

ROMA Connect

Service Overview

Issue 03
Date 2024-05-16



Copyright © Huawei Technologies Co., Ltd. 2024. All rights reserved.

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

Trademarks and Permissions



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

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

Notice

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

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

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:

<https://www.huawei.com/en/psirt/vul-response-process>

For vulnerability information, enterprise customers can visit the following web page:

<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

Contents

1 Infographics.....	1
2 What Is ROMA Connect?.....	3
3 Application Scenarios.....	9
3.1 Smart Campus Integration.....	9
3.2 Industrial Internet Integration.....	11
3.3 Application & Data Integration of Corporation Groups.....	12
4 Edition Differences.....	14
5 Edition Differences (New Version).....	20
6 Supported Data and Protocols.....	23
7 Quotas.....	27
8 Constraint.....	30
9 Billing.....	36
10 Billing (New Version).....	37
11 Permissions.....	39
12 Basic Concepts.....	45
13 Related Services.....	48

1 Infographics

Too many applications, not enough coordination?
ROMA Connect is here to help.

ROMA

In today's quest for efficiency, new applications are constantly introduced, while legacy ones are still in use.

Witness miscoordination:
Finance: "It's great to meet you!"
HR: "Excited to join the team!"
Travel: "Excited to join the team!"

But since everyone is built differently and has different databases, data can't flow freely between them.

Here's my expense form.
Travel: "Here's my expense form."
Finance: "Oh, not this form!"

Reconstructing legacy applications for every new application costs time, money, and patience.

Instead, ROMA Connect bridges the gap between your application systems.
"I'm here to save the world."

What is ROMA Connect?

This new-generation platform seamlessly connects data, APIs, messages, and devices. Cloud native, low-code/no-code helps you quickly bridge legacy systems and cloud-native applications in your digital journey.

ROMA Connect gives you 4 big boosts:

- 01 FDI**
Flexibly, quickly, and unerringly integrates and converges various types of data (text/message/API/database).
- 02 APIC**
Encapsulates existing backend services (APIs), data sources, and custom functions into RESTful APIs.
- 03 MQS**
Provides secure and standard message channels for unified cross-network access.
- 04 LINK**
Provides device access using MQTT, Modbus, and OPC UA protocols for quick device management on the cloud.

Getting Started

With ROMA Connect, everything syncs up effortlessly. Firstly, it translates for different application systems.

Here, my expense form.
Travel: "Here, my expense form."
ROMA Connect: "Forwarding it now."
Finance: "Here's your expense form!"

Here, the expense form.
ROMA Connect: "Here's your expense form!"
Finance: "Here's your expense form!"

ROMA Connect also functions as a message broadcast station to sync information between different systems.

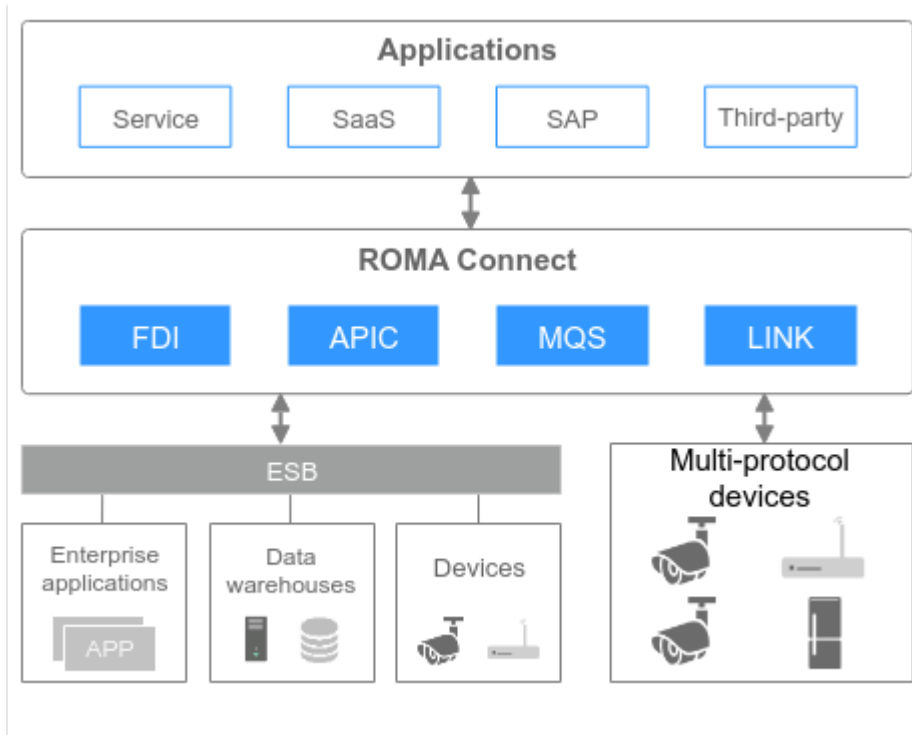
Oh yeah, News on board!
HR: "Oh yeah, News on board!"
ROMA Connect: "Copy that."
Finance: "Salary info received."
Admin: "Salary permissions disabled."

Huawei Cloud | Everything as a Service
For more information, please visit: <https://cloud.huawei.com/en/roma>
Huawei Cloud may collect and use your personal information for the purposes of providing and improving our services. For more information, please visit: <https://cloud.huawei.com/en/roma>

2 What Is ROMA Connect?

Enterprises are posed many challenges in their way to digital transformation. For example, device data is difficult to integrate, data in different formats cannot be transmitted or integrated, data and backend services cannot be shared with partners with ease, and there is no secure information channel for cloud and on-premises applications across difference networks. ROMA Connect is a full-stack application and data integration platform, the culmination of Huawei's digital transformation practices. It focuses on application and data connections and applies to various common use cases of enterprises. ROMA Connect provides lightweight message, data, API, model, and device integration to simplify the enterprise cloudification flow and support cross-regional integration for cloud and on-premises applications, helping enterprises achieve digital transformation.

Figure 2-1 ROMA Connect overall architecture



ROMA Connect consists of four components: data integration (FDI, short for Fast Data Integration), service integration (APIC, short for API Connect), message integration (MQS, short for Message Queue Service), and device integration (LINK).

FDI

FDI is a data integration component of ROMA Connect. FDI supports flexible, fast, and non-intrusive data integration between multiple data sources, such as text, messages, APIs, and relational and non-relational data. It implements data integration across equipment rooms, data centers, and clouds, and supports automatic deployment, O&M, and monitoring of integrated data.

For example, if an enterprise and its partners use different data sources, it is difficult to achieve effective information transmission. FDI provides multiple methods to convert mainstream data source formats such as MySQL, Kafka, and API.

Table 2-1 FDI functions

Function	Description
Lifecycle management of data integration tasks	FDI allows you to modify data integration task information and view running reports, run logs, and status of data integration tasks.
Flexible data reading and writing	<ul style="list-style-type: none"> • Reads and writes various types of data by fragment, such as MySQL data, text files, messages, and APIs. • Supports automatic recovery of tasks when the service is restored after an unexpected interruption occurs. • Supports task scheduling, monitoring, and resumable reading.
Reliable data transmission channel	FDI can continuously monitor data in data channels and supports concurrent execution of more than 100 threads. It monitors the message queue in real time and writes data to the target queue in real time.

Function	Description
Task scheduling	<p>FDI provides comprehensive, flexible, and highly available task scheduling services and supports data integration through APIs or messages. It schedules tasks based on time and data volume rules. FDI assigns tasks to the plug-ins based on the task configuration, and monitors and records the task execution status.</p> <p>Enterprises can select different data integration modes to suit their service requirements.</p> <ul style="list-style-type: none"> • Incremental real-time integration is applicable to scenarios in which data changes need to be monitored in real time, for example, collecting real-time parameters of devices on the production line. • Full real-time integration is ideal for scenarios in which all historical data needs to be monitored in real time, for example, collecting statistics on the supplier shipments. • Incremental scheduled integration is ideal for scenarios in which data changes need to be monitored for a period of time. For example, enterprises use new production policies to verify whether production efficiency meets expectations. • Full scheduled integration is ideal for scenarios in which all historical data needs to be monitored for a period of time, for example, collecting statistics on the number of vehicles entering or leaving a campus during peak and off-peak hours.
Alarms and monitoring	<p>FDI monitors the running status of data integration tasks and processes abnormal tasks to ensure service running.</p>

APIC

APIC is an API integration component of ROMA Connect. It opens data and backend services as APIs to simplify data sharing and service provisioning and reduce the cost on interconnection between enterprises. APIC provides SDKs and sample code in different programming languages to simplify the process of opening up backend services as APIs.

For example, if a company headquarters integrates its IT system with those of its branches in different regions, it is too complex to directly access each other's database and information disclosure may occur. If APIs are used to access databases and security is enhanced for API call, cross-network and cross-regional collaboration can be achieved.

Table 2-2 APIC functions

Function	Description
API lifecycle management	The lifecycle of an API involves creating, publishing, removing, and deleting the API.
Simple debugging tool	APIC provides an inline debugging tool to simplify API development and reduce maintenance costs.
Version management	An API can be published in different environments to meet version upgrade requirements.
Request throttling	Request throttling controls the maximum number of times an API can be called by a user or an app within a time period. The throttling can be accurate to the second, minute, hour, or day. Special applications can be configured so that they are not controlled by request throttling policies.
Monitoring statistics	APIC provides real-time, visualized API monitoring in terms of requests and errors.
Environment variables	When an API is published to different environments, the specified header parameters and special values are added to the API call request header to distinguish different environments. During publication, the variable is replaced with the environment variable value to ensure that the definition of the API does not change.
Custom backend	The custom backend supports data APIs and function APIs. <ul style="list-style-type: none"> • A custom data API allows enterprises to connect a database to APIC as a backend service and convert data service capabilities into REST APIs. • A custom function API is similar to a simplified function service. You can compile custom scripts or functions on the APIC backend as a backend service for the frontend to invoke.

MQS

MQS is a message integration component of ROMA Connect. MQS based on Kafka uses a unified message access mechanism to provide enterprises with secure and standard message channels for cross-network access.

For example, if an enterprise and its partners use different message systems, interconnection between the message systems is costly, and message transmission after the interconnection may not be reliable or secure. To address these issues, the Kafka protocol can be used for communication between the enterprise and its partners. In this way, MQS functions as a message transfer station to provide secure and reliable message transmission. Specifically, the enterprise can create

multiple topics, set the permission for each partner to subscribe to these topics, and publish messages to the topics. Then, partners can subscribe to the topics to obtain messages.

Table 2-3 MQS functions

Function	Description
Basic functions of Kafka	MQS supports topic management and message publishing and subscription after being connected to the client. It also supports visualized operations on the ROMA Connect console, including topic creation and management, user management, permission configuration, and message query.
Monitoring and alarming	MQS allows you to configure monitoring metrics from multiple dimensions, such as instances, nodes, topics, and consumer groups. After monitoring data is reported to Cloud Eye, enterprises can view the monitoring data on the Cloud Eye console. In addition, MQS allows you to configure alarm rules so that alarms can be generated if an exception occurs.
Message viewing	MQS provides a visualized message query function, which allows you to view the message data stored in topics on the console and view the message body more intuitively and conveniently.

LINK

LINK is a component of ROMA Connect for device integration. LINK uses the standard Message Queue Telemetry Transport (MQTT) protocol to connect devices, helping enterprises quickly and easily manage devices on the cloud.

In industrial scenarios, device information and parameters involved in the production process are scattered. If a fault occurs in a production line, it takes a long time to manually collect information and parameters for each device. LINK connects devices to IT systems or big data platforms, and uploads information such as device running status to these platforms so that enterprise customers can see information about all devices graphically and therefore quickly locate faults. In addition, enterprise customers can configure the upper thresholds for device parameters to rule engines of LINK. If real-time parameters of a device are close to the upper thresholds, an alarm notification is sent to users to remind them to stop the device and perform maintenance.

Table 2-4 LINK functions

Function	Description
Publishing and subscribing to messages	LINK supports the standard MQTT protocol. Enterprises can use open-source device SDKs based on this MQTT protocol to easily connect devices to the cloud for message publishing and subscription.
Message exchange between devices and backend applications	<p>You can configure a rule engine on the LINK console to enable a device to communicate with other devices, backends, and other cloud services.</p> <p>LINK supports rule engines to forward data to MQS. Third-party services obtain data through MQS to implement asynchronous message communication between devices and third-party services.</p>
Low-latency access for massive numbers of devices	LINK supports horizontal expansion of brokers and persistent connections of millions of devices.
Two-way synchronization between devices and applications	<p>LINK supports profile definition and binds the profile with a device shadow. This allows users to implement two-way synchronization of configuration data and status data between devices and applications.</p> <p>On the one hand, users can set configuration parameters to the device shadow through APIs. When a device is online or goes online, the configuration parameters can be obtained from the device shadow. On the other hand, devices can report their statuses to the device shadow. When you query the device status, you only need to query the device shadow instead of communicating with the device.</p>
Secure information transmission	LINK provides authorization certification for devices and applications and bidirectional binding authorization for topics to ensure device security and uniqueness. It provides TLS-based data transmission channels for secure message transmission.

3 Application Scenarios

[Smart Campus Integration](#)

[Industrial Internet Integration](#)

[Application & Data Integration of Corporation Groups](#)

3.1 Smart Campus Integration

Many difficulties are encountered in smart campus management:

- Customized management systems hinder information collection and sharing.
Buildings in a campus have different structures. Enterprises can customize subsystems for each building to collect all information on each one. However, after customization, the differences between subsystems hinder information collection and sharing, resulting in difficulty in information transmission. This reduces the "smart" level of a campus.
- Diversified devices and complex data collection make it difficult to implement system linkage.
In scenarios such as vehicle entrance and exit management, visitor registration, and campus asset management, it is difficult to implement linkage management due to the complexity of data collection and centralization.
- The status of important devices cannot be remotely monitored in real time, so warnings cannot be generated.
For example, faulty street lamps cannot be alerted and must be manually repaired in a traditional campus, resulting in passive maintenance.

ROMA Connect has a complete set of integration solutions involving devices, data, and services to help enterprises build smart campuses.

- **Efficient interconnection with various devices from different vendors**
Information about devices from different vendors, such as cameras, turnstiles, and air conditioners, is sent to LINK using the standard MQTT protocol. In addition, LINK is connected to multiple IoT platforms, eliminating the need to collect data from each platform separately.

- Data base construction for providing standard data services**


FDI and MQS quickly integrate all data and open the data to different backend services of an enterprise. For example, vehicle data in turnstile systems, device status in asset management systems, and switch-on/switch-off and device information in street lamp systems are transmitted to backend services in real time or in asynchronous mode for analysis and linkage management.

In addition, the high scalability design provided by ROMA Connect supports huge data transmission and storage on the campus network, improving data transmission efficiency.
- Integration of IT, OT, and AI for building an intelligent operation center**

ROMA Connect provides a channel for data integration and sharing. Enterprises, then, can use the enterprise-grade AI, video analysis, and big data services to build a smart campus.
- Centralized and distributed architecture for supporting campus services**

Enterprises holding large campuses often need to manage multiple campuses. The centralized and distributed architecture of ROMA Connect helps these enterprises integrate data from multiple campuses onto the same platform and assists them in managing the distributed and centralized operations based on actual conditions.

Figure 3-1 Smart campus integration



3.2 Industrial Internet Integration

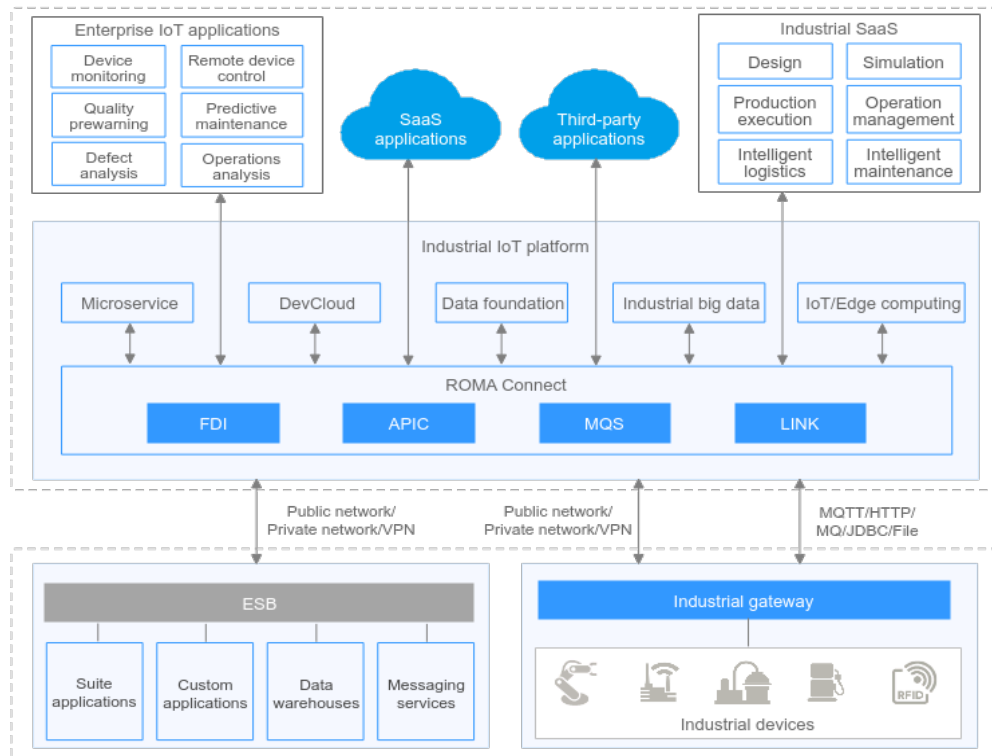
There are several typical problems in the digital transformation of the manufacturing industry:

- **Difficulty in integrating device and environment data**
To monitor and manage production devices of various brands and types in real time, device and environment data need to be collected and uploaded. However, such devices use different data formats and database standards, which makes it difficult to integrate device and environment data.
- **Difficulty in preventing device faults**
In a factory, any machine fault may have a huge impact on the entire assembly production line.
- **Difficulty in optimizing production strategies and decision-making of enterprises**
Different formats of collected data result in difficult data analysis. Therefore, it is a challenge for enterprises to optimize existing production strategies based on the collected data and to determine whether to execute new production strategies.

Leveraging the enterprise-class big data analysis solution, ROMA Connect helps the manufacturing industry transform to IoT integration through data collection and integration, and finally achieves the "smart" vision.

- **Device digitalization and integration**
ROMA Connect uses different methods, such as MQTT, and gateway, to connect various types of devices to enterprise backends, implementing bidirectional communication.
- **Fault prediction and alarming**
Information about all devices is integrated on the ROMA Connect console for real-time monitoring and prewarning analysis. Once parameters of a device become abnormal, ROMA Connect generates an alarm on the console and notifies the owner of repairing the device. If the real-time status of a device deviates from the normal data range, a notification is sent to device maintenance personnel to repair the device in a timely manner.
- **Data conversion and analysis**
ROMA Connect FDI imports data generated by industry SaaS services to ROMA Connect and transmits the data to **MapReduce Service (MRS)**, helping enterprises analyze big data and optimize production strategies.

Figure 3-2 Industrial Internet integration



3.3 Application & Data Integration of Corporation Groups

The integration between a parent company and its subsidiaries and between the corporate group and its partners faces bottlenecks:

- **Geographical differences**
The headquarters, branches, and partners are located in different regions and use different time zones. This reduces the timeliness and reliability of data.
- **Different cloud services**
The cloud services used by the headquarters, branches, and partners are different. Therefore, it is difficult to invoke different cloud services.
- **Network differences**
The networks used by the headquarters, branches, and partners are different. Therefore, interconnection between the public networks, private networks, and VPNs is difficult.

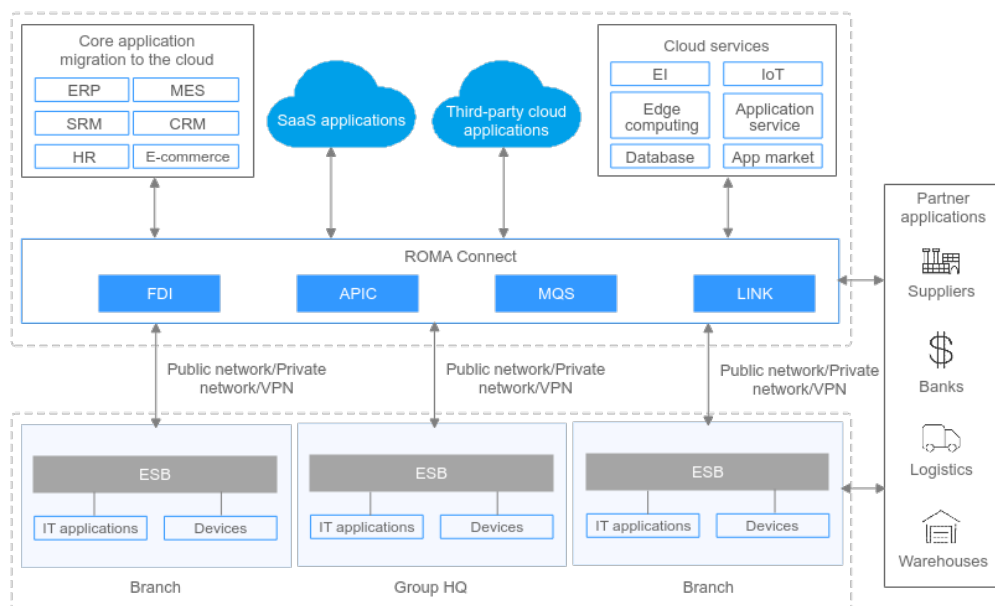
ROMA Connect helps corporation groups implement integration between the headquarters and branches and between the groups and their partners. As shown in [Figure 3-3](#), ROMA Connect supports the following scenarios:

- **Cross-regional integration:** The headquarters, branches, and partners located in different regions transmit their device information, data, and messages to ROMA Connect. ROMA Connect performs operations such as device information visualization, alarm monitoring, data conversion, and message transmission to streamline regional restrictions, implement integration and

governance for regional businesses and share group information, ensuring the reliability of service integration.

- **Cross-cloud integration:** APIC converts SaaS applications and third-party cloud applications into API data. Then, enterprises call these APIs to integrate different cloud applications, ensuring seamless interconnection between services on the cloud.
- **Cross-network integration:** ROMA Connect is used to implement secure cross-network interconnection with partners' service systems. Enterprises upload data and information required by partners to ROMA Connect. ROMA Connect then converts the data formats and integrates data based on the partners' requirements. After an enterprise integrates data and messages, partners can access ROMA Connect to obtain related information.

Figure 3-3 Application & data integration of corporation groups



Application & data integration through ROMA Connect brings the following benefits to enterprises:

- Builds a unified platform for managing multiple cloud services and applications, simplifying management processes and helping enterprises achieve digital transformation.
- Enables information sharing between headquarters, branches, and partners.
- Supports large-scale integrated services, distributed deployment, automatic scaling, and low latency, ensuring service performance and reliability.

4 Edition Differences

This section lists the specifications of ROMA Connect and its components. Use ROMA Connect according to the specifications to reduce system exceptions.

Edition Specifications

The following table lists the ROMA Connect instance specifications in each edition.

Table 4-1 Instance edition specifications

Instance Edition	Number of Systems Supported	Number of Connections Supported	Applicability
Basic	5 to 10	25	Small enterprises
Professional	10 to 20	80	Small- and medium-sized enterprises
Enterprise	20 to 30	200	Medium- and large-sized enterprises
Platinum	More than 30	800	Large enterprises

The numbers of connections and systems listed are for reference only. For details about the number of resources (such as data integration tasks, APIs, and message topics) that can be created, see [Quota Limits](#). To ensure the performance of ROMA Connect, create and use resources within the specified specifications.

- **Number of systems:** A system refers to a user's service system, and the number of systems refers to the number of service systems interconnecting with a ROMA Connect instance. You can set up multiple connections between a service system and a ROMA Connect instance.
- **Number of connections:** A connection refers to an interaction between a service system and ROMA Connect. The number of connections varies depending on the functional module in ROMA Connect that you want to

connect. The following table describes the mappings between the number of resources and the number of connections.

Table 4-2 Mappings between the number of resources and the number of connections

Function	Mapping
FDI	Two FDI tasks in the running state occupy one connection.
APIC	<ul style="list-style-type: none">• Ten hosting APIs (APIs not published by custom backends) occupy one connection.• Five function backends or data backends occupy one connection.
MQS	Three topics occupy one connection.
LINK	1000 devices occupy one connection.

FDI Specifications

The following table lists the read and write performance of each data source when a single task is running in an instance (for reference only). The running performance of a single task is also affected by factors such as the network bandwidth and data source server performance. When multiple tasks are running concurrently in an instance, the performance deteriorates compared with that of a single running task as multiple tasks preempt CPU and memory resources.

- **Common tasks**

The following table lists the reference performance of different types of data sources of common data integration tasks supported by ROMA Connect.

Data Source	Read Rate (MB/s)	Write Rate (MB/s)
MRS Hive	5	2
MRS HDFS	5	2
DWS	5	2
MySQL	6	3
Oracle	6	2
Kafka	10	8
SQL Server	6	3
PostgreSQL	4	2
Gauss100	6	3
FTP	5	3

Data Source	Read Rate (MB/s)	Write Rate (MB/s)
OBS	6	3
MongoDB	0.8	0.3
Redis	/	2
HANA	6	3
API	/	/

 **NOTE**

- When the DWS data source is used at the destination, the larger the destination tables, the slower the write.
 - The write and read rates of an API data source are directly related to the server API response speed.
 - In the performance test, a message of 1 KB is used. In actual application scenarios, the rate is calculated based on 1 KB for messages within this limit.
- **Composite tasks**

The following table lists the reference performance of composite data integration tasks supported by ROMA Connect.

Table 4-3 Data integration from Oracle to DWS

Test Condition			Test Result
Number of table fields (columns)	Number of inserted data records	Data size (KB)	E2E rate (MB/s)
12	1 million	1	1.2
50	1 million	1	0.8
100	1 million	1	0.4
200	1 million	1	0.2

APIC Specifications

The following table lists the APIC specifications supported by a ROMA Connect instance.

 NOTE

The APIC specifications are obtained by testing in the following conditions:

- Connection protocol: HTTPS
- Connection: persistent connection
- Concurrency: greater than or equal to 1000
- Authentication: none
- Size of the returned data: 1 KB
- Bandwidth: 10 MB/s
- Average backend response latency: less than or equal to 10 ms

Table 4-4 APIC specifications

Instance Edition	API Forwarding (TPS)	Function API (TPS)	Data API (TPS)
Basic	4,000	400	400
Professional	6,000	600	600
Enterprise	8,000	800	800
Platinum	10,000	1,000	1,000

 NOTE

Instances of minimal specifications require 10 MB/s bandwidth to meet performance requirements. The bandwidth size required is dynamically increased for the instance specifications, number of requests, and request and response body sizes.

MQS Specifications

Open-source compatibility: ROMA Connect is fully compatible with open-source Kafka 1.1.0 and 2.7 and their APIs. It has all message processing features of native Kafka.

The following table lists the MQS specifications supported by a ROMA Connect instance. When selecting the specifications, you are advised to reserve 30% of the bandwidth to ensure stable running of your applications.

 NOTE

The MQS specifications are obtained by testing in the following conditions:

- Connection: intranet
- Authentication: none
- Data size: 1 KB
- Disk type: SSD

Table 4-5 MQS specifications

Instance Edition	Bandwidth	TPS (High - Throughput)	TPS (Synchronous Replication)	Maximum Number of Partitions	Storage	Specifications
Basic	100 MB/s	100,000	60,000	600	600 GB	Recommended for up to 3000 client connections, 60 consumer groups, and service traffic of 70 MB/s.
Professional	300 MB/s	300,000	150,000	900	1200 GB	Recommended for up to 10,000 client connections, 300 consumer groups, and service traffic of 210 MB/s.
Enterprise	600 MB/s	600,000	300,000	1800	2400 GB	Recommended for up to 20,000 client connections, 600 consumer groups, and 420 MB/s service traffic.
Platinum	1200 MB/s	1.2 million	400,000	1800	4800 GB	Recommended for up to 20,000 client connections, 600 consumer groups, and service traffic of 840 MB/s.

LINK Specifications

ROMA Connect supports device access using MQTT 3.1 and MQTT 3.1.1. The following table lists the LINK specifications supported by an instance.

NOTE

The LINK specifications are obtained by testing in the following conditions:

- Upstream message
 - Connection: intranet
 - Message size: 500 bytes
 - Message destination: MQS topic
- Downstream message
 - Connection: intranet
 - Message size: 500 bytes
 - Delivery mode: Use the demos downloaded from the console to call data plane APIs for message delivery.

Table 4-6 LINK specifications

Instance Edition	Upstream Message	Downstream Message
Basic	10,000 TPS for 20,000 online devices	1000 TPS for 20,000 online devices
Professional	20,000 TPS for 40,000 online devices	1500 TPS for 40,000 online devices
Enterprise	40,000 TPS for 100,000 online devices	2000 TPS for 100,000 online devices
Platinum	100,000 TPS for 450,000 online devices	5000 TPS for 450,000 online devices

5 Edition Differences (New Version)

Recommended Edition Specifications

The following table lists the recommended specifications for .

Table 5-1 Recommended specifications

Number of Systems Supported	Estimated Number of RCUs	Applicability
Less than 5	5	Small enterprises
5 to 10	20	Small enterprises
10 to 20	40	Small- and medium-sized enterprises
20 to 30	80	Medium- and large-sized enterprises
More than 30	150	Large enterprises

The numbers of RCUs and systems listed are for reference only. For details about the number of resources (such as application flows, APIs, and message topics) that can be created, see [Quota Limits](#). To ensure the performance of ROMA Connect, create and use resources within the specified specifications.

- **Number of systems:** A system refers to a user's service system, and the number of systems refers to the number of service systems interconnecting with a ROMA Connect instance. Multiple connections can be established between a service system and a ROMA Connect instance.
- **Number of RCUs:** The ROMA Compute Unit (RCU) is the capability compute unit of the new version. Each RCU can be allocated to different integration capabilities, including FDI, APIC, MQS, LINK, and composable applications. The performance specifications of each integration capability depend on the number of RCUs allocated. More RCUs indicate higher specifications.

Note that minimum RCUs and adjustments vary with the integration capability.

Specifications

The specifications of each integration capability of the new version vary with the number of allocated RCUs. [Table 5-2](#) lists the specifications of each RCU.

NOTE

- The following performance data comes from lab tests and may differ from that in actual service scenarios.
- Once FDI is enabled, the number of FDI tasks for both **message and non-message data** is fixed to 1000 and does not change with RCUs.

Table 5-2 Specifications

Integration Capability	Indicator	Specifications/ RCU	Minimum RCUs	Minimum Adjustment
Composite Application (Flow Running)	Number of running flows	30	2	1
APIC (API Management)	API forwarding	1000 TPS	2	1
	Number of hosted APIs	100		
APIC (Data/Function API)	Data/Function API forwarding	100 TPS	2	1
	Number of custom backends	50		
	API traffic	0.24 Mbit/s		
FDI	Non-message data traffic	About 2 Mbit/s, up to 4 Mbit/s per task	4	2
	Message data traffic	About 5 Mbit/s, up to 10 Mbit/s per task		
MQS	Message traffic	30 Mbit/s	3	1 (Increasing only)
	Disk	500 GB		
	Number of partitions	250		
LINK	Number of online devices	10,000	2	1

Integration Capability	Indicator	Specifications/RCU	Minimum RCUs	Minimum Adjustment
	Number of concurrent device messages	2,500 TPS		

 NOTE

The MQS specifications are obtained by testing in the following conditions:

- Connection: intranet
- Authentication: SASL_SSL
- Data size: 1 KB/10 KB
- Disk type: ultra-high I/O, SSD

The LINK specifications are obtained by testing in the following conditions:

- Upstream message
 - Connection: intranet
 - Message size: 512 bytes
 - Message destination: MQS topic

The FDI specifications are obtained by testing in the following conditions:

- Three concurrent tasks in each of the following scenarios:
 - MySQL to MySQL
 - OBS to OBS
 - Kafka to Kafka
- Data read and write: 1 million records
- Data migration volume: 1 GB

The APIC specifications are obtained by testing in the following conditions:

- Backend type: shubao
 - Connection: persistent/short connection
 - Connection protocol: HTTP/HTTPS
 - Concurrency: at least 4000
 - Authentication: none
 - Size of the returned data: 1 KB
 - Bandwidth: 10 MB/s
 - Average backend response latency: 10 ms
- Backend Type: LiveData
 - Connection: persistent/short connection
 - Connection protocol: HTTP/HTTPS
 - Concurrency: at least 400
 - Authentication: none
 - Size of the returned data: 1 KB
 - Bandwidth: 10 MB/s
 - Average backend response latency: 10 ms

6 Supported Data and Protocols

FDI

Table 6-1 lists the data sources supported by FDI tasks.

Table 6-1 Data sources supported by FDI

Data Source	Version	Common Task Source	Common Task Destination	Composite Task Source	Composite Task Destination
API	-	Yes	Yes	No	No
ActiveMQ	5.15.9	Yes	Yes	No	No
ArtemisMQ	2.9.0	Yes	Yes	No	No
DB2	9.7	Yes	Yes	No	No
DIS	-	Yes	Yes	No	No
DWS	1.3.4	Yes	Yes	No	No
DM	-	Yes	Yes	No	No
FTP	-	Yes	Yes	No	No
Gauss100	FusionInsight_LibrA_V100R003C20, FusionInsight_LibrA_V300R001C00	Yes	Yes	No	No
GaussDB(for MySQL)	2.0.15.6	Yes	Yes	No	No

Data Source	Version	Common Task Source	Common Task Destination	Composite Task Source	Composite Task Destination
HL7	2.1, 2.2, 2.3, 2.3.1, 2.4, 2.5, 2.6, 2.7, 2.8, 2.8.1	Yes	Yes	No	No
HANA	1.0	Yes	Yes	Yes	Yes
IBM MQ	9.1	Yes	Yes	No	No
Kafka	1.1.0, 2.3.0	Yes	Yes	No	Yes
LDAP	-	Yes	No	No	No
MongoDB	3.4	Yes	Yes	No	No
MQS	N/A	Yes	Yes	No	No
MRS Hive	MRS 3.**	Yes	Yes	No	No
MRS HDFS	MRS 3.**	Yes	Yes	No	No
MRS HBase	MRS 3.**	Yes	Yes	No	No
MRS Kafka	MRS 3.**	Yes	Yes	No	No
MySQL	5.7, 8.0	Yes	Yes	Yes	Yes
OBS	3	Yes	Yes	No	No
Oracle	11.2g (not recommended), 12.1g (not recommended), 12.2g, 19c	Yes	Yes	Yes	Yes
PostgreSQL	11	Yes	Yes	Yes	Yes
RabbitMQ	3.6.10	Yes	Yes	No	No
RocketMQ	4.7.0	Yes	Yes	No	No
Redis	3.0.7, 4.0.11	No	Yes	No	No
SAP	SAP Java Connector 3.0.19	Yes	No	No	No
SNMP	v1, v2, or v3	Yes	No	No	No

Data Source	Version	Common Task Source	Common Task Destination	Composite Task Source	Composite Task Destination
SQL Server	2014, 2019, and 2022	Yes	Yes	Yes	Yes
WebSocket	-	Yes	No	No	No
Custom data source	-	Yes	Yes	No	No

APIC

- APIC creates and opens APIs, supporting the following request protocols: RESTful, SOAP, and WebSocket.
- [Table 6-2](#) lists the data sources supported by APIC custom backends.

Table 6-2 Data sources supported by custom backends

Data Source	Version
DWS	1.3.4
Gauss100	FusionInsight_LibrA_V100R003C20, FusionInsight_LibrA_V300R001C00
HANA	1.0
HIVE	2.3.2
MongoDB	3.4
MySQL	5.7 and 8.0
MRS HBase	MRS 3.*.*
MRS Hive	MRS 3.*.*
Oracle	11.2g (not recommended), 12.1g (not recommended), 12.2g, 19c
PostgreSQL	11.0
Redis	3.0.7, 4.0.11
SQL Server	2014, 2019, and 2022

MQS

[Table 6-3](#) lists the message types supported by MQS.

Table 6-3 Message types supported by MQS

Message Type	Version
Kafka	1.1.0, 2.7

LINK

Table 6-4 lists the device access protocols supported by LINK.

Table 6-4 Access protocols supported by LINK

Message Type	Version
MQTT	3.1, 3.1.1
Modbus	-
OPC UA	-

7 Quotas

Quota Limits

A quota refers to the maximum number of resources that you can create in a ROMA Connect instance. The following table lists the resource quotas.

 **NOTE**

The maximum quota may be slightly exceeded in case of high concurrency, but resource usage will not be affected.

Table 7-1 Resource quotas

Component	Resource	Maximum Quota	Modifiable
Integration application	Number of integration applications	100	No
Data source	Number of data sources	500	No
FDI	Number of data integration tasks	1,000	No
APIC	Number of APIs	<ul style="list-style-type: none">• Basic: 250• Professional: 800• Enterprise: 2000• Platinum: 8000	Yes
	Number of API groups	1500	Yes

Component	Resource	Maximum Quota	Modifiable
	Number of APIs in a single API group	<ul style="list-style-type: none"> • Basic: 250 • Professional: 800 • Enterprise: 2000 • Platinum: 5000 	Yes
	Number of environment variables in a single API group	50	Yes
	Number of request throttling requests	300	Yes
	Number of access control policies	100	Yes
	Number of environments	10	Yes
	Number of signature keys	200	Yes
	Number of load balance channels	200	Yes
	Number of ECSs in a load balance channel	10	Yes
	Number of custom authorizers	50	Yes
	Number of custom backends	<ul style="list-style-type: none"> • Basic: 125 • Professional: 400 • Enterprise: 1000 • Platinum: 4000 	Yes
	Number of client quota policies	20,000	Yes
LINK	Number of product templates	100	Yes
	Number of products	500	Yes

Component	Resource	Maximum Quota	Modifiable
	Number of devices	1,000	Yes
	Number of rules	2,000	No

8 Constraint

FDI

Table 8-1 Constraints

Function	Constraints
Data synchronization	<ul style="list-style-type: none"> ● A single data record cannot exceed 8 MB. ● The fields in time format support down-to-second precision. ● A table name cannot contain hyphens (-) or number signs (#). ● Modifying table structures after a task is started leads to task failure. In this case, restart the task. ● The destination does not support tables whose mapping fields are all primary keys. ● Up to 800 MB of files can be collected by concurrent tasks. <p>NOTE This constraint applies only to OBS, FTP, and MRS HDFS data sources. For example, if two OBS tasks and two FTP tasks are concurrently executed, the total size of files to be collected from the four tasks cannot exceed 800 MB.</p>
FTP data source	<p>If parsing is enabled, each file cannot exceed 200 MB (files exceeding 200 MB will be automatically skipped) and up to 1,500,000 data records can be parsed. If parsing is disabled, each file cannot exceed 6 MB and up to 20,000 files can be collected.</p> <p>NOTE Statistics on multiple files synchronization between FTP data sources indicate the number of files synchronized this time.</p>
OBS data source	<p>If parsing is enabled, each file cannot exceed 200 MB (files exceeding 200 MB will be automatically skipped). If parsing is disabled, each file cannot exceed 10 MB.</p>

Function	Constraints
MRS data source	<ul style="list-style-type: none"> Only MRS clusters authenticated by Kerberos can be connected. Only structured data is supported.
MRS Hive data source	<ul style="list-style-type: none"> Hive supports only RCFile and TextFile read and write. When MRS Hive serves as the source, only tables of up to 1 million records can be synchronized.
API data source	<ul style="list-style-type: none"> If the server does not respond in 60 seconds, an error is reported during task execution. When the API serves as the source, every request can read up to 20 MB data. Otherwise, paging must be enabled. When the API serves as the source, only constant parameters are supported. Dynamic parameter transfer is not supported. When the API serves as the destination, the source data cannot be mapped to the destination headers.
Kafka data source	SASL can be used only to connect ROMA Connect MQS. The Kafka service developed by users themselves supports only plaintext connection.
Oracle database	<ul style="list-style-type: none"> Only the following field types are supported. Fields support only uppercase letters. CHAR, VARCHAR, DATE, NUMBER, FLOAT, LONG, NCHAR, NVARCHAR2, RAW, TIMESTAMP The system time difference between the Oracle system and the ROMA Connect server must be less than 2 minutes.
SQL Server database	Only the following field types are supported: BIT, CHAR, DATE, DATETIME, DECIMAL, FLOAT, IMAGE, INT, MONEY, NUMERIC
MySQL database	Only the following field types are supported: INT, BIGINT, TINYINT, MEDIUMINT, FLOAT, DOUBLE, DECIMAL, CHAR, VARCHAR, TINYTEXT, TEXT, MEDIUMTEXT, LONGTEXT, DATETIME, TIMESTAMP, SMALLINT, YEAR, BINARY, JSON
PostgreSQL/ DWS database	<p>Only the following field types are supported: BOOL, CIDR, CIRCLE, DATE, NUMERIC, FLOAT4, FLOAT8, MONEY, PATH, POINT, INT, TIMESTAMP, TIMETZ, UUID, VARBIT, VARCHAR</p> <p>For better write performance, destination data sources do not support the batch number and constant features by default. To enable these two features, contact technical support.</p>
Redis database	When the Redis data source is at the destination, to write a time field (datetime/date) into the yyyy-MM-dd HH:mm:ss format, set the destination field type to string
DIS database	Each channel supports only one task to collect source data.

Function	Constraints
WebSocket database	When you create a data integration task and set Parse to Yes , Parsing Path in Metadata must be configured. Otherwise, the task will fail.
Relational database	A maximum of 10 million data records can be synchronized.
Composite task (CDC)	<ul style="list-style-type: none"> ● Source Scheduled: MySQL, Oracle, SQL Server, PostgreSQL, and HANA. Real-time: MySQL, Oracle, and SQL Server ● Destination Scheduled: MySQL, Oracle, PostgreSQL, SQL Server, and HANA Real-time: MySQL, Oracle, PostgreSQL, Kafka, and SQL Server ● The destination table must have a primary key. Otherwise, data synchronization will be affected. ● The Oracle data source at the source can contain only tables with uppercase table names and field names. ● The Oracle data source at the destination cannot contain tables with lowercase field names. ● When you modify a composite task and add a source table to it, the source table must contain data. ● Each table name can include up to 64 characters for composite tasks. ● Automatic tasks map the first 2,000 source/destination tables and will fail if delayed over 1 minute by performance, load, or network issues. If that happens, try manual mapping. ● Binary fields are not supported when defining real-time composite tasks. ● A MySQL schema task supports a maximum of 10 groups of mappings, and a maximum of 2,000 tables are supported.
Flow task	<ul style="list-style-type: none"> ● Destination tables cannot be cleared each time a task is executed. ● Constants cannot be specified for the destination fields. ● When you create a task with multiple destinations, the value of Batch Number Format of the first connection applies to all the connections. For example, if UUID is set for the first connection while yyyyMMddHHmmss is set for the second and the third, UUID is used for all the three connections.

APIC

Table 8-2 Constraints

Function	Constraint
Data API response body size	The data body returned by a data API cannot exceed 10 MB.
Number of data records returned by a data API	By default, a data API obtains 2,000 records from the database. The excessive records cannot be returned.
Paging of data API results	When paging is enabled, a maximum of 2,000 data records can be obtained at a time.
Custom backend header	X-Apig-Mode is taken up and not recommended for services.
Request body size of a hosting API for transparent transmission	The request body of a hosting API cannot exceed 2 GB.
Function API HTTP Client request timeout	The maximum timeout period is 30s and cannot be changed.
Cross-domain request	If the OPTIONS request is accessed using an IP address, the inbound IP address of ROMA Connect cannot be mapped. To map the inbound IP address, domain name access is required.
Signature authentication request body size	The size of the request body can be configured in the instance configurations, ranging from 1 to 9536 MB. However, during App authentication development, only requests whose body does not exceed 12 MB can be accessed. Otherwise, the signature will fail.
Sandbox memory calculation	The APIC sandbox memory size cannot be accurately calculated due to the underlying JVM. It is an approximate value.
New and modified resources	It takes 5 to 10 seconds for a new or modified APIC resource to take effect.

MQS

Table 8-3 Constraints

Function	Constraint
Message size	The maximum size is 10 MB.
Faulty node	If some nodes in the instance are faulty, topic management (such as creation and deletion) cannot be performed.
Topic import	<ul style="list-style-type: none">• Only XLSX, XLS, and CSV files can be imported.• The description in the files to be imported cannot start with an equal sign (=). Newline characters contained in the description will be escaped.• The number of topics in a file to be imported cannot exceed 100.
Topic export	Only XLSX, XLS, and CSV files can be exported.
Number of messages to be queried	A maximum of 500 messages can be queried at a time.
Number of connections of each IP address	A maximum of 1,000 connections can be created for each client IP address through private networks, and the same constraint applies to each instance through public networks.
Topic aging time	When you create or modify a topic on the console, the maximum aging time is 168 hours.

LINK

Table 8-4 Constraints

Function	Constraint
Maximum size of a message reported by a device	512 KB
Maximum size of a message delivered by a command	
File types supported by device import and export	CSV

Function	Constraint
File types supported by product import and export	
File types supported by rule import and export	
Maximum size of a device import file	200 MB
Maximum size of a product import file	
Maximum size of a rule import file	
Server MQTT QoS levels	Only QoS 0 and QoS 1 are supported. QoS 2 is not supported.
Modbus device usage	Command delivery is not supported.
Device access protocol	Only MQTT, OPC UA, and Modbus are supported.

9 Billing

Billing Item

ROMA Connect is billed based on ROMA Connect instance editions and service duration.

Table 9-1 ROMA Connect billing item

Billing Item	Description
ROMA Connect instance	Billed based on the instance edition, required duration, and the number of instances you have purchased

Billing Mode

ROMA Connect supports the yearly/monthly billing mode. You can pay for your instance at a time based on the required duration.

Renewal

ROMA Connect is billed on a yearly/monthly basis. Before a purchased ROMA Connect instance expires, you can renew it to extend its validity period, or you can set auto-renewal rules for the ROMA Connect instance. For more information about renewing resource packages, see [Renewal Management](#).

Expiration and Overdue Payment

ROMA Connect is billed on a yearly/monthly basis. If you do not renew your subscription after a ROMA Connect instance expires, the cloud service platform provides a retention period. For details about the retention period and rules, see [Stopping Services and Releasing Resources](#). If you do not renew your subscription when the retention period ends, data in the ROMA Connect instance will be deleted and cannot be restored.

10 Billing (New Version)

Billing Item

The new version is billed on the number of RCUs and the usage duration.

Table 10-1 ROMA Connect billing items

Billing Item	Description
RCU	ROMA Computing Unit, billed by quantity and duration.
Bandwidth (optional)	Public network bandwidth is required for accessing ROMA Connect from the public network, and is billed based on the bandwidth size and duration.

Billing Mode

Pay-per-use billing and packages are available for new-version ROMA Connect instances.

- **Pay-per-use:** The instance is billed based on the RCU quantity and duration (by hour). If the duration is less than an hour, billing is based on the actual usage.
- **Packages:** A package contains a specified RCU-hour quota. A ROMA Connect instance of the new version consumes the RCU-hours in the package first, and the overflow (if any) is billed on a pay-per-use basis. The package quota is valid only in one effective month. The remaining quota expires and cannot be transferred to the next month.

If you select **1,000 RCU Hours/Month** for a five-RCU instance, you can use the instance for 200 hours in the effective month. After the 200 hours, the instance is billed in the pay-per-use mode. If there are multiple instances, the total number of RCUs are calculated.

Renewal

You can renew your resources on the [Renewals](#) page of the management console. For details, see [Renewal Management](#).

Expiration and Overdue Payment

With pay-per-use billing, fees are deducted every hour. If your account balance is insufficient to pay for the expense occurred in the previous hour, your account will be in arrears, and the cloud service platform has a retention period. If the instances are renewed within the retention period, you can continue to use ROMA Connect. Otherwise, services will be stopped and instance resources will be released after the retention period. For details about the retention period and rules, see [Stopping Services and Releasing Resources](#).

ROMA Connect cannot be used after arrears. It is recommended that you renew your account promptly.

11 Permissions

If you need to assign different permissions to employees in your enterprise to access your Huawei Cloud ROMA Connect resources, IAM is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you secure access to your resources.

With IAM, you can use your Huawei Cloud account to create IAM users for your employees, and assign permissions to the users to control their access to specific resource types. For example, some software developers in your enterprise need to use ROMA Connect resources but should not be allowed to delete them or perform any high-risk operations. In this scenario, you can create IAM users for the software developers and grant them only the permissions required for using ROMA Connect resources.

If your Huawei Cloud account does not need individual IAM users for permissions management, you may skip this chapter.

IAM is free of charge. You pay only for the resources in your account. For more information about IAM, see [IAM Service Overview](#).

ROMA Connect Permissions

By default, new IAM users do not have any permissions assigned. To assign permissions to these new users, add them to one or more groups, and attach permissions policies or roles to these groups.

ROMA Connect is a project-level service deployed and accessed in specific physical regions. To assign permissions to a user group, specify the scope as region-specific projects and select projects for the permissions to take effect. If **All projects** is selected, the permissions will take effect for the user group in all region-specific projects. When accessing ROMA Connect, the users need to switch to a region where they have been authorized to use this service.

You can grant users permissions by using roles and policies.

- **Roles:** A type of coarse-grained authorization mechanism first provided by IAM to define permissions related to user responsibilities. There are only a limited number of roles for granting permissions to users. When using roles to grant permissions, you may also need to assign other roles on which the

permissions depend. However, roles are not an ideal choice for fine-grained authorization and secure access control.

- Policies: A type of fine-grained authorization mechanism lately provided by IAM to define permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization and secure access control. For example, regarding ROMA Connect services, you can grant IAM users only the permissions for managing instance resources. Most policies define permissions based on APIs. For details, see [supported actions of fine-grained policies](#).

Table 11-1 lists all the system roles supported by ROMA Connect.

Table 11-1 ROMA Connect system permissions

Role/Policy Name	Description	Type	Dependency
ROMA FullAccess	Full permissions for ROMA Connect. Users granted these permissions can operate and use all ROMA Connect instances.	System-defined policy	None
ROMA CommonOperations	Common user permissions for ROMA Connect. This does not include permissions to create, modify, and delete instances.	System-defined policy	None
ROMA ReadOnlyAccess	Read-only permissions for ROMA Connect. Users granted these permissions can only view ROMA Connect data.	System-defined policy	None
ROMA Administrator	All permissions for ROMA Connect. Users granted these permissions can operate and use all ROMA Connect instances.	System-defined policy	None

Table 11-2 lists the common operations supported by each system-defined policy of ROMA Connect. Select the proper system-defined policies as required.

Table 11-2 Common operations supported by each system-defined policy or role of ROMA Connect

Operation	ROMA FullAccess	ROMA CommonOperations	ROMA ReadOnlyAccess	ROMA Administrator
Creating a ROMA Connect instance	√	x	x	√
Querying instance information	√	√	√	√
Modifying a ROMA Connect instance	√	x	x	√
Deleting a ROMA Connect instance	√	x	x	√
Operating resources in an instance	√	√	x	√

Integration Application Permissions

ROMA Connect provides strict permissions management for user resources. In one instance, IAM users can view and manage only the integration applications and resources created by themselves. With integration application authorization, IAM users can share applications and resources with other IAM users under the same account.

Table 11-3 Application permissions

Permission	FDI	APIC	MQS	LINK
read	View data sources of applications.	View, debug, and export APIs of applications.	View and export topics of applications.	View and export devices, products, and rules of applications, as well as debug devices.

Permission	FDI	APIC	MQS	LINK
modify	Create and edit data sources of applications.	Create, edit, release, take APIs offline, and import APIs of applications.	Create and edit topics of applications.	Create, edit, and import devices, products, and rules of applications, as well as reset device and product passwords.
delete	Delete data sources of applications.	Delete APIs of applications.	Delete topics of applications.	Delete devices, products, and rules of applications, product properties, device topics, as well as rule data sources and destinations.
access	N/A	Configure authorization, access control, request throttling, and signature key binding for APIs of applications.	Configure permissions for topics of applications.	Deliver commands to and forcibly take offline devices, as well as configure plug-ins for devices that use the OPC UA or Modbus protocol.
admin	Application administrator permissions.			

ABM Permissions

Table 11-4 lists all the system policies or roles supported by ABM.

Table 11-4 ABM system permissions

Role/Policy Name	Description	Type	Dependency
ROMA ABM FullAccess	Full permissions for ROMA ABM	System-defined policy	None
ROMA ABM CommonOperations	Common user permissions for ROMA ABM. This does not include permissions to create, modify, and delete instances.	System-defined policy	None
ROMA ABM ReadOnlyAccess	Read-only permissions for ROMA ABM	System-defined policy	None

ROMA Connect Business Flow Permissions

Table 11-5 shows all system permissions of a business flow.

Table 11-5 ROMA Connect business flow permissions

System Policy/Policy Name	Description	Type	Dependency
ROMA BFS CommonOperations	Operation (such as create) permissions	System-defined policy	None
ROMA BFS FullAccess	Administrator permissions	System-defined policy	None
ROMA BFS ReadOnlyAccess	Read-only permissions	System-defined policy	None

ROMA Connect Business Flow Diagram Permissions

ROMA Connect implements strict permission management on service flow diagrams. IAM users can view and manage only the service flow diagrams created by themselves. By default, IAM users cannot view the service flow diagrams created by other IAM users. ROMA Connect business flows can be shared with other IAM users under the same account to view and use.

Table 11-6 ROMA Connect business flow diagram permissions

Permission	Business Flow
read	You can view the authorized business flow diagram.
modify	You can modify and save the authorized service flow diagram.
delete	You can delete the authorized business flow diagram.
operate	You can start, stop, restart, and manually schedule the authorized business flow diagram.
admin	Administrator permissions for business flow diagrams.

12 Basic Concepts

RCU

The ROMA Compute Unit (RCU) is the capability compute unit of the new edition ROMA Connect. Each RCU can be allocated to different integration capabilities, including FDI, APIC, MQS, LINK, and composable applications. The performance specifications of each integration capability depend on the number of RCUs allocated. More RCUs indicate higher specifications.

Connector

A connector is a custom data source plug-in. ROMA Connect supports common data source types, such as relational databases, big data storage, semi-structured storage, and message systems. If the data source types supported by ROMA Connect cannot meet your data integration requirements, you can develop a read/write plug-in to connect to ROMA Connect through a standard RESTful API to enable ROMA Connect to read and write these data sources.

Environment

An environment refers to the usage scope of an API. You can call an API only after you publish it in an environment. You can publish APIs in different custom environments, such as the development environment and test environment. RELEASE is the default environment for formal publishing.

Environment Variable

Environment variables are specific to environments. You can create environment variables in different environments to call different backend services by using the same API.

Load Balance Channel

A load balance channel allows ROMA Connect to access ECSs in the same VPC and use the backend services deployed on the ECSs to expose APIs. In addition, the load balance channel can balance access requests sent to backend services.

Producer

A producer is a party that publishes messages into topics. The messages will be then delivered to other systems for processing.

Consumer

A consumer is a party that subscribes to messages from topics. The ultimate purpose of subscribing to messages is to process the message content. For example, in a log integration scenario, the alarm monitoring platform functions as a consumer to subscribe to log messages from topics, identify alarm logs, and send alarm messages or emails.

Partition

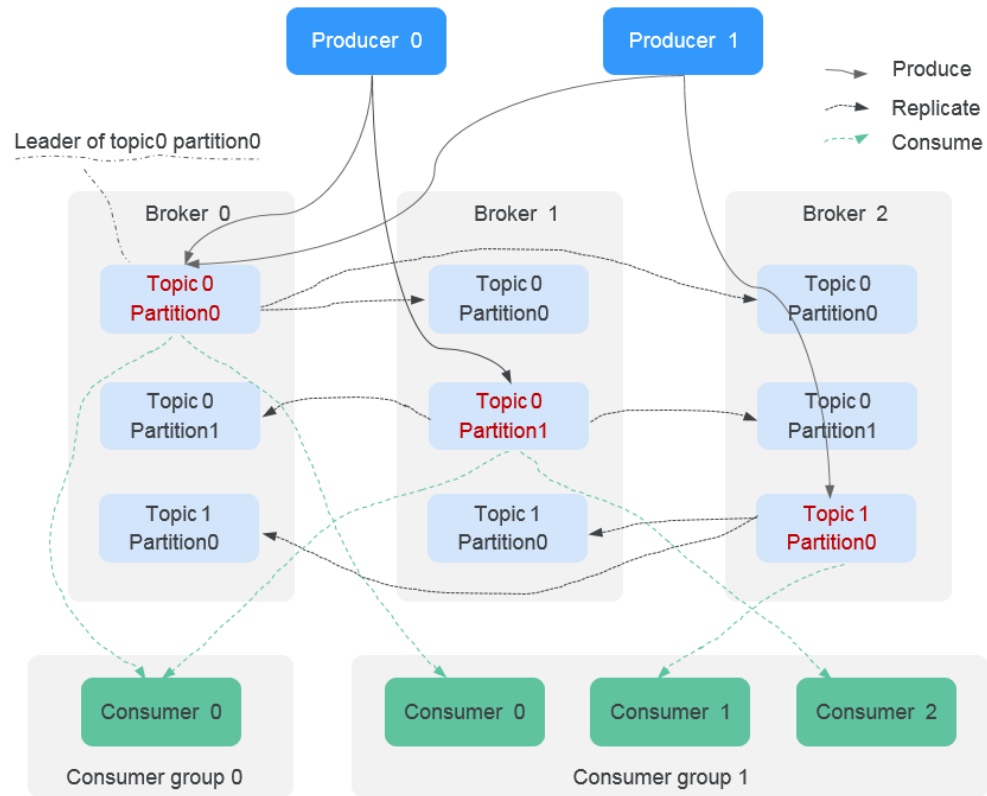
A topic is a place holder of your messages in Kafka and is further divided into partitions. Messages are stored in different partitions in a distributed manner, implementing horizontal expansion and high availability of Kafka.

Replica

To improve message reliability, each partition of Kafka has multiple replicas to back up messages. Each replica stores all data of a partition and synchronizes messages with other replicas. A partition has one replica as the leader which handles the creation and retrieval of all messages. The rest replicas are followers which replicate the leader.

The topic is a logical concept, whereas the partition and broker are physical concepts. The following figure shows the relationship between partitions, brokers, and topics of Kafka based on the message production and consumption directions.

Figure 12-1 Kafka message flow



Topic

A topic is a model for publishing and subscribing to messages in a message queue. Messages are produced, consumed, and managed based on topics. A producer publishes a message to a topic. Multiple consumers subscribe to the topic. The producer does not have a direct relationship with the consumers.

Product

A product is a collection of devices with the same capabilities or features. Each device belongs to a product. You can define a product to determine the functions and attributes of a device.

Thing Model

A thing model defines the service capabilities of a device, that is, what the device can do and what information the device can provide for external systems. After the capabilities of a device are divided into multiple thing model services, define the attributes, commands, and command fields of each thing model service.

Rule Engine

A rule engine allows you to configure forwarding rules so that data reported by devices can be forwarded to other cloud services for storage or further analysis.

13 Related Services

Virtual Private Cloud (VPC)

VPC allows you to create private, isolated virtual networks on a cloud platform. You can configure the CIDR block, subnets, and security groups, assign EIPs, and allocate bandwidth for a VPC.

ROMA Connect runs in a VPC and uses the VPC to manage IP addresses and bandwidth. When you create a ROMA Connect instance, you need to associate it with a VPC, subnet, and security group. To enable public network access for the instance, bind an EIP to the instance.

Data Ingestion Service (DIS)

DIS builds data intake streams for custom applications capable of processing or analyzing streaming data. It provides efficient collection, transmission, and distribution of real-time data from the IoT and the Internet.

ROMA Connect uses DIS as the source and destination of data integration tasks or uses DIS as the destination of the LINK rule engine to forward data from devices to DIS.

Distributed Message Service for Kafka

Distributed Message Service for Kafka (Kafka for short) is a message queuing service based on Apache Kafka. Apache Kafka is a distributed message middleware that features high throughput, data persistence, horizontal scalability, and stream data processing.

ROMA Connect uses Kafka as the source and destination of data integration tasks or uses Kafka as the destination of the LINK rule engine to forward data from devices to Kafka.

MapReduce Service (MRS)

MRS is a cloud service that deploys and manages Hadoop systems. It provides enterprise-grade big data clusters on the cloud. Tenants can fully control clusters and easily run big data components such as Hadoop, Spark, HBase, Kafka, and Storm.

ROMA Connect uses MRS Hive, MRS HDFS, MRS HBase, or Kafka as the source and destination of data integration tasks.

Object Storage Service (OBS)

OBS is an object-based cloud storage service that provides massive, secure, highly reliable, and cost-effective data storage capabilities for users to store data of any type and size.

ROMA Connect uses OBS as the source and destination of a data integration task. It can also store the data that fails to be converted to OBS during the running of data integration tasks.

Distributed Cache Service (DCS)

DCS is an online, distributed, in-memory cache service. It is reliable, scalable, usable out of the box, and easy to manage. Compatible with Redis, DCS can meet your requirements for high read/write performance and fast data access.

ROMA Connect uses Redis as the destination of data integration tasks or encapsulates Redis into APIs and exposes them to external systems.

FunctionGraph

FunctionGraph hosts event-driven functions in a serverless context while ensuring high availability, high scalability, and zero maintenance. All you need to do is write your code and set conditions.

ROMA Connect uses the functions created in FunctionGraph as the backend services of APIs.

Identity and Access Management (IAM)

IAM provides identity authentication, permissions management, and access control on a cloud platform.

With IAM, you can control access to ROMA Connect.

Cloud Eye

Cloud Eye is a secure, scalable, and integrated monitoring service provided by the cloud platform. With Cloud Eye, you can monitor your ROMA Connect service and configure alarm rules and notifications.

Cloud Trace Service (CTS)

CTS generates traces to enable you to get a history of operations performed on cloud service resources. The content of a trace includes ROMA Connect operation requests sent using the management console or open APIs as well as the operation results. You can view all generated traces to query, audit, and backtrack performed operations.