vCS	Virtual Compute Server
vDAS	Virtual Database Audit System
vDWS	Virtual Data Warehouse Service
VMM	Virtual Machine Monitor
VMX	Virtual Machine Extension
vSSL	Virtual Secure Sockets Layer
vSSL VPN	Virtual Secure Sockets Layer Virtual Private Network
VT-x	Virtualization Technology-x
vWOC	Virtual WAN Optimization Controller
VXLAN	Virtual eXtensible Local Area Network



SANGFOR

