Solution

Getting Started with Intelligent Data Insight

Issue	1.0
Date	2025-05-14

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

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

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

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 Solution Overview

Scenarios

This solution is built on the low-code platform Dify. It combines the inference capabilities of DeepSeek models with backend services to implement a workflow system that transforms users' natural language questions into intelligent data queries, analysis, and visualized feedback. The results are presented in a user-friendly manner, creating an automated data insight assistant for businesses. It is suitable for enterprises that require intelligent decision support in processes such as sales analysis, customer service optimization, and compliance audits. It balances cost and data security to enhance efficient business decision-making.

Solution Architecture

This solution helps you quickly deploy an intelligent data analytics platform. It transforms natural language questions into structured data queries, analytics, and visualization.

Figure 1-1 Solution architecture

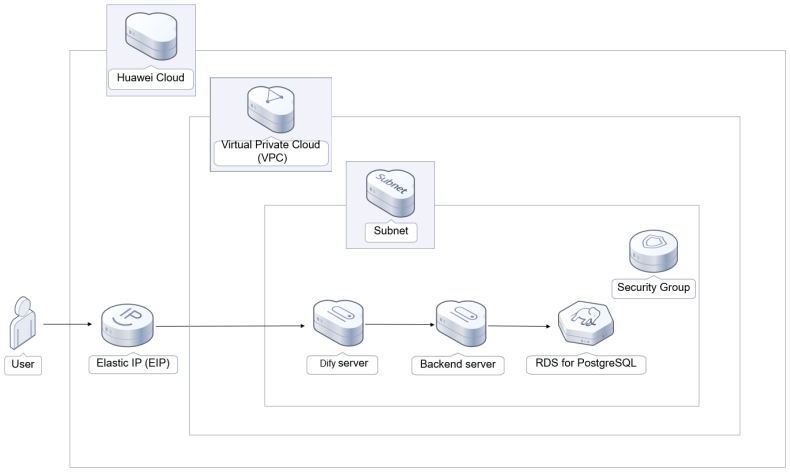

This solution will:

- Create two **Huawei Cloud Flexus X Instances (FlexusX)**, one for deploying Dify and the other for backend services.

- Create two **EIPs** and associate them with the FlexusX instances to enable access to and from the Internet.
- Create an RDS that supports both MySQL and PostgreSQL to provide data sources for intelligent data insight.

Advantages

- Zero-code interaction and low threshold

  Users can directly ask questions in natural language without mastering SQL, Python, or data analysis skills, lowering the usage threshold.
- Intelligent analysis

  Large AI models automatically interpret data and provide a complete report with data overviews, detailed analyses, and visualized conclusions, replacing manual analysis and document writing.
- Ecosystem integration

  The Dify platform allows you to export and publish APIs, which can be quickly embedded into your service systems.

Notes and Constraints

- Before deploying this solution, make sure you have created a HUAWEI ID capable of accessing the target region and enabled Huawei Cloud services.
- If you choose the yearly/monthly billing mode, ensure that your account balance is sufficient for automatic payment during one-click resource deployment. If you do not have sufficient balance, you can enter the **Billing Center** to manually pay for the order.
- If you choose to use IAM agencies to deploy resources, ensure that your Huawei Cloud account has sufficient IAM permissions. For details, refer to **(Optional) Creating the rf_admin_trust Agency**. If you use a HUAWEI ID or an IAM user under the **admin** user group, you do not need to select an agency; the permissions of the currently logged-in user will be used for deployment.

2 Resource Planning and Costs

This solution deploys the resources listed in the following table. The costs are only estimates and may differ from the final prices. For details, see **Pricing Details**.

Table 2-1 Resource planning and costs (pay-per-use)

Huawei Cloud Service	Example Configuration	Quantity	Estimated Monthly Cost
Huawei Cloud Flexus X Instance	• Pay-per-use: $0.30 USD/hour • Region: CN-Hong Kong • Specifications: FlexusX | Performance mode (disabled) | x1.8u.16g | 8 vCPUs | 16 GB • Image: Ubuntu 22.04 server 64bit • System disk: high I/O | 100 GB	1	$216.38 USD
Huawei Cloud Flexus X Instance	• Pay-per-use: $0.08 USD/hour • Region: CN-Hong Kong • Specifications: FlexusX | Performance mode (disabled) | x1.2u.4g | 2 vCPUs | 4 GB • Image: Ubuntu 22.04 server 64bit • System disk: high I/O | 40 GB	1	$54.93 USD

Huawei Cloud Service	Example Configuration	Quantity	Estimated Monthly Cost
RDS for PostgreSQL	• Pay-per-use: $0.11 USD/hour • Region: CN-Hong Kong • Billing mode: pay-per-use • Specifications: rds.pg.n1.large.2 (2 vCPUs | 4 GB) • Storage: SSD | 40 GB • DB engine: version 12	1	$76.32 USD
Elastic IP (EIP)	• Region: CN-Hong Kong • Billing mode: pay-per-use $0.16 USD/GB/hour • Routing type: dynamic BGP • Billed by: traffic • Bandwidth: 300 Mbit/s	2	$0.32 USD/GB/hour
Total	-	-	**$347.63 USD + EIP fee**

Table 2-2 Resource planning and costs (yearly/monthly)

Huawei Cloud Service	Example Configuration	Quantity	Estimated Monthly Cost
Huawei Cloud Flexus X Instance	• Yearly/Monthly • Region: CN-Hong Kong • Specifications: FlexusX | Performance mode (disabled) | x1.8u.16g | 8 vCPUs | 16 GB • Image: Ubuntu 22.04 server 64bit • System disk: high I/O | 100 GB	1	$159.54 USD

Huawei Cloud Service	Example Configuration	Quantity	Estimated Monthly Cost
Huawei Cloud Flexus X Instance	• Yearly/Monthly • Region: CN-Hong Kong • Specifications: FlexusX | Performance mode (disabled) | x1.2u.4g | 2 vCPUs | 4 GB • Image: Ubuntu 22.04 server 64bit • System disk: high I/O | 40 GB	1	$40.72 USD
RDS for PostgreSQL	• Yearly/Monthly • Region: CN-Hong Kong • Billing mode: pay-per-use • Specifications: rds.pg.n1.large.2 (2 vCPUs | 4 GB) • Storage: SSD | 40 GB • DB engine: version 12	1	$50.56 USD
Elastic IP (EIP)	• Region: CN-Hong Kong • Billing mode: pay-per-use $0.16 USD/GB/hour • Routing type: dynamic BGP • Billed by: traffic • Bandwidth: 300 Mbit/s	2	$0.32 USD/GB/hour
Total	-	-	**$250.82 USD + EIP fee**

3 Implementation Procedure

3.1 Preparations

If you use the HUAWEI ID created when you first used Huawei Cloud, skip these preparations. If you use an IAM user account, check if you are in the **admin** user group. If not, **grant relevant permissions** to your account and complete the following preparations.

(Optional) Creating the rf_admin_trust Agency

Step 1 Visit the Huawei Cloud official website, log in to the **console**, hover over the account name, and choose **Identity and Access Management**.

Figure 3-1 Huawei Cloud console

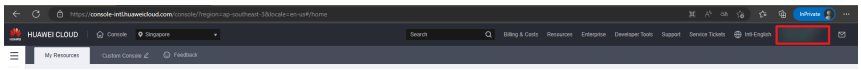

Figure 3-2 Identity and Access Management

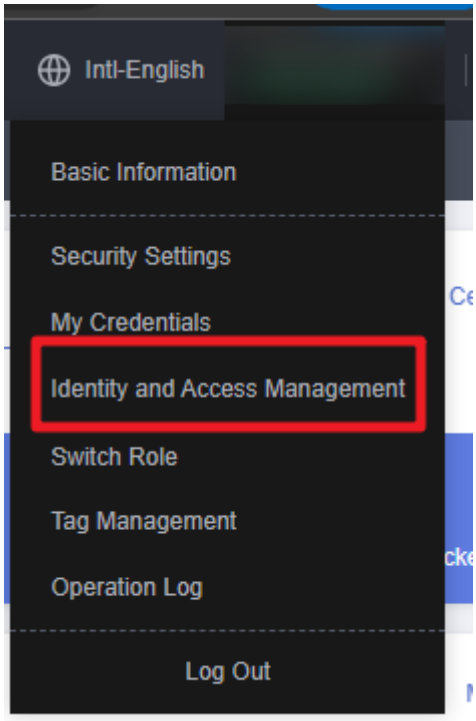

Step 2 In the navigation pane on the left, choose **Agencies**. On the displayed page, search for the **rf_admin_trust** agency.

Figure 3-3 Agencies

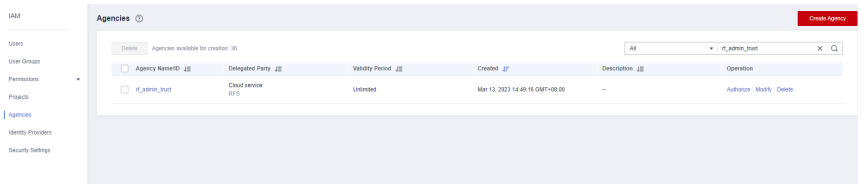

- If the agency is found, skip the following steps.
- If the agency is not found, perform the following steps.

Step 3 Click **Create Agency** in the upper right corner of the page. On the displayed page, set **Agency Name** to **rf_admin_trust**, **Agency Type** to **Cloud service**, and **Cloud Service** to **RFS**, and click **Next**.

Figure 3-4 Creating an agency

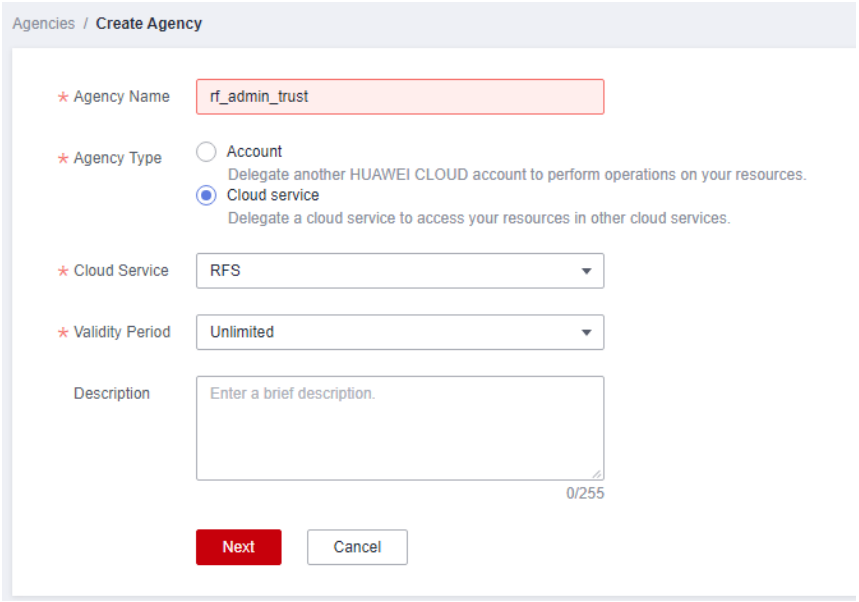

Step 4 Click **Authorize**.

Figure 3-5 Authorizing an agency

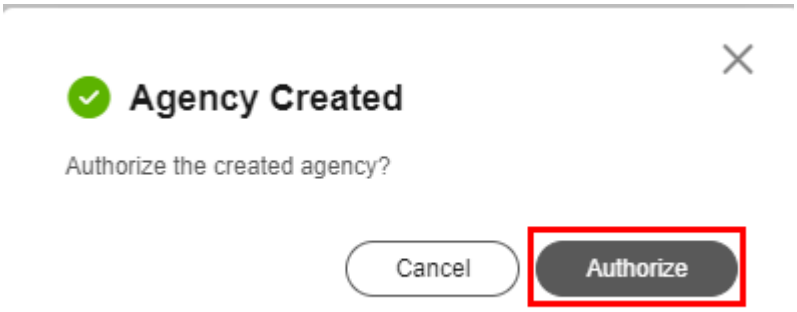

Step 5 Search for **Tenant Administrator**, select it in the search results, and click **Next**.

Figure 3-6 Selecting a policy

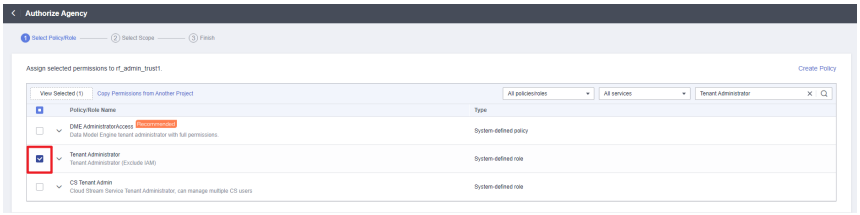

Step 6 Select **All resources** for **Scope** and click **OK**.

Figure 3-7 Selecting a scope

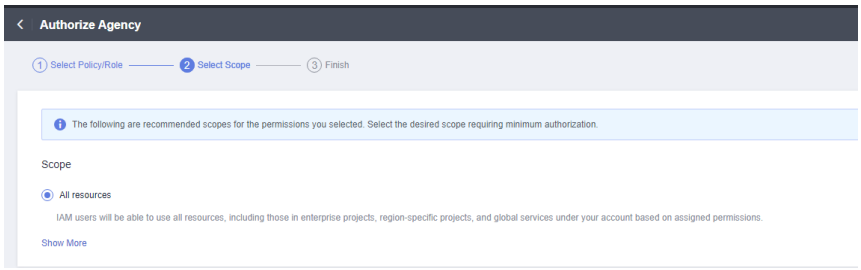

Step 7 Check that the **rf_admin_trust** agency is displayed in the agency list.

Figure 3-8 Agencies

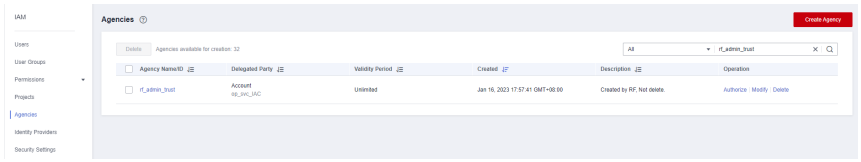

----End

3.2 Rapid Deployment

This section describes how to efficiently deploy the Getting Started with Intelligent Data Insight solution. You can complete quick deployment by following the steps and instructions in this section.

Table 3-1 Parameter descriptions (single-cloud server deployment)

Parameter	Type	Mandatory	Description	Default Value
vpc_name	String	Mandatory	The Virtual Private Cloud (VPC) name. This template uses a newly created VPC. The name must be unique. It can contain up to 57 characters. Only letters, numbers, underscores (_), hyphens (-), and periods (.) are allowed.	intelligent-data-insight_demo

Parameter	Type	Mandatory	Description	Default Value
secgroup_name	String	Mandatory	The security group name. This template uses a newly created security group. For details, see Modifying Security Group Rules. It can contain up to 64 characters. Only letters, numbers, underscores (_), hyphens (-), and periods (.) are allowed.	intelligent-data-insight_demo
rds_name	String	Mandatory	The RDS instance name. It must be unique and contain 4 to 64 characters. It must start with a letter and only letters (case-sensitive), numbers, hyphens (-), and underscores (_) are allowed.	intelligent-data-insight_demo
db_type	String	Mandatory	The RDS engine. Options: MySQL and PostgreSQL.	PostgreSQL

Parameter	Type	Mandatory	Description	Default Value
rds_flavor	String	Mandatory	The flavor of RDS instances. Single-node RDS instances are created by default. Default flavor of RDS for MySQL instances: rds.mysql.n1.large.2 (2U4G); default flavor of RDS for PostgreSQL instances: rds.pg.n1.large.2 (2U4G). For details about other flavors, see: **RDS for MySQL Instance Classes** **RDS for PostgreSQL Instance Classes**	rds.pg.n1.large.2

Parameter	Type	Mandatory	Description	Default Value
rds_port	number	Mandatory	The RDS port. RDS for MySQL instances can use database ports 1024 to 65535, excluding 12017 and 33071, which are reserved for RDS system use. Default value: **3306**. RDS for PostgreSQL instances can use database ports 2100 to 9500. Default value: **5432**.	5432
rds_volume_size	number	Mandatory	The storage space of an RDS instance. The default storage disk type is SSD. The value ranges from 40 to 4000 and must be a multiple of 10.	40

Parameter	Type	Mandatory	Description	Default Value
rds_password	String	Mandatory	The initial password used for logging in to the RDS database. It must contain 8 to 32 characters and only letters, numbers, and the following special characters are allowed: ~!@#%^*-_=+?	Empty
db_account_name	String	Mandatory	A user-defined RDS read-only username. It must be unique and contain 4 to 64 characters. It must start with a letter and only letters (case-sensitive), numbers, hyphens (-), and underscores (_) are allowed.	readonly_user

Parameter	Type	Mandatory	Description	Default Value
db_account_password	String	Mandatory	The login password for a user-defined RDS read-only username. It must contain 8 to 32 characters and at least three of the following types of characters: uppercase letters, lowercase letters, numbers, and the following special characters: ~!@#%^*-_=+?,	Empty

Parameter	Type	Mandatory	Description	Default Value
db_name	String	Mandatory	The database name. It can contain up to 63 characters. Only letters, numbers, and underscores (_) are allowed. It cannot start with **pg** or a number and must be different from RDS for PostgreSQL template database names. RDS for PostgreSQL template databases include **postgres**, **template0**, and **template1**.	business_db
dify_version	String	Mandatory	The Dify community edition, supporting v1.1.3, v0.15.3, and v0.15.2.	0.15.3

Parameter	Type	Mandatory	Description	Default Value
dify_ecs_name	String	Mandatory	The name of a Dify cloud server. It must be unique and contain up to 54 characters. Only letters, numbers, underscores (_), hyphens (-), and periods (.) are allowed.	intelligent-data-insight_dify_ecs_demo
dify_ecs_flavor	String	Mandatory	The flavor of Elastic Cloud Server (ECS) or FlexusX. The flavor ID format of a FlexusX is **x1.?u.?g**. For example, the flavor ID of a FlexusX with 2 vCPUs and 4 GiB memory is x1.2u.4g. For details about FlexusX flavors, refer to the console. For details about ECS flavors, see **A Summary List of x86 ECS Specifications**.	x1.8u.16g

Parameter	Type	Mandatory	Description	Default Value
dify_password	String	Mandatory	The password of the Dify cloud server. The password must contain 8 to 26 characters and include at least three of the following types of characters: uppercase letters, lowercase letters, numbers, and special characters (!@$%^-_=+[{}]:,./?). For details about how to change the password, see **Resetting the Password for Logging In to an ECS on the Management Console**. Default administrator account: **root**.	Empty
dify_system_disk_size	number	Mandatory	The system disk size of a Dify cloud server, in GB. The default disk type is high I/O. Value range: 40–1024. The disk cannot be shrunk.	100

Parameter	Type	Mandatory	Description	Default Value
dify_eip_size	number	Mandatory	The EIP bandwidth of a Dify cloud server, in Mbit/s. EIPs are billed by traffic. Value range: 1–300.	300
db_query_ecs_name	String	Mandatory	The name of the database interface cloud server. It must be unique and contain up to 54 characters. Only letters, numbers, underscores (_), hyphens (-), and periods (.) are allowed.	intelligent-data-insight_db_query_demo

Parameter	Type	Mandatory	Description	Default Value
db_query_ecs_flavor_id	String	Mandatory	The flavor of ECS or FlexusX. The flavor ID format of a FlexusX is **x1.?u.?g**. For example, the flavor ID of a FlexusX with 2 vCPUs and 4 GiB memory is x1.2u.4g. For details about FlexusX flavors, refer to the console. For details about ECS flavors, see **A Summary List of x86 ECS Specifications**.	x1.2u.4g

Parameter	Type	Mandatory	Description	Default Value
db_query_password	number	Mandatory	The password of the database interface cloud server. The password must contain 8 to 26 characters and at least the following types of characters: letters, numbers, and special characters. The following special characters are allowed: !@$ %^-_=+ [{}]:,./?. For details about how to change the password, see **Resetting the Password for Logging In to an ECS on the Management Console**. Default administrator account: **root**.	Empty
db_query_eip_size	number	Mandatory	The EIP bandwidth of the database interface cloud server, in Mbit/s. EIPs are billed by traffic. Value range: 1–300.	300

Parameter	Type	Mandatory	Description	Default Value
charging_mode	String	Mandatory	The billing mode. Options: **postPaid** (pay-per-use) and **prePaid** (yearly/monthly). Default value: **postPaid**.	postPaid
charge_period_unit	String	Mandatory	The unit of the subscription term. This parameter is only mandatory when **charging_mode** is set to **prePaid** (yearly/monthly). Options: **month** and **year**.	month
charge_period	number	Mandatory	The subscription term. This parameter is only mandatory when **charging_mode** is set to **prePaid** (yearly/monthly). Value range: 1–3 (**charging_unit** set to **year**); 1–9 (**charging_unit** set to **month**).	1

Step 1 Log in to **Huawei Cloud Solution Best Practices**, choose **Getting Started with Intelligent Data Insight**, and click **Deploy**. The **Create Stack** page is displayed.

Figure 3-9 Selecting a solution

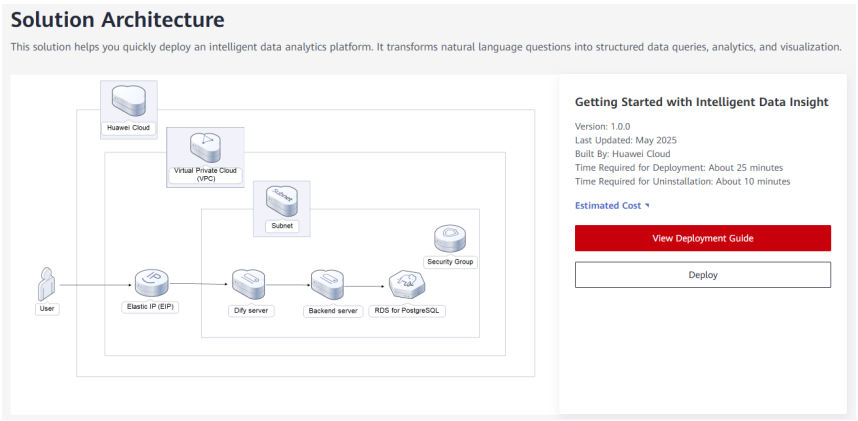

Step 2 On the **Select Template** page, click **Next**.

Figure 3-10 Selecting a template

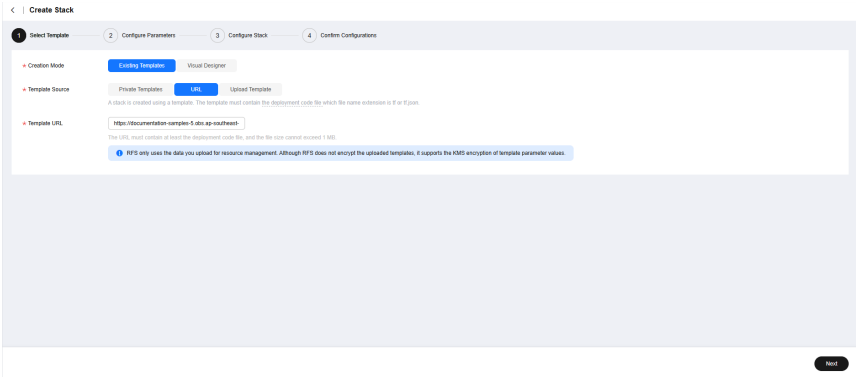

Step 3 On the **Configure Parameters** page, set parameters according to **Table 3-1** and click **Next**.

Figure 3-11 Setting parameters

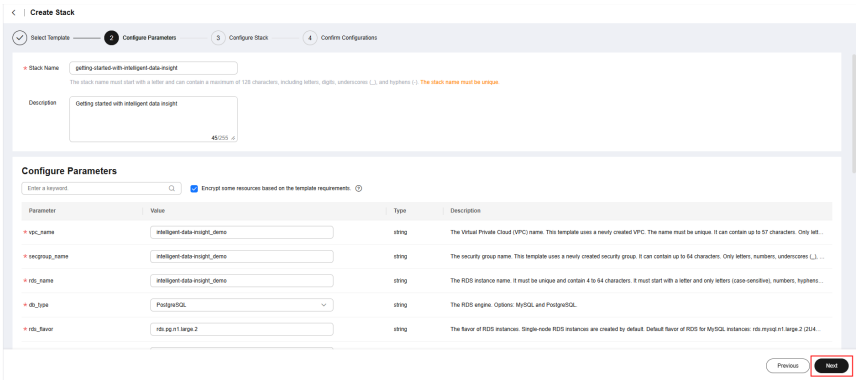

Step 4 On the **Configure Stack** page, select **rf_admin_trust** from the **Agency** drop-down list and click **Next**. This step is optional if you use a master account (HUAWEI ID) or use an IAM user in the **admin** user group.

Figure 3-12 Selecting an agency

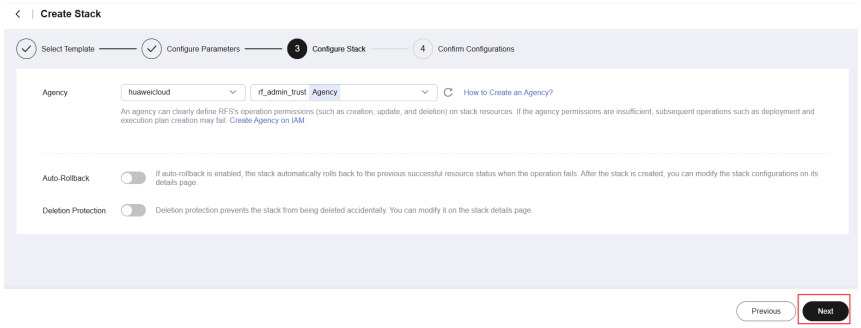

Step 5 On the **Confirm Configurations** page, click **Create Execution Plan**.

Figure 3-13 Confirming configurations

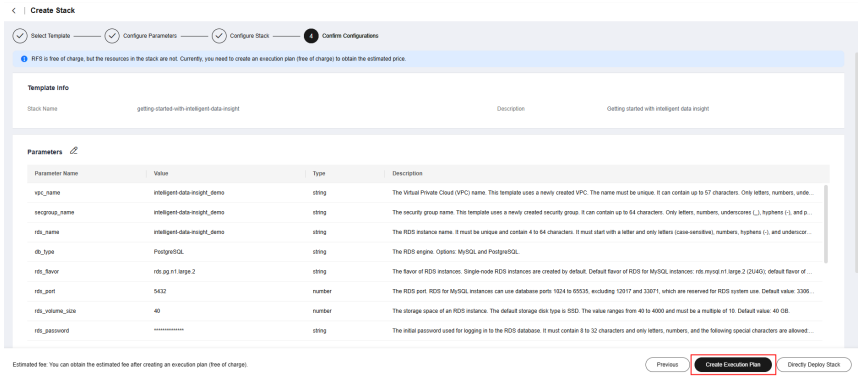

Step 6 In the **Create Execution Plan** dialog box that appears, enter an execution plan name and click **OK**.

Figure 3-14 Creating an execution plan

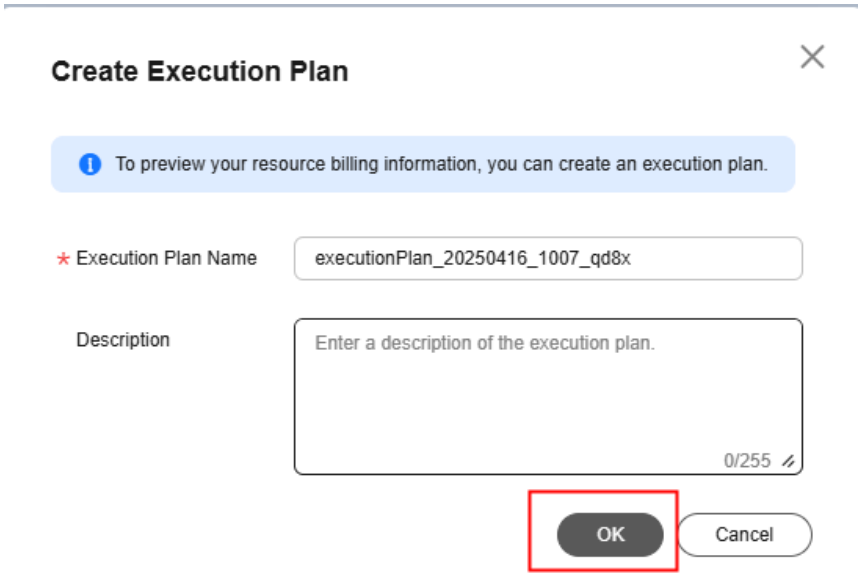

Step 7 Click **Deploy** in the **Operation** column. In the displayed dialog box, click **Execute**.

Figure 3-15 Execution plan created

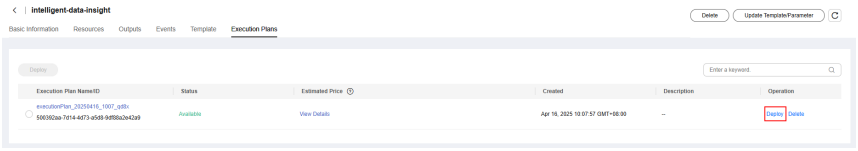

Figure 3-16 Deploying the execution plan

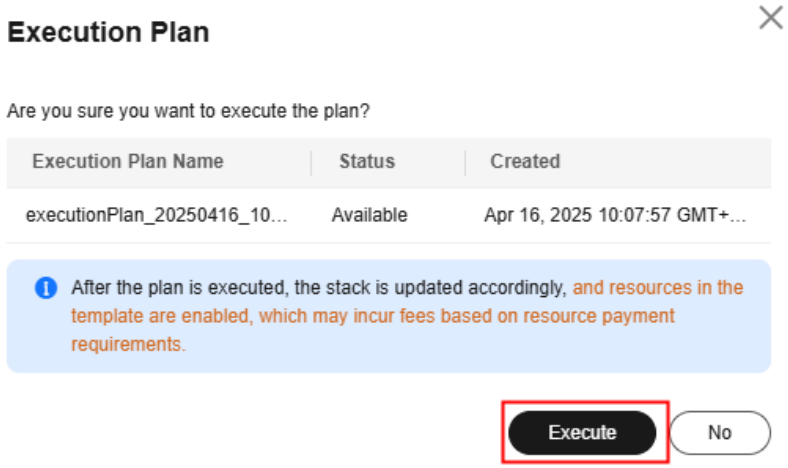

Step 8 (Optional) If you select the yearly/monthly billing mode and your account balance is insufficient, log in to the Billing Center to manually pay for the pending order. Refer to the corresponding table in **2 Resource Planning and Costs** for the total cost of one-click cloud service deployment under the yearly/monthly billing mode.

Step 9 Click the **Events** tab and check whether the message "Apply required resource success." is displayed in the **Description** column. If it is, the solution is successfully deployed.

Figure 3-17 Solution deployed

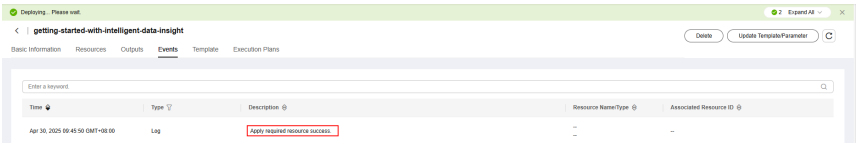

Step 10 Refresh the page, click the **Outputs** tab, and view the access links (Dify platform access link and backend database query link) and link for importing the Dify workflow DSL file in the **Value** column. After the stack is successfully deployed, wait approximately 5 to 10 minutes (affected by network fluctuations) before you can log in to Dify.

Figure 3-18 Descriptions

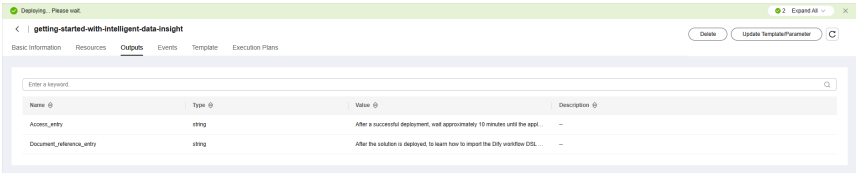

----End

3.3 Getting Started

(Optional) Modifying Security Group Rules

NOTICE

- This solution uses port 80 to access Dify. By default, all traffic is allowed. Refer to **Modifying Security Group Rules** to configure the IP address whitelist.
- This solution uses port 22 to remotely log in to the cloud server over SSH. If you need to remotely log in to the cloud server, configure an IP address whitelist by referring to **Modifying Security Group Rules** so you can access the service properly.
- After the solution is successfully deployed, it takes approximately 5 minutes to initialize the environment. The deployment duration varies depending on the network and bandwidth. You can access the environment only after the deployment is complete.

A security group is a collection of access control rules for traffic to and from cloud resources, such as cloud servers, containers, and databases. Cloud resources associated with the same security group have the same security requirements and are mutually trusted within a VPC.

You can modify security group rules, for example, by adding, modifying, or deleting a TCP port, as follows:

- Adding a security group rule: **Add an inbound rule** and enable a TCP port if needed.
- Modifying a security group rule: Inappropriate security group settings may introduce serious security risks. You can **modify security group rules** to ensure the network security of instances like ECSs.
- Deleting a security group rule: If the source or destination IP address of an inbound or outbound security group rule changes, or a port needs to be disabled, you can **delete the security group rule**.

Logging In to the Dify Platform

Step 1 Log in to the Dify platform: Enter the access address provided in **step 10 of the quick deployment process** to access Dify. Create an administrator account when you first log in by entering the email address, username, and password.

Figure 3-19 Creating an administrator account

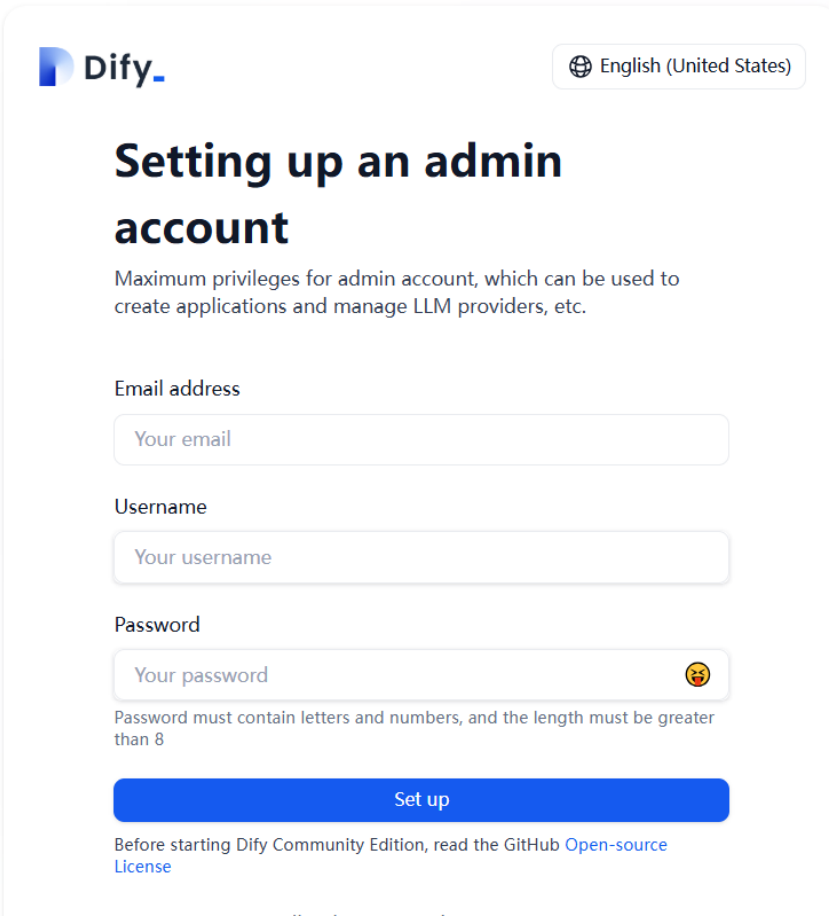

Step 2 Open your browser, access the Dify platform page, and enter the email address and password you configured in step 1 to log in to the Dify platform.

Figure 3-20 Logging in to the Dify platform

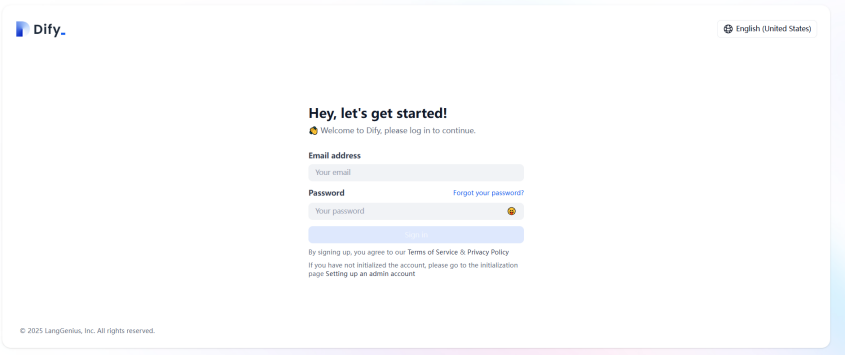

Figure 3-21 Dify platform

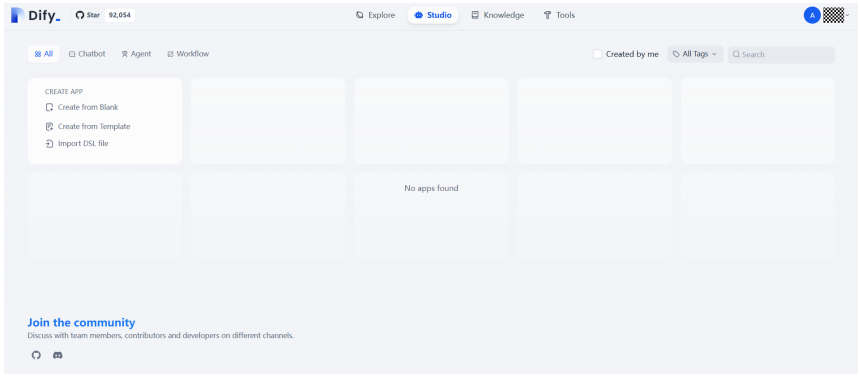

----End

Interconnecting with DeepSeek Models

Refer to "Interconnecting with DeepSeek Models" in **Getting Started**.

Creating an Intelligent Data Insight Workflow

Step 1 Log in to the Dify platform and click **Create from Blank** under **CREATE APP**. On the displayed page, click **Chatflow**, set **App Name & Icon**, and click **Create** in the lower part.

Figure 3-22 Clicking Create from Blank

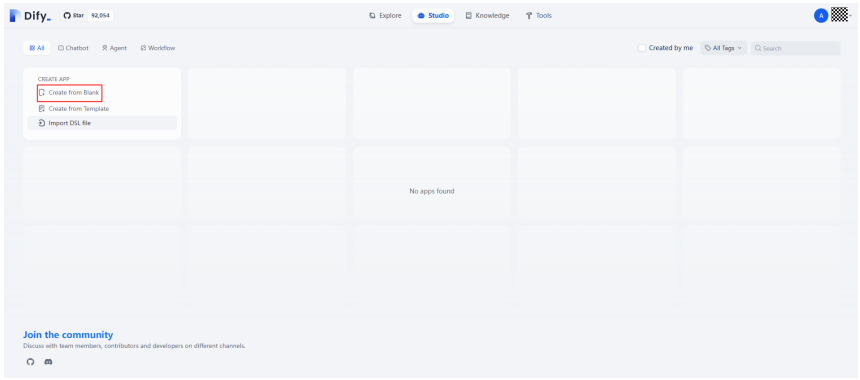

Figure 3-23 Creating a chat workflow

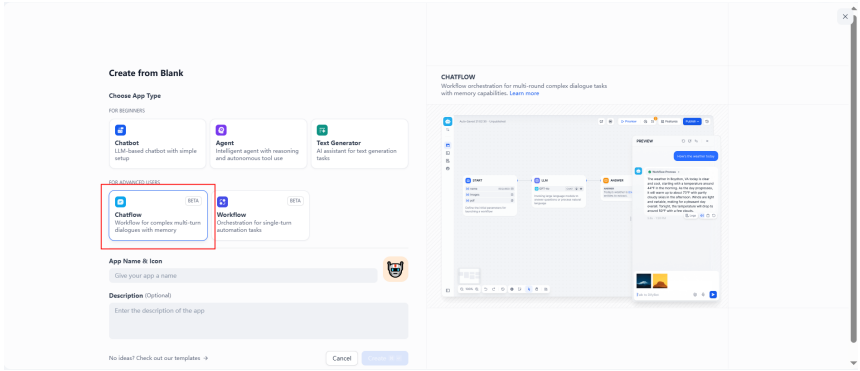

Step 2 Configure an LLM node: Select the connected model, set the context, and enter prompts in **SYSTEM**. You are advised to attach the database table structure and example SQL statements.

Figure 3-24 Configuring an LLM node

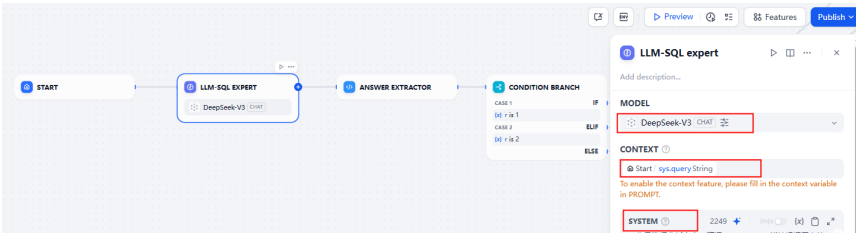

----End

Connecting Dify to a Database

Step 1 Create an HTTP request node and enter the API address exposed by the database server provided in **step 10 of the quick deployment process**.

Figure 3-25 Creating an HTTP request node

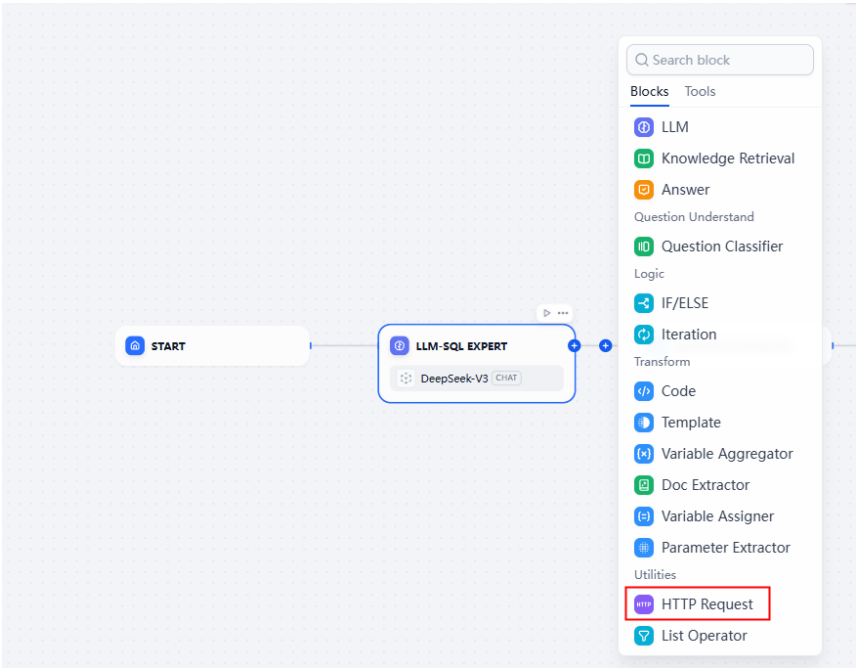

Step 2 Enter the API information of the database server: In **API**, select **POST** and enter the API address exposed by the database server; in **HEADERS**, set **Key** to **Content-Type** and **Value** to **application/json**; in **BODY**, select **JSON** and enter JSON information as shown in Figure 11.

Figure 3-26 Configuring the HTTP request node

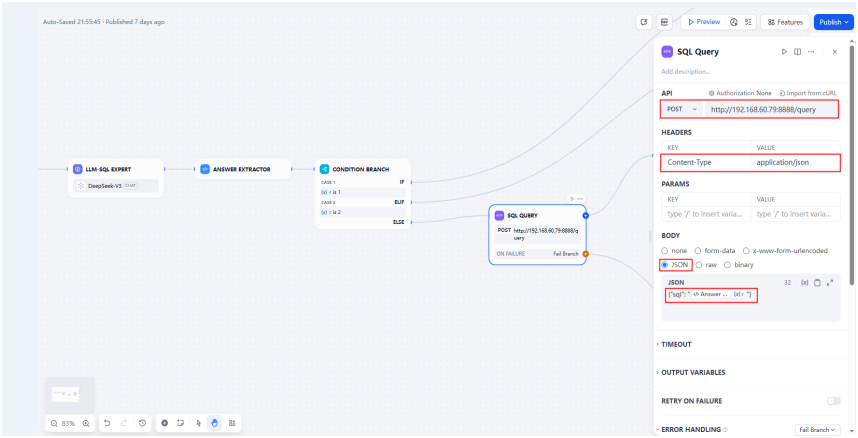

----End

Importing Data into the Database and Testing the Connectivity

Step 1 Log in to the **RDS console**. In the navigation pane on the left, choose **Instances**. On the displayed page, select the RDS for PostgreSQL instance created during one-click deployment and click **Log In** in its **Operation** column.

Figure 3-27 RDS console

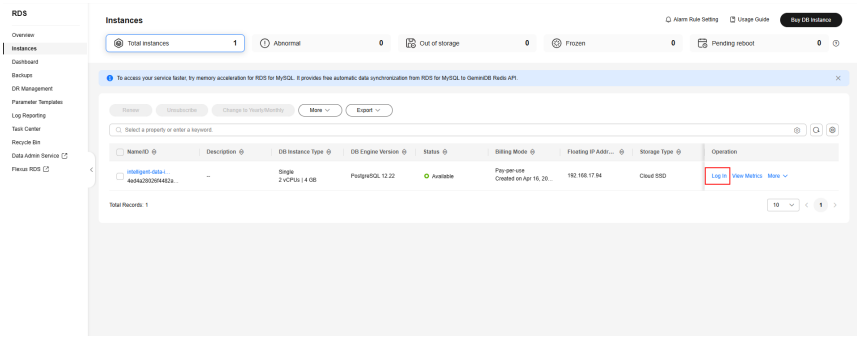

Figure 3-28 Logging in to the RDS for PostgreSQL instance as user root

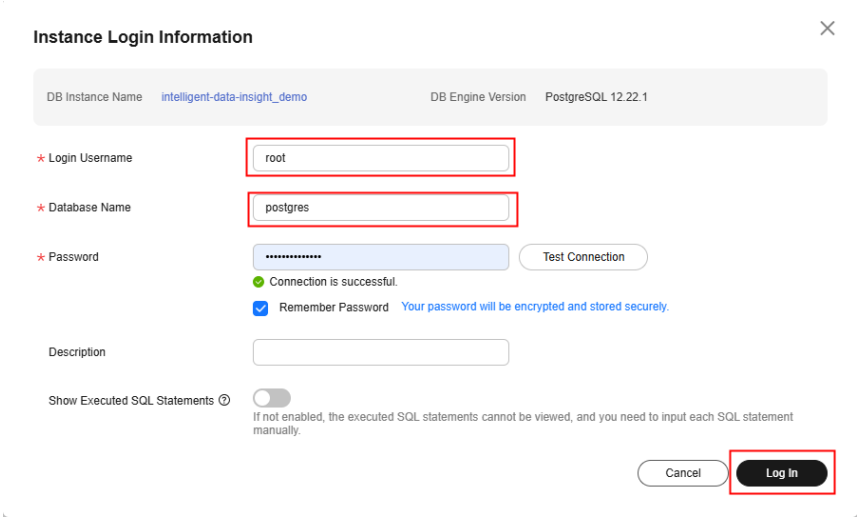

Step 2 If PostgreSQL is selected, grant the read-only permission to user **readonly_user**. If MySQL is selected, skip this step.

Figure 3-29 Setting the default transaction of user readonly_user to read-only

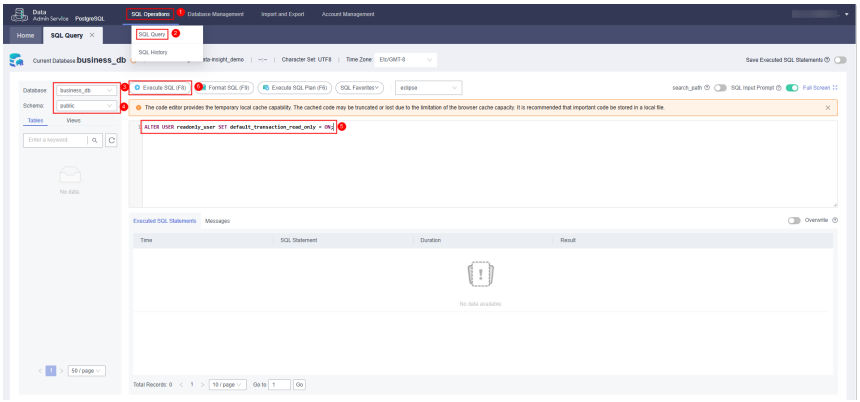

Figure 3-30 Checking whether root permissions have been optimized; if the following SQL statements do not report errors, optimization has been successful

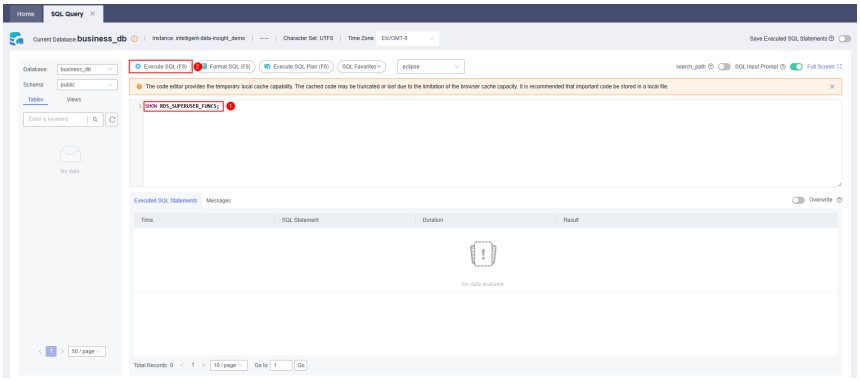

Figure 3-31 Performing this operation

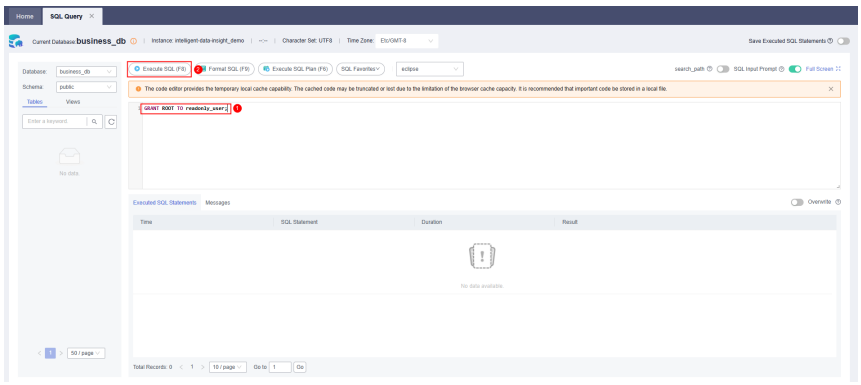

Step 3 Import the prepared SQL file into the preset database **business_db**. On the top menu, click **Import and Export** and select **Import**. On the displayed page, click **Create Task**. In the dialog box that appears, select an existing SQL file or a SQL file in an OBS bucket and click **Create**. (Note: You can also use DRS to transfer data to the database.)

Figure 3-32 Clicking Import and Export and selecting Import

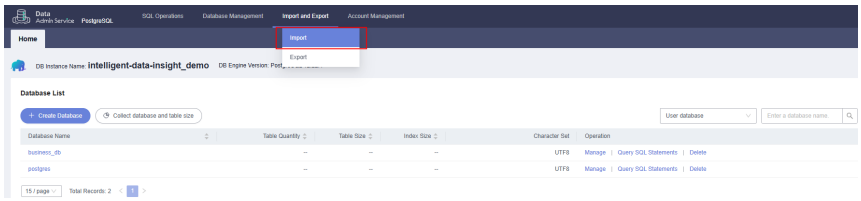

Figure 3-33 Importing the SQL file

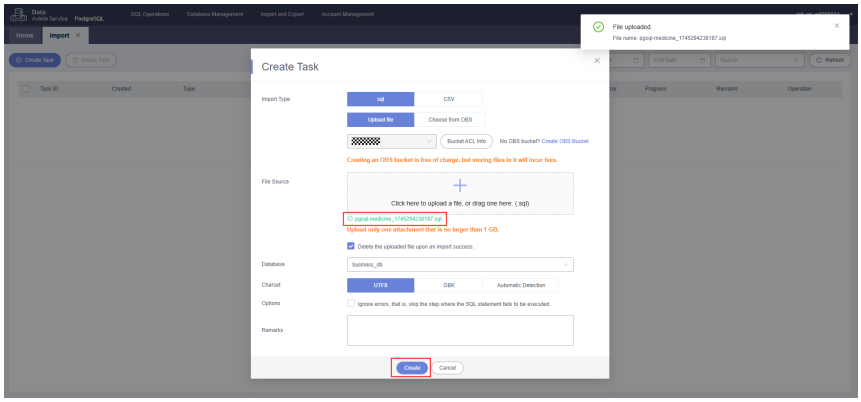

Figure 3-34 Data imported

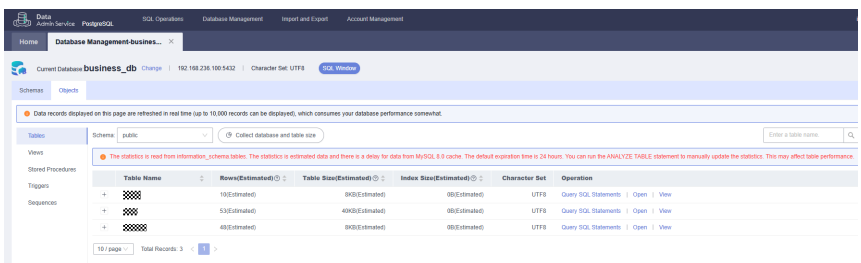

Step 4 Test the database connection on the HTTP request node of the Dify workflow. Click the button for testing the HTTP request node.

Figure 3-35 Clicking the button for testing the HTTP request node

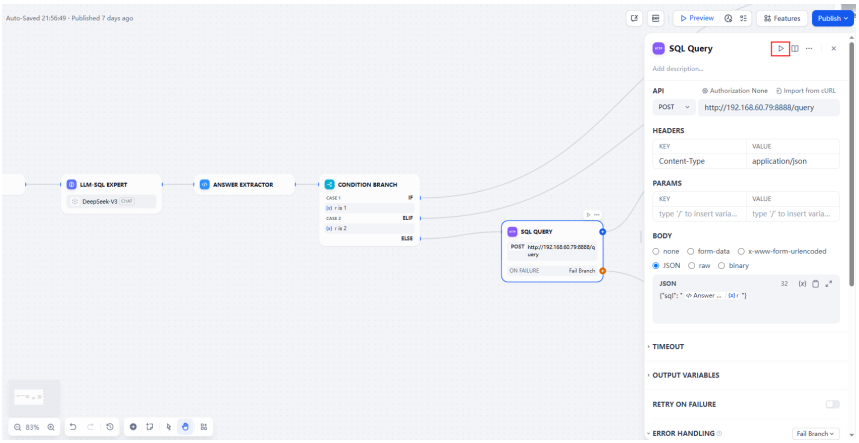

Figure 3-36 Entering and running test SQL statements

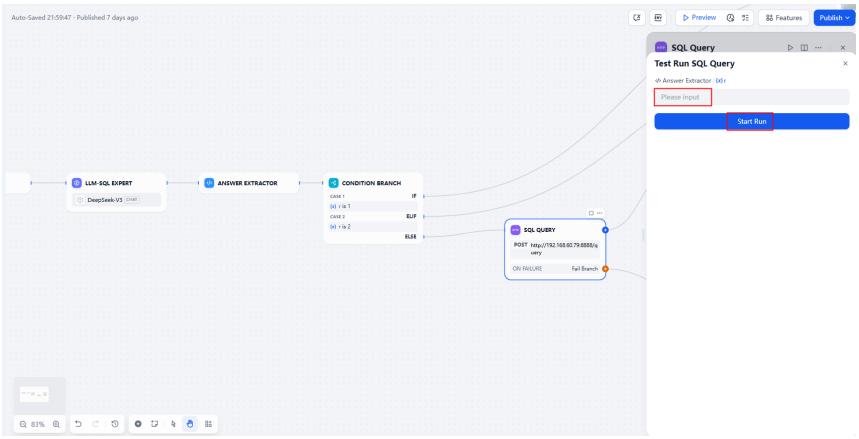

Step 5 Create a data analytics LLM node, configure the reply node content, and complete the setup of the entire intelligent data insight workflow.

Figure 3-37 Creating an LLM node

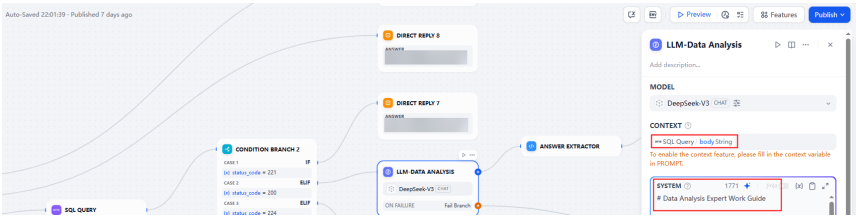

Figure 3-38 Configuring the reply node by selecting the output result of the data analytics LLM node

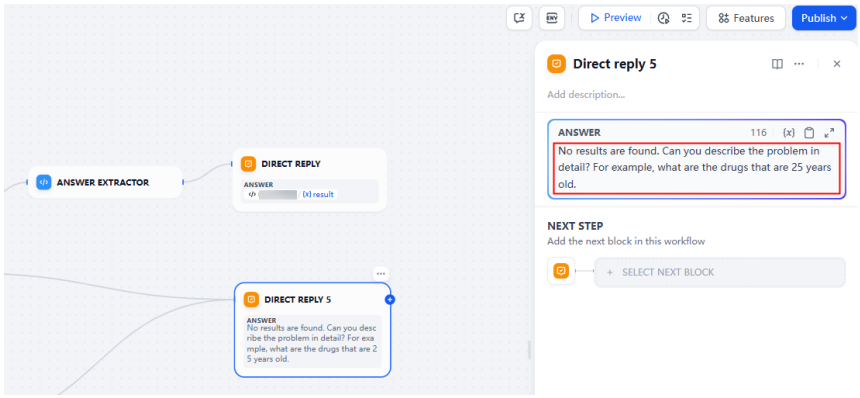

Step 6 Publish the workflow to launch the intelligent data insight dialog page.

Figure 3-39 Publishing a workflow

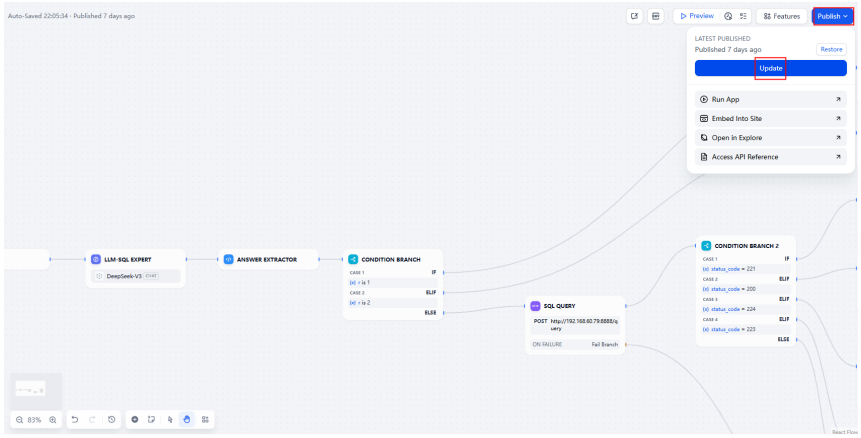

Figure 3-40 Gaining intelligent data insights

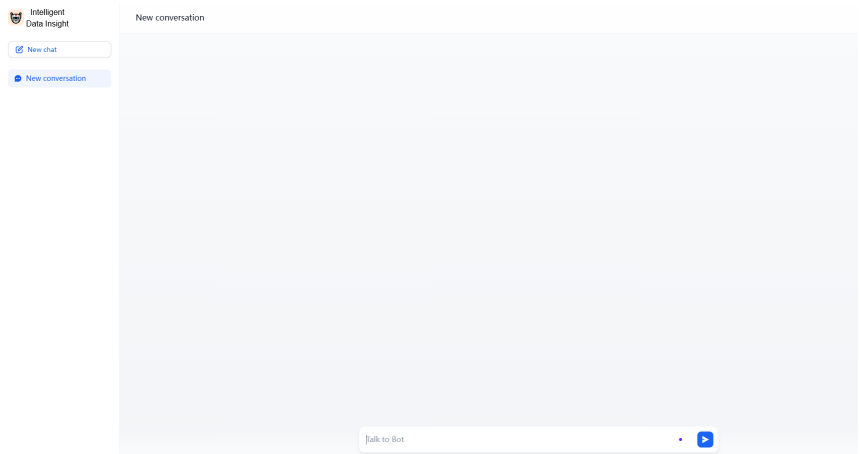

----End

3.4 Quick Uninstallation

Step 1 Before deleting a stack, ensure that the database access service has been stopped. If there are database connections, disconnect the sessions first before attempting to delete the stack. For details about how to terminate the database sessions, refer to steps 2 and 3.

Step 2 Log in to the **RDS console**. In the navigation pane on the left, choose **Instances**. On the displayed page, click the name of the database instance.

Step 3 In the navigation pane on the left, choose **DBA Assistant** > **Real-Time Diagnosis**. On the displayed page, all real-time sessions of the current database are displayed. Select all sessions with the database name **business_db** and the username **readonly_user**. Click **Kill Session**. In the **Kill Session** dialog box that appears, click **OK** to terminate the sessions.

Figure 3-41 Terminating sessions

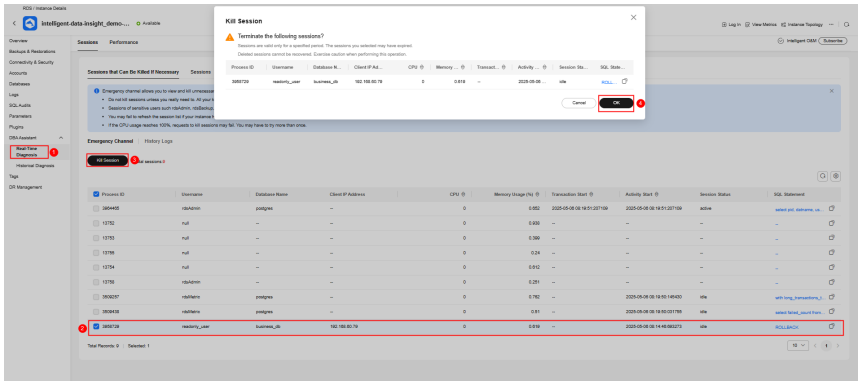

Step 4 Log in to the **Stacks page of the RFS console**, find the resource stack created for the solution, and click **Delete** in its **Operation** column.

Figure 3-42 Deleting the stack

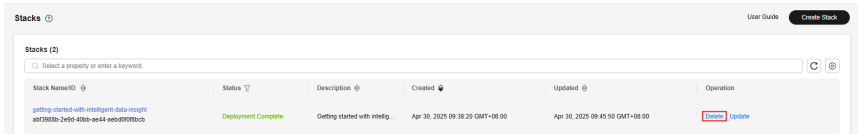

Step 5 In the **Delete Stack** dialog box that appears, set **When Deleted** to **Delete resource**, enter **Delete** in the box, and click **OK**.

Figure 3-43 Confirming the deletion

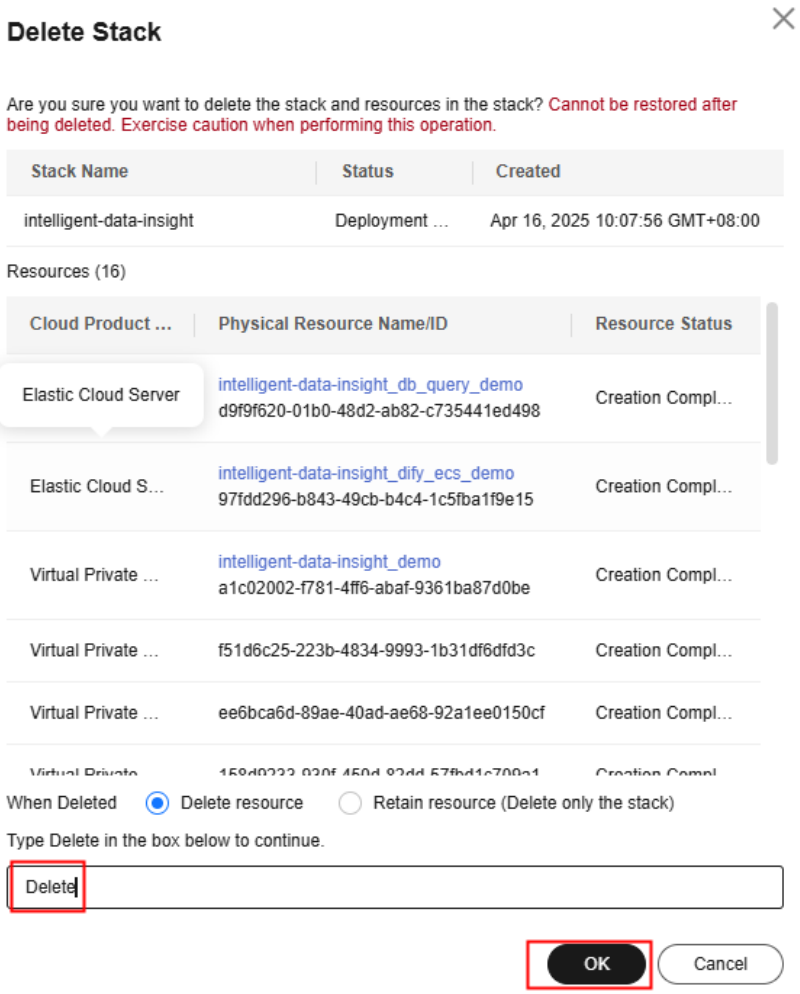

----End

4 Appendix

Terms

- Huawei Cloud Flexus X Instance: It is a next-gen flexible cloud server service tailored for small and medium-sized enterprises (SMEs) and developers. It offers similar functions to ECS but with added features like flexible vCPU memory ratios, live specification changes, and performance mode support.
- Virtual Private Cloud (VPC): An isolated and private virtual network environment that users can apply for on Huawei Cloud. VPC, along with cloud services like **Elastic Public IP (EIP)**, **Cloud Connect**, and **Direct Connect**, enables your cloud resources to securely communicate with each other over the Internet and on-premises networks.
- Elastic IP (EIP): enables your cloud resources to communicate with the Internet using static public IP addresses and scalable bandwidths. EIPs can be bound to or unbound from ECSs, Bare Metal Servers (BMSs), virtual IP addresses, load balancers, and NAT gateways, to access to or be accessed from the public network.

5 Change History

Table 5-1 Change history

Released On	Change Description
2025-05-14	This issue is the first official release.

