Table of Contents

What is NineData	1
NineData Overall Product Architecture	4
Replication System Functional Architecture	7
Dedicated Cluster Overview	9
Introduction to Data Replication	11
Migrating Real- time Incremental Data from Sybase to PostgreSQL	13
Migrating Full Data from Sybase to PostgreSQL	24
Bidirectional Replication Between Sybase and PostgreSQL	34
Full Data Migration from PostgreSQL to Sybase	47
Real-Time Incremental Data Migration from PostgreSQL to Sybase	55
Processing Data via ETL	65
Import Replication Object Configurations via Configuration Files	70
Introduction to database comparison	75
Data comparison	84
Schema comparison	90
Adding Sybase Data Source (Standalone Architecture)	95
Adding Sybase Data Source (Master-Slave Architecture)	98
Adding PostgreSQL Data Source (Primary-Secondary Architecture)	101
Adding PostgreSQL Data Source (Single- node Architecture)	104
Introduction to account management	107
Managing Organizations	111
Creating Users	115
Editing Users	117
Removing Users	119

Configuring User Permissions	120
Managing Roles	121
View audit logs	124
Configure Message Notifications	126



What is NineData

NineData is a cloud-native data management service that integrates multiple modules including database DevOps, data replication, data backup, and data comparison. It supports enterprise data management under various architectures such as on-premises data centers, hybrid clouds (combining on-premises databases with cloud databases), and multi-cloud (business architectures composed of databases from multiple cloud vendors), significantly reducing the operational difficulty and cost of enterprise data.

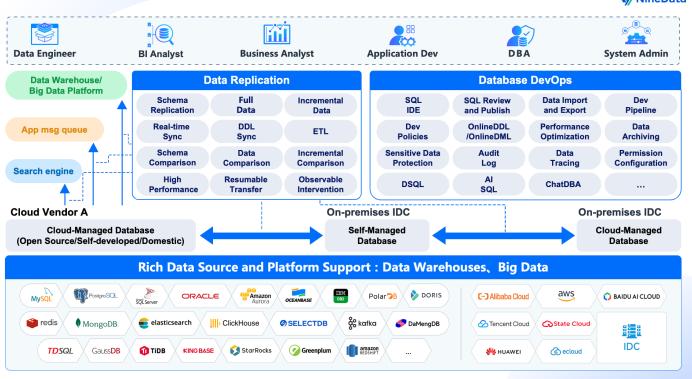
Background

With the deepening development of cloud computing, IDC self-built, hybrid cloud, and multi-cloud architectures have become the mainstream IT architectures for many enterprises. In such IT architectures, achieving unified and standardized management of various data sources under IDC self-built, hybrid cloud, and multi-cloud architectures has become a new challenge for enterprises.

NineData upgrades the concept of data management and introduces a set of multi-cloud, fully managed data management services. This service includes database DevOps, data replication, data backup, and data comparison functions. Through these functions, core application scenarios such as daily database development, secure data access, production database changes and releases, database backup and recovery, database localization migration, data migration, disaster recovery, active-active replication, data warehouse, and data lake construction can be easily achieved.

Functional Overview





Function	Description
Database DevOps	Cloud-native database DevOps tool for online query and management of data, supporting multiple database types, providing capabilities such as data querying, database DevOps, SQL scheduled tasks, and data export. A team collaboration mechanism that enables efficient and secure SQL development anytime, anywhere.
Backup and Recovery	Supports backup and recovery of databases in multiple environments and types. It has capabilities such as sub-second RPO, fine-grained data recovery, and online query of backup data.
Data Replication	Supports real-time and batch data replication between homogeneous and heterogeneous data sources. It realizes business scenarios such as data migration, remote disaster recovery, active-active database replication, and data integration for data warehouses and data lakes.
Data Comparison	Compares the structure definition and data consistency of two data sources. When inconsistencies occur, it can automatically generate data change SQL statements, which can be executed on the target end to repair the differing

Function	Description
	data, ensuring consistency between the two data sources. Supported database objects for data comparison include tables, stored procedures, functions, triggers, views, etc.

Advantages

- Seamless Integration Across Environments and Clouds: Supports seamless flow and
 management of data between multiple cloud vendors and self-built databases. Supports
 cloud platforms such as AWS, Alibaba Cloud, Huawei Cloud, and Tencent Cloud. Based on
 self-developed network technology, data sources in VPC networks can be securely accessed
 through private networks. Through NineData's exclusive clusters and data gateways, secure
 access and management of self-built data sources are realized.
- Rich Data Source Support: Supports unified management of homogeneous and
 heterogeneous data sources. As enterprises evolve, a single database architecture may not
 meet business requirements, and managing data from various types of databases typically
 requires many different tools, inadvertently raising the threshold and cost of operations.
 With the NineData platform, it is possible to easily manage different database types without
 the complexity of multiple tools.
- End-to-End Data Security: Data, as a core production element of an enterprise, holds the lifeline of the enterprise. NineData provides end-to-end data encryption capabilities around the entire lifecycle of data, effectively preventing unauthorized data queries and operations, and mitigating the business impact of data leaks.
- High Availability: All services adopt a high-availability architecture with multiple nodes. In
 the event of a failure on any node, requests or tasks can be quickly switched to normal
 nodes for execution. Furthermore, long-running tasks such as replication, backup, and
 comparison support seamless resumption from the point of interruption, so once an
 abnormality occurs in the running node, data transmission can continue after node
 switching.

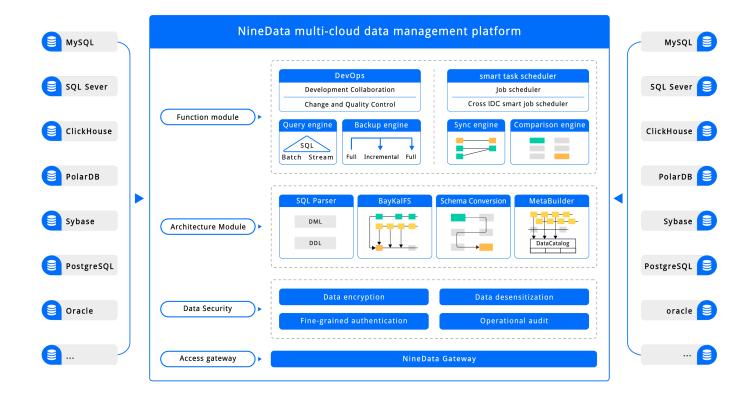
A

NineData Overall Product Architecture

This chapter introduces the overall product architecture and core functional components of NineData.

Architecture Overview

NineData provides a distributed, easily scalable, highly reliable, and multi-tenant product architecture. The product architecture diagram is as follows.



The entire product architecture is divided into five layers:

- Access Layer: The user interaction layer, providing tasks creation, management, and operation for data management. Mainly provided in the form of an interactive console.
- **Function Engine Layer**: Responsible for specific data management tasks. It includes DevOps, query engines, backup engines, synchronization engines, comparison engines, and scheduling engines.

- **Core Component Layer**: Common layer components that serve the upper-level function engines, providing capabilities such as log parsing, data transformation, and data storage engines.
- **Data Security Layer**: Provides data security protection capabilities covering the entire data lifecycle and throughout the data management process. Components include fine-grained access control, sensitive data management and masking, data encryption, and secure operation audit.
- **Data Source Connection Layer**: Responsible for providing access to various data sources for the upper-level engines.

Functional Components

Engine and Component	Description	
DevOps and Query Engine	 DevOps: Check whether SQL requests have execution permissions. Dynamically mask sensitive data. Provide enterprise-level database collaborative change capabilities. 	
	Query Engine: Receives SQL requests issued by users.	
Scheduling Engine	Responsible for overall platform task and resource scheduling. • Resource Scheduling: Dynamically adjust resource levels and online/offline service nodes based on cluster load pressure and node health conditions. • Task Scheduling: Tasks are scheduled to the most suitable nodes based on the geographical location of task data sources, network access quality, and resource load of each node. If a task encounters an exception, the scheduling engine automatically reschedules the task to healthy nodes to ensure stable execution.	

Backup Engine	Responsible for executing backup and recovery tasks and periodically pushing task execution status, progress, and exceptions.
Synchronization Engine	Responsible for executing data replication tasks, automatically scheduling related dependent tasks in a pipeline manner based on user-configured replication types. It also periodically pushes task execution status, progress, and exceptions.
Comparison Engine	Responsible for executing structure and data comparison tasks between two data sources and periodically pushing task execution status, progress, and exceptions.
SQL Parser	SQL parsing module. Responsible for parsing SQL statements and logs throughout the platform, assisting engines in tasks such as data reading and secure data access.
Sensitive Data Management	Responsible for providing sensitive data metadata and masking functions for query engines, backup engines, synchronization engines, and comparison engines. It includes built-in algorithms for detecting sensitive data, which automatically scans sensitive fields in data sources and tags them when enabled.

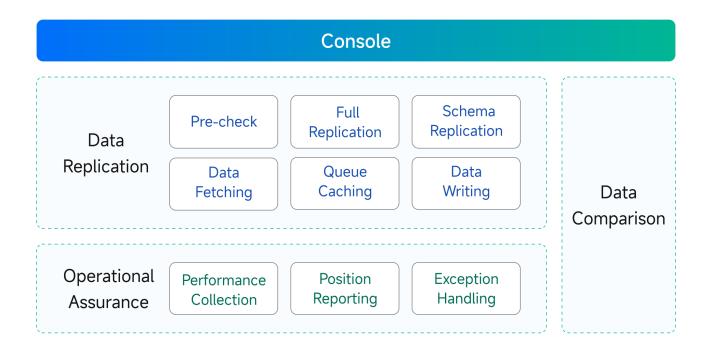
Ħ

Replication System Functional Architecture

This chapter introduces the product architecture of the NineData replication system.

Architecture Overview

The core capability of the NineData replication system is the ability to replicate data between two data sources. To achieve this core capability, the overall functional architecture is as shown in the following diagram.



The core modules of the overall data replication include:

- **Console**: The user interface where users create and manage replication tasks. The current console provides end-to-end configuration of data replication tasks. It also offers comprehensive task monitoring and operational capabilities.
- **Pre-check**: To ensure the successful initiation of replication/comparison tasks, certain prerequisites need to be met. To allow users to identify and resolve any issues beforehand,

- the replication platform provides pre-check capabilities to detect prerequisites before replication/comparison tasks.
- Schema Replication: Databases are divided into relational databases and NoSQL databases. In relational databases, data storage and management are accomplished through schema objects. To achieve scenarios such as data migration, disaster recovery, and activeactive databases, the replication system provides the ability to replicate schema objects.
- Full Data Replication: This module is mainly responsible for replicating historical full data. During the full data replication phase, the replication system migrates all existing data from the source database to the target database. To expedite the full replication access for new data sources, the replication system provides a plugin architecture and defines its own intermediate type. In the full replication framework, the replication system maps the source data source type to the custom intermediate data type, and then maps it from the intermediate data type to the target data type. This design allows for the reuse of existing data type conversion capabilities for supported data sources by the replication system. This N-by-M topology approach enables implementation through the combination of N and M.
- Data Capture, Data Writing, and Queue Buffering: These modules are responsible for
 real-time capture and replication of database change logs, enabling incremental data
 replication based on these capabilities. The data capture module needs to be compatible
 with different database log collection and listening protocols to achieve end-to-end data
 replication capabilities including real-time listening, collection, parsing, processing, and
 writing of logs.
- Stability Assurance: The other capabilities in the architecture diagram are provided to
 ensure the long-term stable operation of replication tasks. This includes core metric
 monitoring, anomaly diagnosis and handling mechanisms, data quality assurance systems,
 etc., providing observable and intervenable capabilities for replication tasks. In addition, the
 replication system provides a scheduling framework based on K8S, which can automatically
 connect various replication steps and provide fully automated task scheduling processes.

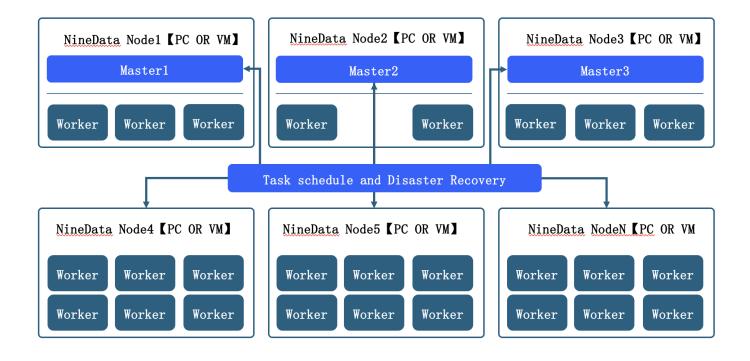
Dedicated Cluster Overview

In addition to providing SaaS services, NineData also offers a dedicated cluster feature for users with high security requirements.

Feature Introduction

The dedicated cluster mode of NineData is a service mode that privatizes NineData services. Compared to the SaaS mode of NineData services, the entire set of services will be established in a user-provided server cluster, ensuring data remains within the dedicated cluster. All data flow will occur within the dedicated cluster environment, suitable for users with high security and compliance requirements.

Functional Architecture



Advantages

Network Isolation: The network environment of NineData's dedicated cluster is independent
of other environments, including NineData's public cloud. Data sources within the dedicated
cluster cannot be directly accessed, but interaction with the servers in the dedicated cluster,

- such as task configuration information and status information, can be done through the NineData console without involving any user data.
- High Availability: When the number of hosts within the dedicated cluster is greater than 1, it supports high availability of tasks within the cluster. In the event of an abnormal node status, all tasks in that node are automatically transferred to healthy nodes for execution.
- Task Scheduling: When the number of hosts within the dedicated cluster is greater than 1, the system automatically assesses the status of each node and intelligently schedules tasks to nodes with lower utilization rates.
- Data Security: Supports high-security authentication mechanisms between clusters and data sources, as well as encrypted data transmission.
- Operations and Maintenance Capabilities: Provides functions such as host status monitoring and anomaly alerts. It also supports operations such as adding new hosts to existing clusters or removing hosts from clusters.

Introduction to Data Replication

NineData data replication supports offline and real-time replication between various homogeneous and heterogeneous data sources. It is suitable for various business scenarios such as data migration, database scaling, database version upgrades, disaster recovery, active-active deployments across different locations, data warehousing, and data lake integration.

Product Advantages

Compared to traditional data replication products, NineData's data replication product has significant advantages in terms of security, performance, cost, and other aspects.

- ETL Support: Data replication supports ETL (Extract, Transform, Load) functionality, enabling data transformation during replication from the source, processing it into the desired state, and then replicating the processed data to the destination.
- Easy to Use: Supporting a step-by-step guided task configuration, while automatically and smoothly initiating each replication stage based on the user-configured replication type. No learning curve, simple and user-friendly.
- High Performance: Combining real-time log capture, intelligent sharding, transactional concurrency, and hotspot merging technologies to achieve strong replication performance with incremental replication latency reduced to seconds.
- High Reliability: Comprehensive monitoring and alerting mechanisms, built-in task anomaly
 detection, self-repair mechanisms, and intervention capabilities enhance task availability.
 Mechanisms such as incremental replication, multi-version metadata, and seamless task
 restart in case of anomalies ensure task reliability.
- Strong Consistency: Innovative breakpoint resumption technology and transactional consistency synchronization capability effectively ensure data consistency in replication. Coupled with NineData's data comparison feature, it ensures data quality.
- Multi-Database: Provides data replication capabilities for various databases for enterprises, supporting replication between locally deployed and cloud vendor databases. It also supports various common databases, including MySQL, Sybase, SQLServer, PostgreSQL, Oracle, Redis, MongoDB, ClickHouse, Doris, SelectDB, Elasticsearch, Kafka, Redshift, Greenplum, SingleStore, StarRocks, and more.

Use Cases

- Database Migration: Required for cross-regional, cross-cloud, or cross-data-source data migration, as well as migrating on-premises databases to the cloud, and other scenarios.
- Real-time Data Warehouse Synchronization: Supports real-time integration of data from multiple sources and environments into a unified data warehouse for analysis.
- Cross-Cloud and Cross-Region Disaster Recovery: Supports continuous data synchronization between business centers and disaster recovery centers. In the event of a failure in the primary region, you can switch user requests to the disaster recovery region to achieve cross-cloud and cross-region disaster recovery.
- Active-Active Deployments Across Different Locations: Supports bidirectional real-time synchronization between business nodes in multiple regions, ensuring global data consistency.
 - When any unit fails, simply switch the traffic of that unit to other units to achieve second-level business recovery, effectively ensuring high service availability.
 - Distribute business traffic to various business units based on a certain dimension of the business. For example, divide traffic of each unit based on the user's region, achieving nearby access for users, reducing network latency, and improving user experience.
 Meanwhile, having business units distributed in different regions can effectively solve the problem of infrastructure limitations in a single region hindering business expansion.

Migrating Real-time Incremental Data from Sybase to PostgreSQL

Migrating Real-time Incremental Data from Sybase to PostgreSQL

NineData data replication supports replicating data from Sybase to PostgreSQL, supporting both full replication and incremental replication.

Prerequisites

- Source and destination data sources have been added to NineData.
- Tables must be manually created in the destination data source (PostgreSQL), and the table structure must match that of the source data source (Sybase) that needs to be replicated.
- Versions of source and destination data sources are as follows:

Source Data Source	Destination Data Source
Sybase 16.0, 15.7	PostgreSQL 16, 15, 14

You must have the following permissions on both the source and destination data sources:

Replication Type	Sybase	PostgreSQL
One-Way Full and	replication_role,	All Privileges on Schema, All
Incremental Replication	sso_role	Privileges on All Tables

If PostgreSQL is in a master-slave architecture, the virtual IP (VIP) of PostgreSQL must be
used when adding the data source to ensure that NineData accesses the master even after a
failover.

Precautions

- Evaluate the performance of the source and destination data sources before performing data synchronization. It is recommended to execute data synchronization during off-peak hours. Otherwise, full data initialization will consume certain read and write resources of both the source and destination data sources, leading to increased database load.
- Ensure that each table in the synchronized objects has a primary key or unique constraint, and column names are unique to avoid duplicating the same data during synchronization.

Operation Steps

- 1. Log in to the NineData console.
- 2. Click on **Replication** in the left sidebar navigation.
- 3. On the **Replication** page, click on the **Create Replication** button in the upper right corner.
- 4. On the **Source & Target** tab, configure according to the table below, and then click **Next**.

Parameter	Description
Name	Enter a meaningful name for the data synchronization task to facilitate later retrieval and management. Supports up to 64 characters.
Source	The Sybase data source where the synchronization objects are located.
Target	The PostgreSQL data source that receives the synchronized objects.
Replication Mode	Select Uni-directional.
Туре	Check Full and Incremental .
If target table already exists	 Choose according to your needs. Pre-Check Error and Stop Task: Stops the task if data is detected in the target table during the precheck phase.

Parameter	Description
	 Ignore the existing data and append: Ignores the data detected in the target table during the precheck phase and appends other data.
	 Clear the existing data before write: Deletes the data detected in the target table during the precheck phase and rewrites it.

5. On the **Objects** tab, configure the following parameters, and then click **Next**.

Parameter	Description
Replication Objects	Select the content to be copied. You can choose All Objects to copy all contents of the source database, or choose Customized Object , select the content to be copied in the Source Object list, and click > to add it to the Target Object list on the right.

6. (Optional Step) The data replication feature provides configuration file templates for batch import. Click Import Config at the top right corner of the Objects tab, and select Download Template to download the configuration file template to your local machine. After editing, click Upload to upload the configuration file for batch import.

Parameter	Description
source_table_name	Name of the source table containing the objects to be synchronized.
destination_table_name	Name of the target table receiving the synchronized objects.
source_schema_name	Name of the source schema containing the objects to be synchronized.
destination_schema_name	Name of the target schema receiving the synchronized

Parameter	Description
	objects.
source_database_name	Name of the source database containing the objects to be synchronized.
target_database_name	Name of the target database receiving the synchronized objects.
column_list	List of columns to be synchronized.
extra_configuration	Additional configuration information, where you can configure the following: Column Mapping: column_name, destination_column_name Column Value: column_value Data Filtering: filter_condition

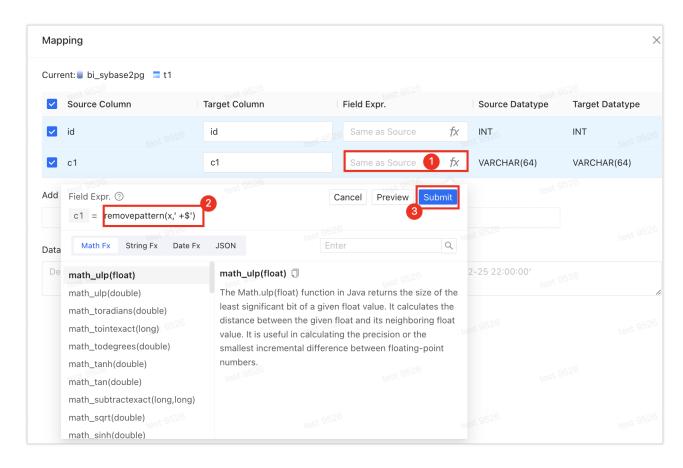


• An example of the extra_configuration content is as follows:

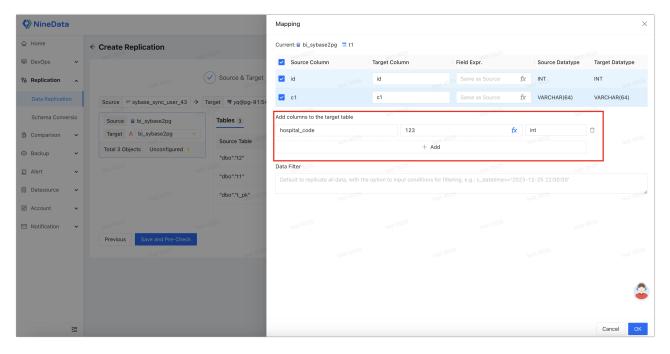
```
"column_name": "created_time", // Original column name to
be mapped
   "destination_column_name": "migrated_time", // Target
column name mapped to "migrated_time"
   "column_value": "current_timestamp()", // Change the
column value to the current timestamp
   "filter_condition": "id != 0" // Only synchronize rows
where ID is not 0.
}
```

• For a complete example of the configuration file, refer to the downloaded template.

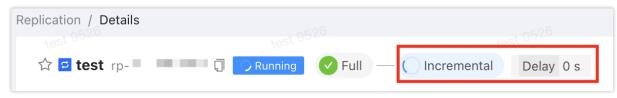
- 7. In the **Mapping** tab, the system defaults to selecting the same name database from the target data source. If it does not exist, you need to manually select the target database, then perform the following sub-steps as needed, and finally click **Save and Pre-Check**.
 - Using Field Expr.: To remove trailing spaces from a field name, click on Mapping & Filtering on the right side of the target table, find the field you want to adjust, click the blank box under the Field Expr. column, enter removepattern(x,' +\$') in the Enter text box, click Submit, and then click OK.



Adding a new column in the target table: To add a column named hospital_code with
a constant value, click on Mapping & Filtering on the right side of the target table, then
click Add columns to the target table under Add, fill in the column name, column
expression, and column type sequentially, then click OK.



- Refresh Metadata: If there are updates in the source and target data sources during the mapping configuration stage, click the Refresh Metadata button at the top right corner of the page to re-fetch the information from the source and target data sources.
- Setting Data Filter: Click on Mapping & Filtering on the right side of the page, and in the text box under Data Filter, enter emp_no>=10005, so data with emp_no less than 10005 will not be synchronized to the target data source.
- 8. On the **Pre-check** tab, wait for the system to complete the precheck. After the precheck passes, click **Launch**.
 - You can check Enable data consistency comparison. After the synchronization task is completed, the data consistency comparison based on the source data source is automatically started to ensure the consistency of data on both ends. Depending on the Type you choose, Enable data consistency comparison will start as follows:
 - Full: Starts after full replication is completed.
 - Full+Incremental, Incremental: Starts when the incremental data is first
 consistent with the source data source and Delay is 0 seconds. You can click View
 Details to view the synchronization delay on the Details page.



• If the precheck fails, you need to click **Details** in the **Actions** column on the right side of the target check item to troubleshoot the cause of the failure, manually fix it, and then

- click Check Again to perform the precheck again until it passes.
- For check items with **Result** as **Warning**, you can repair or ignore them according to the specific situation.
- 9. On the **Launch** page, when prompted **Launch Successfully**, the synchronization task starts running. At this point, you can perform the following operations:
 - Click View Details to view the execution status of each stage of the synchronization task.
 - Click **Back to list** to return to the **Replication** task list page.

View Synchronization Results

- 1. Log in to the NineData console.
- 2. Click on **Replication** > **Data Replication** in the left navigation bar.
- 3. On the **Replication** page, click on the **Task ID** of the target synchronization task. The page explanation is as follows:



Index	Function	Description
1	Sync Delay	The synchronization delay between the source data source and the target data source. 0 seconds indicates no delay between the two ends. At this point, you can choose to switch the business to the target data source for smooth migration.
2	Configure Alerts	After configuring alerts, the system will notify you through your chosen method in case of task failure.

Index	Function	Description
3	More	 Pause: Pause the task. Only tasks with the status Running can be paused. Duplicate: Create a new replication task with the same configuration as the current task. Terminate: End tasks that are incomplete or in listening mode (i.e., incremental synchronization). Once the task is terminated, it cannot be restarted. Please proceed with caution. If the synchronization objects include triggers, the trigger replication option will pop up. Please select as needed. Delete: Delete the task. Once the task is deleted, it cannot be recovered. Please proceed with caution.
4	Full Replication (Displayed in scenarios with full replication)	 Display the progress and detailed information of full replication. Click on Monitor on the right side of the page to view various monitoring indicators during full replication. Click on Log on the right side of the page to view the execution log of full replication. Click on C on the right side of the page to view the latest information.
5	Incremental Replication (Displayed in scenarios with incremental replication)	 Display various monitoring indicators for incremental replication. Click on Log on the right side of the page to view the execution log of incremental replication. Click on C on the right side of the page to view the latest information.

Index	Function	Description
6	Modify Object	 Display modification records of synchronized objects. Click on Modify Objects on the right side of the page to configure synchronized objects. Click on C on the right side of the page to view the latest information.
7	Data Comparison	Display the comparison results between the source data source and the target data source. If you haven't enabled data comparison, click on Enable Comparison on the page. Click on Re-compare on the right side of the page to reinitiate the comparison between the current source and target endpoints. Click on Stop on the right side of the page to immediately stop the comparison task once it has started. Click on Log on the right side of the page to view the execution log of consistency comparison. Click on Monitor (Displayed only in data comparison) on the right side of the page to view the trend graph of RPS (records per second) comparison. Click on Details to view earlier records. Click on in the Actions column on the right side of the comparison list (Displayed under the Data tab only in case of inconsistencies) to view the comparison details between the source and target endpoints. Click on in the Actions column on the right side of the comparison list (Displayed in case of inconsistencies) to generate change SQL. You

Index	Function	Description
		can directly copy this SQL and execute it on the target data source to modify the inconsistent content.
8	View Reverse	Displayed only in bidirectional replication tasks, click to view the replication details from the target data source to the source data source.
9	Expand	Display detailed information of the current replication task.

Appendix: Precheck Item List

Check Item	Check Content
Source object existence check	Check whether the table on the target database currently exists on the source database
Source datasource connection check	Check the status of the source datasource, database connectable, and verify the username and password
Target datasource connection check	Check the status of the target datasource, database connectable, and verify the username and password
Target databse privilege check	Check whether the account privileges of the target database meet the requirements
Source databse privilege check	Check whether the account privileges of the source database meet the requirements
Target database data existence check	Check whether the object to be replicated not empty in the target database

Check Item	Check Content
Objects with the same name in the target database existence check	Check whether the object to be replicated already exists in the target database
Cyclic replication check	Check for replication loops

Migrating Full Data from Sybase to PostgreSQL

Migrating Full Data from Sybase to PostgreSQL

NineData data replication supports replicating full data from Sybase to PostgreSQL. This chapter introduces the configuration method.

Prerequisites

- Source and destination data sources have been added to NineData.
- Tables must be manually created in the destination data source (PostgreSQL), and the table structure must match that of the source data source (Sybase) that needs to be replicated.
- Versions of source and destination data sources are as follows:

Source Data Source	Destination Data Source
Sybase 16.0, 15.7	PostgreSQL 16, 15, 14

You must have the following permissions on both the source and destination data sources:

Replication Type	Sybase	PostgreSQL
One-Way Full Replication	SELECT permission on migration objects	All Privileges on Schema, All Privileges on All Tables

• If PostgreSQL is in a master-slave architecture, the virtual IP (VIP) of PostgreSQL must be used when adding the data source to ensure that NineData accesses the master even after a failover.

Precautions

- Evaluate the performance of the source and destination data sources before performing data synchronization. It is recommended to execute data synchronization during off-peak hours. Otherwise, full data initialization will consume certain read and write resources of both the source and destination data sources, leading to increased database load.
- Ensure that each table in the synchronized objects has a primary key or unique constraint, and column names are unique to avoid duplicating the same data during synchronization.

Operation Steps

- 1. Log in to the NineData console.
- 2. Click on **Replication** in the left navigation bar.
- 3. On the **Replication** page, click on the **Create Replication** button in the upper right corner.
- 4. On the **Source & Target** tab, configure according to the table below, and then click **Next**.

Parameter	Description
Name	Enter the name of the data synchronization task. For ease of later retrieval and management, please use a meaningful name. Supports up to 64 characters.
Source	The Sybase data source where the synchronization objects are located.
Target	The PostgreSQL data source that receives the synchronized objects.
Replication Mode	Select Uni-directional .
Туре	Keep the default configuration, which is to select Full .
If target table already exists	 Pre-Check Error and Stop Task: Stops the task if data is detected in the target table during the precheck phase. Ignore the existing data and append: Ignores the data detected in the target table during the precheck phase and

Parameter	Description
	appends other data.
	 Clear the existing data before write: Deletes the data
	detected in the target table during the precheck phase and
	rewrites it.

5. On the **Objects** tab, configure the following parameters, and then click **Next**.

Parameter	Description
Replication Objects	Select the content to be copied. You can choose All Objects to copy all contents of the source database, or choose Customized Object , select the content to be copied in the Source Object list, and click > to add it to the Target Object list on the right.

6. (Optional Step) The data replication feature provides configuration file templates for batch import. Click Import Config at the top right corner of the Objects tab, and select Download Template to download the configuration file template to your local machine. After editing, click Upload to upload the configuration file for batch import.

Parameter	Description
source_table_name	Name of the source table containing the objects to be synchronized.
destination_table_name	Name of the target table receiving the synchronized objects.
source_schema_name	Name of the source schema containing the objects to be synchronized.
destination_schema_name	Name of the target schema receiving the synchronized objects.

Parameter	Description
source_database_name	Name of the source database containing the objects to be synchronized.
target_database_name	Name of the target database receiving the synchronized objects.
column_list	List of columns to be synchronized.
extra_configuration	Additional configuration information, where you can configure the following: • Column Mapping: column_name, destination_column_name • Column Value: column_value • Data Filtering: filter_condition

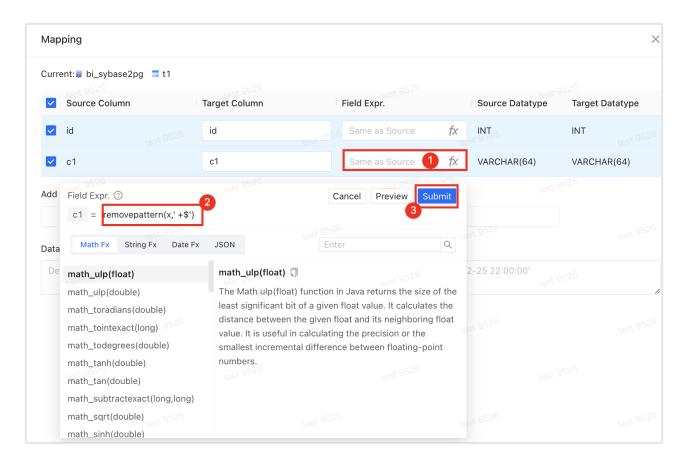
О ТІР

• An example of the extra_configuration content is as follows:

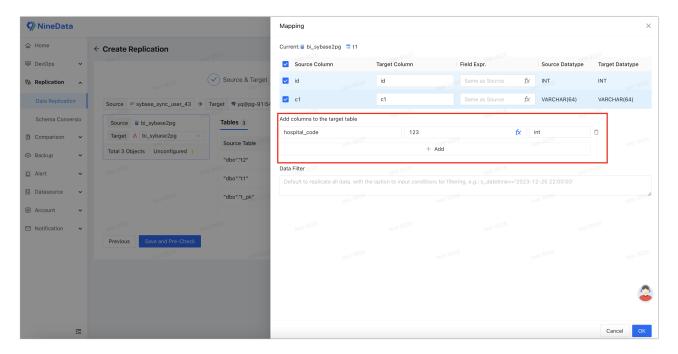
```
"column_name": "created_time", // Original column name to
be mapped
   "destination_column_name": "migrated_time", // Target
column name mapped to "migrated_time"
   "column_value": "current_timestamp()", // Change the
column value to the current timestamp
   "filter_condition": "id != 0" // Only synchronize rows
where ID is not 0.
}
```

• For a complete example of the configuration file, refer to the downloaded template.

- 7. In the **Mapping** tab, the system defaults to selecting the same name database from the target data source. If it does not exist, you need to manually select the target database, then perform the following sub-steps as needed, and finally click **Save and Pre-Check**.
 - Using Field Expr.: To remove trailing spaces from a field name, click on Mapping & Filtering on the right side of the target table, find the field you want to adjust, click the blank box under the Field Expr. column, enter removepattern(x,' +\$') in the Enter text box, click Submit, and then click OK.



Adding a new column in the target table: To add a column named [hospital_code] with
a constant value, click on Mapping & Filtering on the right side of the target table, then
click Add columns to the target table under Add, fill in the column name, column
expression, and column type sequentially, then click OK.

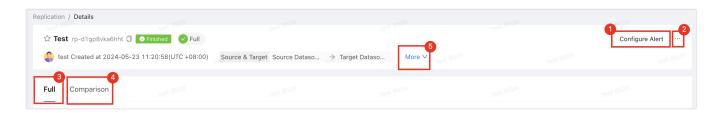


- Refresh Metadata: If there are updates in the source and target data sources during the mapping configuration stage, click the Refresh Metadata button at the top right corner of the page to re-fetch the information from the source and target data sources.
- Setting Data Filter: Click on Mapping & Filtering on the right side of the page, and in the text box under Data Filter, enter emp_no>=10005, so data with emp_no less than 10005 will not be synchronized to the target data source.
- 8. On the **Pre-check** tab, wait for the system to complete the precheck. After the precheck passes, click **Launch**.
 - You can check Enable data consistency comparison. After the synchronization task is completed, the data consistency comparison based on the source data source is automatically started to ensure the consistency of data on both ends. The comparison task will start automatically after full replication is completed.
 - If the precheck fails, you need to click the **Details** in the **Actions** column on the right side of the target check item to troubleshoot the cause of the failure, manually fix it, and then click **Check Again** to perform the precheck again until it passes.
 - For check items with **Result** as **Warning**, you can repair or ignore them according to the specific situation.
- 9. On the **Launch** page, when prompted **Launch Successfully**, the synchronization task starts running. At this point, you can perform the following operations:

- Click View Details to view the execution status of each stage of the synchronization task.
- Click **Back to list** to return to the **Replication** task list page.

View Synchronization Results

- 1. Log in to the NineData console.
- 2. Click on **Replication** > **Data Replication** in the left navigation bar.
- 3. On the **Replication** page, click on the **Task ID** of the target synchronization task, and the page will show as follows.



Index	Function	Description
1	Configure Alerts	After configuring alerts, the system will notify you through your chosen method in case of task failure.
2	More	 Pause: Pause the task. Only tasks with the status Running can be paused. Duplicate: Create a new replication task with the same configuration as the current task. Terminate: End tasks that are incomplete or in listening mode (i.e., incremental synchronization). Once the task is terminated, it cannot be restarted. Please proceed with caution. If the synchronization objects include triggers, the trigger replication option will pop up. Please select as needed. Delete: Delete the task. Once the task is deleted, it cannot be recovered. Please proceed with

Index	Function	Description
		caution.
3	Full Replication (Displayed in scenarios with full replication)	 Display the progress and detailed information of full replication. Click on Monitor on the right side of the page to view various monitoring indicators during full replication. Click on Log on the right side of the page to view the execution log of full replication. Click on C on the right side of the page to view the latest information.
4	Data Comparison	Display the comparison results between the source data source and the target data source. If you haven't enabled data comparison, click on Enable Comparison on the page. Click on Re-compare on the right side of the page to reinitiate the comparison between the current source and target endpoints. Click on Stop on the right side of the page to immediately stop the comparison task once it has started. Click on Log on the right side of the page to view the execution log of consistency comparison. Click on Monitor (Displayed only in data comparison) on the right side of the page to view the trend graph of RPS (records per second) comparison. Click on Details to view earlier records. Click on in the Actions column on the right side of the comparison list (Displayed under the Data tab only in case of inconsistencies) to view the comparison details between the source and target endpoints.

Index	Function	Description
		 Click on in the Actions column on the right side of the comparison list (Displayed in case of inconsistencies) to generate change SQL. You can directly copy this SQL and execute it on the target data source to modify the inconsistent content.
5	Expand	Display detailed information of the current replication task.

Appendix: Precheck Item List

Check Item	Check Content
Source object existence check	Check whether the table on the target database currently exists on the source database
Source datasource connection check	Check the status of the source datasource, database connectable, and verify the username and password
Target datasource connection check	Check the status of the target datasource, database connectable, and verify the username and password
Target databse privilege check	Check whether the account privileges of the target database meet the requirements
Source databse privilege check	Check whether the account privileges of the source database meet the requirements
Target database data existence check	Check whether the object to be replicated not empty in the target database

Check Item	Check Content
Objects with the same name in the target database existence check	Check whether the object to be replicated already exists in the target database
Cyclic replication check	Check for replication loops

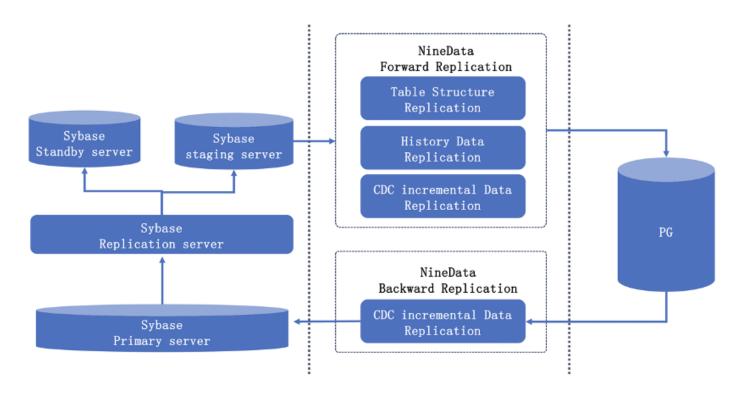


Bidirectional Replication Between Sybase and PostgreSQL

Bidirectional Replication Between Sybase and PostgreSQL

NineData data replication supports replicating data from Sybase to PostgreSQL, while data newly written to PostgreSQL will also be streamed back to Sybase in real-time, achieving real-time consistency between both ends.

Introduction to the Bidirectional Replication Architecture from Sybase to PostgreSQL



As shown in the diagram above, if a bidirectional replication is not established between Sybase instances, the link for data flowing back from PostgreSQL to Sybase must point to the Sybase primary node. The specific replication process is as follows:

- 1. **Forward replication**: The forward replication link copies data from Sybase to PostgreSQL. In this link, replication data initialization can be performed.
- 2. **Backward replication**: Since a bidirectional replication link is not established between the Sybase replica and the Sybase primary, data flowing back from PostgreSQL cannot be

synchronized back to the Sybase primary. Therefore, the backward replication link needs to copy data from PostgreSQL to Sybase primary. In the backward flow link, only incremental data replication is required.

Prerequisites

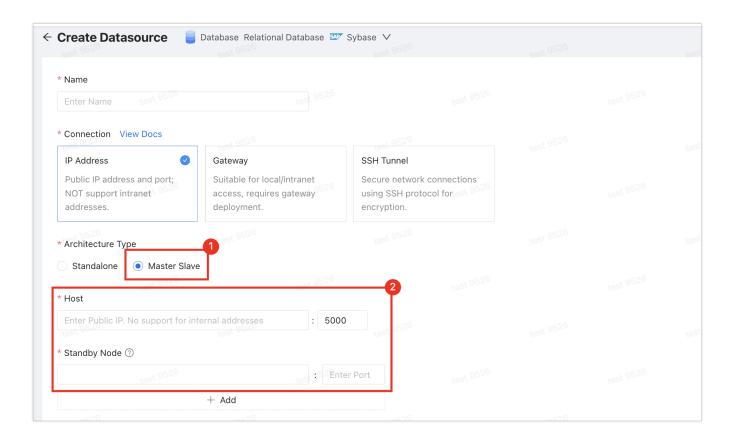
- Source and destination data sources have been added to NineData.
- Tables must be manually created in the destination data source (PostgreSQL), and the table structure must match that of the source data source (Sybase) that needs to be replicated.
- Versions of source and destination data sources are as follows:

Source Data Source	Destination Data Source
Sybase 16.0, 15.7	PostgreSQL 16, 15, 14

• You must have the following permissions on both the source and destination data sources:

Replication Type	Sybase	PostgreSQL
Bidirectional Full & Incremental Replication	sa_role	Database Owner, Table Owner, All Privileges on Schema, Replication Role

- If PostgreSQL is in a master-slave architecture, the virtual IP (VIP) of PostgreSQL must be used when adding the data source to ensure that NineData accesses the master even after a failover.
- If Sybase data source is in a master-slave architecture, both the primary server and the staging server connection addresses must be entered when entering the data source.
 Otherwise, replication from PostgreSQL back to Sybase will fail.



• The session_replication_role parameter in the destination data source (PostgreSQL) has been set to replica. If not set, execute the following statement in the destination data source:

```
SET session_replication_role = 'replica';
```

Precautions

- Evaluate the performance of the source and destination data sources before performing data synchronization. It is recommended to execute data synchronization during off-peak hours.
 Otherwise, full data initialization will consume certain read and write resources of both the source and destination data sources, leading to increased database load.
- Ensure that each table in the synchronized objects has a primary key or unique constraint, and column names are unique to avoid duplicating the same data during synchronization.

Step One: Add Multi-Active Tags to All Data Sources Participating in Replication

To prevent data cyclic replication, you need to add multi-active tags to all data sources participating in replication.

- 1. Log in to the NineData console.
- 2. In the left navigation pane, click **Datasource** > **Datasource**.
- 3. Click the ID of the destination data source to open the **Details** page.
- 4. In the Data Source Details area (which includes information such as data source name, ID, creator, and creation time), click **More**.
- 5. Find **Multi-Active Tag** and click the **□** icon on the right.
- 6. Enter the multi-active tag and click **OK**.
 - Multi-active tags can contain 1 to 64 characters.
 - Multi-active tags must be globally unique and cannot be repeated with other multiactive tags.

Step Two: Create Bidirectional Replication Tasks Between Sybase and PostgreSQL

- 1. Log in to the NineData console.
- 2. Click on **Replication** in the left navigation bar.
- 3. On the **Replication** page, click on the **Create Replication** button in the upper right corner.
- 4. On the Source & Target tab, configure according to the table below, and then click Next.

Parameter	Description
Name	Enter the name of the data synchronization task. For ease of later retrieval and management, please use a meaningful name. Supports up to 64 characters.
Source	The Sybase data source where the synchronization objects are located.
Target	The PostgreSQL data source that receives the synchronized objects.
Replication Mode	Select Bi-directional .

Parameter	Description
Forward Type	The replication type from the source data source to the target data source. Keep the default setting. If you only need to perform incremental replication, you can uncheck Full .
Reverse Type	The replication type from the target data source to the source data source. Only Incremental is supported and is selected by default and cannot be turned off.
If target table already exists	 Pre-Check Error and Stop Task: Stop the task when data is detected in the target table during the precheck phase. Ignore the existing data and append: Ignore the data detected in the target table during the precheck phase and append other data. Clear the existing data before write: Delete the data detected in the target table during the precheck phase and rewrite it.

5. On the **Objects** tab, configure the following parameters, and then click **Next**.

Parameter	Description
Replication Objects	Select the content to be copied. You can choose All Objects to copy all contents of the source database, or choose Customized Object , select the content to be copied in the Source Object list, and click > to add it to the Target Object list on the right.

6. (Optional Step) The data replication feature provides configuration file templates for batch import. Click Import Config at the top right corner of the Objects tab, and select Download Template to download the configuration file template to your local machine. After editing, click Upload to upload the configuration file for batch import.

Parameter	Description
source_table_name	Name of the source table containing the objects to be synchronized.
destination_table_name	Name of the target table receiving the synchronized objects.
source_schema_name	Name of the source schema containing the objects to be synchronized.
destination_schema_name	Name of the target schema receiving the synchronized objects.
source_database_name	Name of the source database containing the objects to be synchronized.
target_database_name	Name of the target database receiving the synchronized objects.
column_list	List of columns to be synchronized.
extra_configuration	Additional configuration information, where you can configure the following:



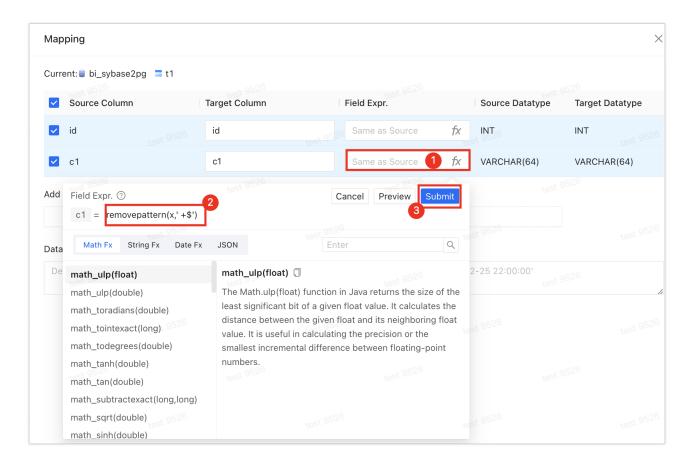
• An example of the extra_configuration content is as follows:

```
{
   "column_name": "created_time", // Original column name to
be mapped
```

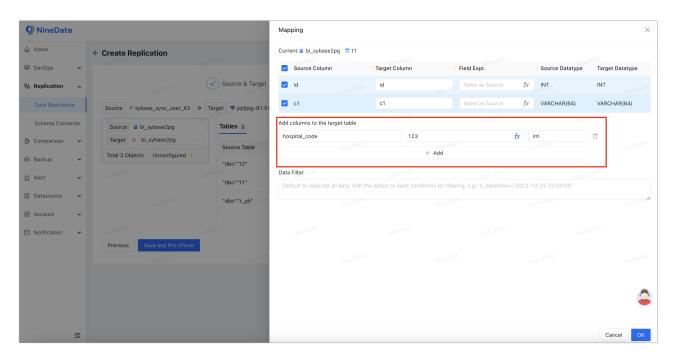
```
"destination_column_name": "migrated_time", // Target
column name mapped to "migrated_time"
    "column_value": "current_timestamp()", // Change the
column value to the current timestamp
    "filter_condition": "id != 0" // Only synchronize rows
where ID is not 0.
}

o For a complete example of the configuration file, refer to the downloaded template.
```

- 7. In the **Mapping** tab, the system defaults to selecting the same name database from the target data source. If it does not exist, you need to manually select the target database, then perform the following sub-steps as needed, and finally click **Save and Pre-Check**.
 - Using Field Expr.: To remove trailing spaces from a field name, click on Mapping & Filtering on the right side of the target table, find the field you want to adjust, click the blank box under the Field Expr. column, enter removepattern(x, ' +\$') in the Enter text box, click Submit, and then click OK.

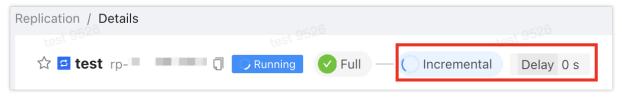


 Adding a new column in the target table: To add a column named (hospital_code) with a constant value, click on Mapping & Filtering on the right side of the target table, then click **Add columns to the target table** under **Add**, fill in the column name, column expression, and column type sequentially, then click **OK**.



- Refresh Metadata: If there are updates in the source and target data sources during
 the mapping configuration stage, click the Refresh Metadata button at the top right
 corner of the page to re-fetch the information from the source and target data sources.
- Setting Data Filter: Click on Mapping & Filtering on the right side of the page, and in the text box under Data Filter, enter emp_no>=10005, so data with emp_no less than 10005 will not be synchronized to the target data source.
- 8. On the **Pre-check** tab, wait for the system to complete the precheck. After the precheck passes, click **Launch**.
 - You can check Enable data consistency comparison. After the synchronization task is completed, the data consistency comparison based on the source data source is automatically started to ensure the consistency of data on both ends. Depending on the Type you choose, the timing of starting Enable data consistency comparison is as follows:
 - Full: Starts after full replication is completed.
 - **Full+Incremental**, **Incremental**: Starts when the incremental data is first consistent with the source data source and **Delay** is set to 0 seconds. You can click

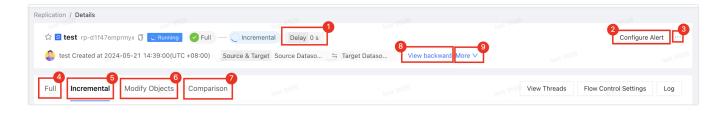
View Details to view the synchronization delay on the Details page.



- If the precheck fails, you need to click the **Details** in the **Actions** column on the right side of the target check item to troubleshoot the cause of the failure, manually fix it, and then click **Check Again** to perform the precheck again until it passes.
- For check items with **Result** as **Warning**, you can repair or ignore them according to the specific situation.
- 9. On the **Launch** page, when prompted **Launch Successfully**, the synchronization task starts running. At this point, you can perform the following operations:
 - Click View Details to view the execution status of each stage of the synchronization task.
 - Click **Back to list** to return to the **Replication** task list page.

View Synchronization Results

- 1. Log in to the NineData console.
- 2. Click on **Replication** > **Data Replication** in the left navigation bar.
- 3. On the **Replication** page, click on the **Task ID** of the target synchronization task. The page explanation is as follows:



Index	Function	Description
1	Sync Delay	The synchronization delay between the source data source and the target data source. 0 seconds indicates no delay between the two ends. At this

Index	Function	Description
		point, you can choose to switch the business to the target data source for smooth migration.
2	Configure Alerts	After configuring alerts, the system will notify you through your chosen method in case of task failure.
3	More	 Pause: Pause the task. Only tasks with the status Running can be paused. Duplicate: Create a new replication task with the same configuration as the current task. Terminate: End tasks that are incomplete or in listening mode (i.e., incremental synchronization). Once the task is terminated, it cannot be restarted. Please proceed with caution. If the synchronization objects include triggers, the trigger replication option will pop up. Please select as needed. Delete: Delete the task. Once the task is deleted, it cannot be recovered. Please proceed with caution.
4	Full Replication (Displayed in scenarios with full replication)	 Display the progress and detailed information of full replication. Click on Monitor on the right side of the page to view various monitoring indicators during full replication. Click on Log on the right side of the page to view the execution log of full replication. Click on C on the right side of the page to view the latest information.
5	Incremental Replication (Displayed in scenarios with	Display various monitoring indicators for incremental replication.

Index	Function	Description
	incremental replication)	 Click on Log on the right side of the page to view the execution log of incremental replication. Click on C on the right side of the page to view the latest information.
6	Modify Object	 Display modification records of synchronized objects. Click on Modify Objects on the right side of the page to configure synchronized objects. Click on ^C on the right side of the page to view the latest information.
7	Data Comparison	Display the comparison results between the source data source and the target data source. If you haven't enabled data comparison, click on Enable Comparison on the page. Click on Re-compare on the right side of the page to reinitiate the comparison between the current source and target endpoints. Click on Stop on the right side of the page to immediately stop the comparison task once it has started. Click on Log on the right side of the page to view the execution log of consistency comparison. Click on Monitor (Displayed only in data comparison) on the right side of the page to view the trend graph of RPS (records per second) comparison. Click on Details to view earlier records. Click on in the Actions column on the right side of the comparison list (Displayed under the Data tab only in case of inconsistencies) to view

Index	Function	Description
		 the comparison details between the source and target endpoints. Click on in the Actions column on the right side of the comparison list (Displayed in case of inconsistencies) to generate change SQL. You can directly copy this SQL and execute it on the target data source to modify the inconsistent content.
8	View Reverse	Displayed only in bidirectional replication tasks, click to view the replication details from the target data source to the source data source.
9	Expand	Display detailed information of the current replication task.

Appendix: Precheck Item List

Check Item	Check Content
Source object existence check	Check whether the table on the target database currently exists on the source database
Source datasource connection check	Check the status of the source datasource, database connectable, and verify the username and password
Target datasource connection check	Check the status of the target datasource, database connectable, and verify the username and password
Target databse privilege check	Check whether the account privileges of the target database meet the requirements

Check Item	Check Content
Source databse privilege check	Check whether the account privileges of the source database meet the requirements
Bi-directional multi-active tag check	Check whether the multi-active tag exists
Target database data existence check	Check whether the object to be replicated not empty in the target database
Objects with the same name in the target database existence check	Check whether the object to be replicated already exists in the target database
Cyclic replication check	Check for replication loops

Full Data Migration from PostgreSQL to Sybase

Full Data Migration from PostgreSQL to Sybase

NineData Data Replication supports full data replication from PostgreSQL to Sybase. This section introduces the configuration process.

Prerequisites

- Both the source data source and the target data source have been added to NineData and are accessible by all deployment servers in NineData's environment.
- The target table in Sybase must be created manually, with a structure identical to the source table in PostgreSQL that needs replication.
- Version compatibility for the source and target databases is shown in the table below.

Source Database	Target Database
PostgreSQL 16, 15, 14	Sybase 16.0, 15.7

• You must have the following permissions on the source and target data sources.

Replication Type	PostgreSQL	Sybase
One-Way Full Replication	SELECT on migration objects	SELECT, TRUNCATE, INSERT

Notes

 Before starting data synchronization, evaluate the performance of both the source and target databases, and consider executing the data synchronization during off-peak hours.

- During the full data initialization, a certain amount of read/write resources from both databases will be occupied, potentially increasing the database load.
- Ensure that each table in the synchronization object has a primary key or unique constraint, and that column names are unique. Otherwise, duplicate data may be synchronized.

Procedure

- 1. Log in to the NineData console.
- 2. Click on **Replication** in the left navigation bar.
- 3. On the **Replication** page, click on **Create Replication** in the top right corner.
- 4. On the Source & Target tab, configure according to the table below and click on Next.

Parameter	Description	
Name	Enter the name of the data synchronization task. To facilitate subsequent search and management, please use meaningful names as much as possible. Up to 64 characters are supported.	
Source	The PostgreSQL data source where the synchronization object is located.	
Target	The Sybase data source that receives the synchronization object.	
Туре	Keep the default configuration, that is, select Full .	
If target table already exists	 Pre-Check Error and Stop Task: Stop the task if data is detected in the target table during the pre-inspection stage. Ignore the existing data and append: Ignore the data if it is detected in the target table during the pre-inspection stage, and append other data. Clear the existing data before write: Delete the data if it is detected in the target table during the pre-inspection stage, and rewrite it. 	

5. On the **Objects** tab, configure the following parameters, and then click on **Next**.

Parameter	Description
Replication Objects	Choose the content to be replicated. You can select All Objects to replicate all content from the source database, or you can select Customized Object , select the content to be replicated in the Source Object list, and click > to add to the right Target Object list.

6. (Optional step) The data replication feature provides a batch import configuration file template. Click on **Import Config** in the top right corner of the **Objects** tab, and then click **Download Template** to download the configuration file template to your local machine. After editing, click on **Upload** to upload the configuration file to achieve batch import.

Parameter	Description
source_table_name	The source table name where the object to be synchronized is located.
destination_table_name	The target table name that receives the synchronized object.
source_schema_name	The source Schema name where the object to be synchronized is located.
destination_schema_name	The target Schema name that receives the synchronized object.
source_database_name	The source database name where the object to be synchronized is located.
target_database_name	The target database name that receives the synchronized object.
column_list	The list of fields to be synchronized.

Parameter	Description	
extra_configuration	Additional configuration information, where you can configure the following information: • Field mapping: column_name, destination_column_name • Field value: column_value • Data filtering: filter_condition	

TIP

• An example of extra_configuration is as follows:

```
"column_name": "created_time", //Specify the original
column name for column name mapping
   "destination_column_name": "migrated_time", //Map the
target column name to "migrated_time"
   "column_value": "current_timestamp()", //Change the field
value of the column to the current timestamp
   "filter_condition": "id != 0" //Only rows with ID not
equal to 0 will be synchronized.
}
```

- For the overall example content of the configuration file, please refer to the downloaded template.
- 7. On the **Mapping** tab, the system defaults to selecting the database with the same name in the target data source. If it does not exist, you need to manually select the target database, and then perform the following sub-steps as needed, and finally click on **Save and Pre-Check**.
 - Refresh Metadata: If there are updates in the source and target data sources during the mapping configuration phase, you can click the Refresh Metadata button in the top

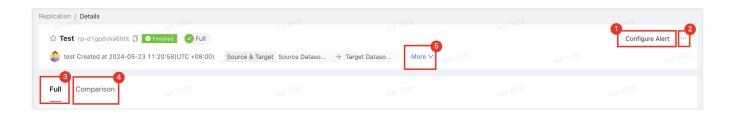
- right corner of the page to refresh the information of the source and target data sources.
- Set Data Filter: Click on Mapping & Filtering on the right side of the page, and enter emp_no>=10005 in the text box below Data Filter, so that data with emp_no less than 10005 will not be synchronized to the target data source.
- 8. On the **Pre-check** tab, wait for the system to complete the pre-inspection. After the pre-inspection is passed, click on **Launch**.

◯ TIP

- You can check Enable data consistency comparison. After the synchronization
 task is completed, automatically start the data consistency comparison based on
 the source data source to ensure data consistency on both sides. The comparison
 task will automatically start after the full replication is completed.
- If the pre-inspection does not pass, you need to click on **Details** in the **Actions** column on the right side of the target inspection item, troubleshoot the cause of the failure, manually repair it, and then click on **Check Again** to re-execute the pre-inspection until it passes.
- If the **Result** is **Warning**, you can repair or ignore it according to the specific situation.
- 9. On the **Launch** page, it prompts **Launch Successfully**, and the synchronization task starts running. At this time, you can perform the following operations:
 - Click on View Details to view the execution of each stage of the synchronization task.
 - Click on Back to list to return to the Replication task list page.

Viewing Synchronization Results

- 1. Log in to the NineData console.
- 2. Click on **Replication** > **Data Replication** in the left navigation bar.
- 3. On the **Replication** page, click on the **Task ID** of the target synchronization task, and the page will show as follows.



Index	Function	Description	
1	Configure Alerts	After configuring alerts, the system will notify you through your chosen method in case of task failure.	
2	More	 Pause: Pause the task. Only tasks with the status Running can be paused. Duplicate: Create a new replication task with the same configuration as the current task. Terminate: End tasks that are incomplete or in listening mode (i.e., incremental synchronization). Once the task is terminated, it cannot be restarted. Please proceed with caution. If the synchronization objects include triggers, the trigger replication option will pop up. Please select as needed. Delete: Delete the task. Once the task is deleted, it cannot be recovered. Please proceed with caution. 	
3	Full Replication (Displayed in scenarios with full replication)	 Display the progress and detailed information of full replication. Click on Monitor on the right side of the page to view various monitoring indicators during full replication. Click on Log on the right side of the page to view the execution log of full replication. Click on C on the right side of the page to view the latest information. 	

Index	Function	Description
4	Data Comparison	Display the comparison results between the source data source and the target data source. If you haven't enabled data comparison, click on Enable Comparison on the page. Click on Re-compare on the right side of the page to reinitiate the comparison between the current source and target endpoints. Click on Stop on the right side of the page to immediately stop the comparison task once it has started. Click on Log on the right side of the page to view the execution log of consistency comparison. Click on Monitor (Displayed only in data comparison) on the right side of the page to view the trend graph of RPS (records per second) comparison. Click on Details to view earlier records. Click on in the Actions column on the right side of the comparison list (Displayed under the Data tab only in case of inconsistencies) to view the comparison details between the source and target endpoints. Click on in the Actions column on the right side of the comparison list (Displayed in case of inconsistencies) to generate change SQL. You can directly copy this SQL and execute it on the target data source to modify the inconsistent content.
5	Expand	Display detailed information of the current replication task.

Appendix: Pre-check List

检查项	检查内容
Source datasource connection check	Verify the source database instance is reachable and the username and password are correct
Target datasource connection check	Verify the target database instance is reachable and the username and password are correct
Target databse privilege check	Check whether the account privileges of the target database meet the requirements
Source databse privilege check	Check whether the account privileges of the source database meet the requirements
Target database data existence check	Check whether the object to be replicated not empty in the target database
Objects with the same name in the target database existence check	Check whether the object to be replicated already exists in the target database

Data replication

Real-Time Incremental Data Migration from PostgreSQL to Sybase

Real-Time Incremental Data Migration from PostgreSQL to Sybase

NineData Data Replication supports data migration from PostgreSQL to Sybase, including both full and incremental replication.

Prerequisites

- The source and target data sources have been added to NineData and are accessible by all deployment servers within the NineData environment.
- The target table in Sybase must be created manually, with a structure identical to the source table in PostgreSQL that needs replication.
- Version compatibility for the source and target databases is shown in the table below.

Source Database	Target Database
PostgreSQL 16, 15, 14	Sybase 16.0, 15.7

• The following permissions are required on the source and target data sources.

Replication Type	PostgreSQL	Sybase
Full Replication + Incremental Replication	SELECT on migration objects	SELECT, INSERT, UPDATE, DELETE

• For incremental replication, configure the following parameters in the postgresql.conf file. If the location of this file is unknown, you can execute the SQL command SHOW config_file; in the psql client to find it.

- The wal_level parameter of the source data source must be set to logical. To confirm the current value, you can execute SHOW wal_level in the client.
- Set the wal_sender_timeout parameter of the source data source to 0. This parameter is used to terminate replication connections that are idle for longer than the specified milliseconds. The default value is 60000 milliseconds. Setting it to 0 disables the timeout mechanism. To confirm the current value, you can execute SHOW wal_sender_timeout in the client.
- The max_replication_slots parameter of the source data source must be greater than 1. This parameter specifies the maximum number of replication slots the server can support. The default value is 10.
- The max_wal_senders parameter of the source data source must be greater than 1. This parameter specifies the maximum number of concurrent connections. The default value is 10. To confirm the current value, you can execute SHOW max_wal_senders in the client.

Notes

- Before starting data synchronization, evaluate the performance of both the source and target databases, and consider executing the data synchronization during off-peak hours.
 During the full data initialization, a certain amount of read/write resources from both databases will be occupied, potentially increasing the database load.
- Ensure that each table in the synchronization object has a primary key or unique constraint, and that column names are unique. Otherwise, duplicate data may be synchronized.

Procedure

- 1. Log in to the NineData console.
- 2. Click on **Replication** in the left navigation bar.
- 3. On the **Replication** page, click on **Create Replication** in the top right corner.
- 4. On the **Source & Target** tab, configure according to the table below and click on **Next**.

Parameter	Description	
Name	Enter the name of the data synchronization task. To facilitate subsequent search and management, please use meaningful names as much as possible. Up to 64 characters are supported.	
Source	The PostgreSQL data source where the synchronization object is located.	
Target	The Sybase data source that receives the synchronization object.	
Туре	Check Full and Incremental .	
If target table already exists	 Choose according to actual needs. Pre-Check Error and Stop Task: Stop the task if data is detected in the target table during the pre-inspection stage. Ignore the existing data and append: Ignore the data if it is detected in the target table during the pre-inspection stage, and append other data. Clear the existing data before write: Delete the data if it is detected in the target table during the pre-inspection stage, and rewrite it. 	

5. On the **Objects** tab, configure the following parameters, and then click on **Next**.

Parameter	Description
Replication Objects	Choose the content to be replicated. You can select All Objects to replicate all content from the source database, or you can select Customized Object , select the content to be replicated in the Source Object list, and click > to add it to the right Target Object list.

6. (Optional step) The data replication feature provides a batch import configuration file template. Click on **Import Config** in the top right corner of the **Objects** tab, and then click

Download Template to download the configuration file template to your local machine. After editing, click on **Upload** to upload the configuration file to achieve batch import.

Parameter	Description
source_table_name	The source table name where the object to be synchronized is located.
destination_table_name	The target table name that receives the synchronized object.
source_schema_name	The source Schema name where the object to be synchronized is located.
destination_schema_name	The target Schema name that receives the synchronized object.
source_database_name	The source database name where the object to be synchronized is located.
target_database_name	The target database name that receives the synchronized object.
column_list	The list of fields to be synchronized.
extra_configuration	Additional configuration information, where you can configure the following information: • Field mapping: column_name, destination_column_name • Field value: column_value • Data filtering: filter_condition



• An example of extra_configuration is as follows:

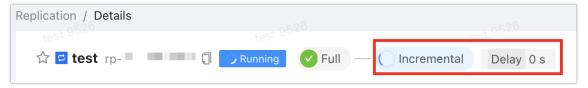
```
"column_name": "created_time", //Specify the original
column name for column name mapping
   "destination_column_name": "migrated_time", //Map the
target column name to "migrated_time"
   "column_value": "current_timestamp()", //Change the field
value of the column to the current timestamp
   "filter_condition": "id != 0" //Only rows with ID not
equal to 0 will be synchronized.
}
```

- For the overall example content of the configuration file, please refer to the downloaded template.
- 7. On the **Mapping** tab, the system defaults to selecting the database with the same name in the target data source. If it does not exist, you need to manually select the target database, and then perform the following sub-steps as needed, and finally click on **Save and Pre-Check**.
 - Refresh Metadata: If there are updates in the source and target data sources during
 the mapping configuration phase, you can click the Refresh Metadata button in the top
 right corner of the page to refresh the information of the source and target data
 sources.
 - Set Data Filter: Click on Mapping & Filtering on the right side of the page, and enter
 emp_no>=10005 in the text box below Data Filter, so that data with emp_no less than
 10005 will not be synchronized to the target data source.
- 8. On the **Pre-check** tab, wait for the system to complete the pre-inspection. After the pre-inspection is passed, click on **Launch**.



- You can check Enable data consistency comparison. After the synchronization
 task is completed, automatically start the data consistency comparison based on
 the source data source to ensure data consistency on both sides. Depending on
 the Type you choose, the timing for starting Enable data consistency
 comparison is as follows:
 - **Full**: Start after full replication is completed.

• **Full+Incremental**, **Incremental**: Start when the incremental data is consistent with the source data source for the first time and **Delay** is 0 seconds. You can click on **View Details** to view the synchronization delay on the **Details** page.



- If the pre-inspection does not pass, you need to click on **Details** in the **Actions** column on the right side of the target inspection item, troubleshoot the cause of the failure, manually fix it, and then click on **Check Again** to re-execute the pre-inspection until it passes.
- If the **Result** is **Warning**, you can fix or ignore it according to the specific situation.
- 9. On the **Launch** page, it prompts **Launch Successfully**, and the synchronization task starts running. At this time, you can perform the following operations:
 - Click on View Details to view the execution status of each stage of the synchronization task.
 - Click on **Back to list** to return to the **Replication** task list page.

Viewing Synchronization Results

- 1. Log in to the NineData console.
- 2. Click on **Replication** > **Data Replication** in the left navigation bar.
- 3. On the **Replication** page, click the **Task ID** of the target synchronization task, and the page description is as follows.

No.	Function	Description
1	Synchronization Delay	The data synchronization delay between the source and target data sources. O seconds indicate no delay between the two ends. At this point, you can choose to switch the business to the target data source to achieve a smooth migration.

No.	Function	Description
2	Configure Alerts	After configuring alerts, the system will notify you through the method you choose when the task fails.
3	More	 Pause: Pause the task, only tasks in running status are selectable. Duplicate: Create a new replication task with the same configuration as the current task. Terminate: End tasks that are not completed or in listening mode (i.e., incremental synchronization). After terminating the task, it cannot be restarted, please operate with caution. If the synchronization object includes triggers, trigger replication options will pop up, please choose as needed. Delete: Delete the task, the task cannot be recovered after deletion, please operate with caution.
4	Full Replication (displayed in scenarios including full replication)	Display the progress and detailed information of full replication. Click on Monitor on the right side of the page: View various monitoring metrics during the full replication process. Click on Log on the right side of the page: View the execution logs of full replication. Click on C on the right side of the page: View the latest information.
5	Incremental Replication (displayed in scenarios including incremental replication)	Display various monitoring metrics of incremental replication. • Click on Log on the right side of the page: View the execution logs of incremental replication.

No.	Function	Description
		 Click on C on the right side of the page: View the latest information.
6	Modify Objects	 Display the modification records of synchronization objects. Click on Modify Objects on the right side of the page to configure synchronization objects. Click on C on the right side of the page: View the latest information.
7	Data Comparison	Display the results of the comparison between the source and target data sources. If you have not enabled data comparison, please click on Enable Comparison in the page. Click on Re-compare on the right side of the page: Re-initiate the comparison between the current source and target data. Click on Stop on the right side of the page: After the comparison task starts, you can click this button to immediately stop the comparison task. Click on Log on the right side of the page: View the execution logs of consistency comparison. Click on Monitor on the right side of the page (only displayed during data comparison): View the trend chart of comparison RPS (records compared per second). Click on Details to view records from earlier times. Click on in the Actions column on the right side of the comparison details between the source and target ends. Click on in the Actions column on the right side of the comparison details between the source and target ends.

No.	Function	Description
		inconsistencies are found): Generate change SQL, you can directly copy this SQL to the target data source for execution, modifying the inconsistent content.
8	Expand	Display detailed information of the current replication task.

Appendix: Pre-check List

Check Item	Description
Source datasource connection check	Verify the source database instance is reachable and the username and password are correct
Target datasource connection check	Verify the target database instance is reachable and the username and password are correct
Target databse privilege check	Check whether the account privileges of the target database meet the requirements
Source databse privilege check	Check whether the account privileges of the source database meet the requirements
Target database data existence check	Check whether the object to be replicated not empty in the target database
Objects with the same name in the target database existence check	Check whether the object to be replicated already exists in the target database
Check wal_level	Verify that the wal_level of the source data source is set to logical

Check Item	Description
Check max_wal_senders	Ensure that max_wal_senders meets the required number of replication connections
Check max_replication_slots	Ensure that max_replication_slots meets the required number of replication slots

n Data replication

Processing Data via ETL

NineData data replication supports ETL (Extract, Transform, Load) functionality. This means it can read data from the source, transform it into the desired state, and then replicate the processed data to the target.

This article uses a full migration from Sybase to PostgreSQL as an example to explain how to configure ETL.

Operation Steps

- 1. Log in to the NineData console.
- 2. Click on **Replication** in the left navigation bar.
- 3. On the **Replication** page, click on the **Create Replication** button in the upper right corner.
- 4. On the **Source & Target** tab, configure according to the table below, and then click **Next**.

Parameter	Description
Name	Enter the name of the data synchronization task. For ease of later retrieval and management, please use a meaningful name. Supports up to 64 characters.
Source	The Sybase data source where the synchronization objects are located.
Target	The PostgreSQL data source that receives the synchronized objects.
Replication Mode	Select Uni-directional .
Туре	Keep the default configuration, which is to select Full .

Parameter	Description
If target table already exists	 Pre-Check Error and Stop Task: Stops the task if data is detected in the target table during the precheck phase. Ignore the existing data and append: Ignores the data detected in the target table during the precheck phase and appends other data. Clear the existing data before write: Deletes the data
	detected in the target table during the precheck phase and rewrites it.

5. On the **Objects** tab, configure the following parameters, and then click **Next**.

Parameter	Description
Replication Objects	Select the content to be copied. You can choose All Objects to copy all contents of the source database, or choose Customized Object , select the content to be copied in the Source Object list, and click > to add it to the Target Object list on the right.

6. (Optional Step) The data replication feature provides configuration file templates for batch import. Click **Import Config** at the top right corner of the **Objects** tab, and select **Download Template** to download the configuration file template to your local machine. After editing, click **Upload** to upload the configuration file for batch import.

Parameter	Description
source_table_name	Name of the source table containing the objects to be synchronized.
destination_table_name	Name of the target table receiving the synchronized objects.
source_schema_name	Name of the source schema containing the objects to be synchronized.

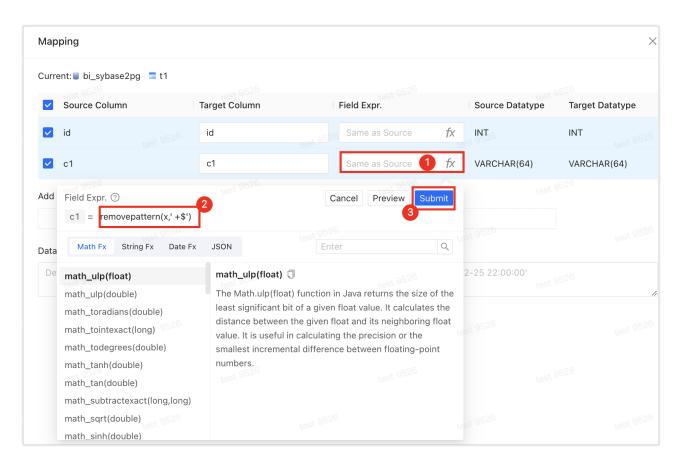
Parameter	Description
destination_schema_name	Name of the target schema receiving the synchronized objects.
source_database_name	Name of the source database containing the objects to be synchronized.
target_database_name	Name of the target database receiving the synchronized objects.
column_list	List of columns to be synchronized.
extra_configuration	Additional configuration information, where you can configure the following: Column Mapping: column_name, destination_column_name Column Value: column_value Data Filtering: filter_condition



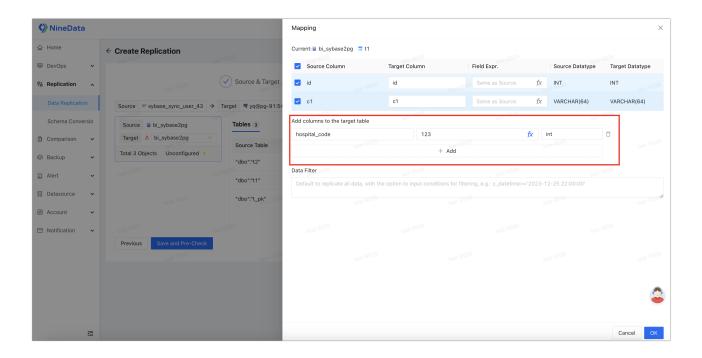
• An example of the extra_configuration content is as follows:

```
"column_name": "created_time", // Original column name to
be mapped
   "destination_column_name": "migrated_time", // Target
column name mapped to "migrated_time"
   "column_value": "current_timestamp()", // Change the
column value to the current timestamp
   "filter_condition": "id != 0" // Only synchronize rows
where ID is not 0.
}
```

- For a complete example of the configuration file, refer to the downloaded template.
- 7. In the **Mapping** tab, the system defaults to selecting the same name database from the target data source. If it does not exist, you need to manually select the target database, then perform the following sub-steps as needed, and finally click **Save and Pre-Check**.
 - Using Field Expr.: To remove trailing spaces from a field name, click on Mapping & Filtering on the right side of the target table, find the field you want to adjust, click the blank box under the Field Expr. column, enter removepattern(x,' +\$') in the Enter text box, click Submit, and then click OK.



Adding a new column in the target table: To add a column named hospital_code with
a constant value, click on Mapping & Filtering on the right side of the target table, then
click Add columns to the target table under Add, fill in the column name, column
expression, and column type sequentially, then click OK.



- Setting Data Filter: Click on Mapping & Filtering on the right side of the page, and in the text box under Data Filter, enter emp_no>=10005, so data with emp_no less than 10005 will not be synchronized to the target data source.
- 8. On the **Pre-check** tab, wait for the system to complete the precheck. After the precheck passes, click **Launch**.
 - You can check Enable data consistency comparison. After the synchronization task is completed, the data consistency comparison based on the source data source is automatically started to ensure the consistency of data on both ends. The comparison task will start automatically after full replication is completed.
 - If the precheck fails, you need to click the **Details** in the **Actions** column on the right side of the target check item to troubleshoot the cause of the failure, manually fix it, and then click **Check Again** to perform the precheck again until it passes.
 - For check items with **Result** as **Warning**, you can repair or ignore them according to the specific situation.
- 9. On the **Launch** page, when prompted **Launch Successfully**, the synchronization task starts running. At this point, you can perform the following operations:
 - Click View Details to view the execution status of each stage of the synchronization task.
 - Click **Back to list** to return to the **Replication** task list page.



Import Replication Object Configurations via Configuration Files

NineData data replication provides configuration file templates for batch import. After customizing the configuration file, you can upload it to NineData to quickly configure replication objects.

This article uses a migration from Sybase to PostgreSQL as an example to explain how to import configuration files.

Operation Steps

- 1. Log in to the NineData console.
- 2. Click on **Replication** in the left navigation bar.
- 3. On the Replication page, click on Create Replication at the top right corner.
- 4. On the **Source & Target** tab, configure the settings according to the table below and click **Next**.

Parameter	Description
Name	Enter the name of the data synchronization task. To facilitate subsequent search and management, use a meaningful name. Up to 64 characters are supported.
Source	The Sybase data source where the objects to be synchronized are located.
Target	The PostgreSQL data source that will receive the synchronized objects.

Parameter	Description
Replication Mode	 Select the replication mode. Uni-directional: Unidirectional data replication task from the source data source to the target data source. Bi-directional: Bidirectional incremental data replication task between the source and target data sources.
Туре	 Displayed only when Replication Mode is Uni-directional. Full: Synchronize all objects and data from the source data source, i.e., full data replication. Incremental: After full synchronization, perform incremental synchronization based on the logs of the source data source.
Forward Type	Type of replication from the source data source to the target data source, displayed only when Replication Mode is Bi-directional . The default selection is Incremental and cannot be disabled.
Forward Incremental Started	 Required when Replication Mode is Bi-directional and Forward Type is Incremental. From Started: Perform incremental replication based on the start time of the current replication task. Customized Time: Select the start time for incremental replication. You can select the time zone according to your business location.
Reverse Type	Type of replication from the target data source to the source data source, displayed only when Replication Mode is Bi-directional . Only Incremental is supported, and the default selection cannot be disabled.
If target table already exists	 Pre-Check Error and Stop Task: Stop the task if data is detected in the target table during the pre-check phase. Ignore the existing data and append: Ignore the data in the target table detected during the pre-check phase and append other data.

Parameter	Description
	 Clear the existing data before write: Delete the data in the target table detected during the pre-check phase and write the new data.

5. On the **Objects** tab, select and configure the replication objects. Since this article focuses on configuring through file import, please directly click **Import Config** at the top right corner, then click **Download Template** to download the configuration file template to your local machine. After editing, click **Upload** to upload the configuration file for batch import.

Parameter	Description
source_table_name	Name of the source table containing the objects to be synchronized.
destination_table_name	Name of the target table receiving the synchronized objects.
source_schema_name	Name of the source schema containing the objects to be synchronized.
destination_schema_name	Name of the target schema receiving the synchronized objects.
source_database_name	Name of the source database containing the objects to be synchronized.
target_database_name	Name of the target database receiving the synchronized objects.
column_list	List of columns to be synchronized.
extra_configuration	Additional configuration information, where you can configure the following:

Parameter	Description
	 Column Mapping: column_name,
	destination_column_name
	Column Value: column_value
	 Data Filtering: filter_condition

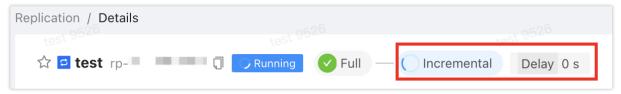
◯ TIP

• An example of the extra_configuration content is as follows:

```
"column_name": "created_time", // Original column name to
be mapped
   "destination_column_name": "migrated_time", // Target
column name mapped to "migrated_time"
   "column_value": "current_timestamp()", // Change the
column value to the current timestamp
   "filter_condition": "id != 0" // Only synchronize rows
where ID is not 0.
}
```

- For a complete example of the configuration file, refer to the downloaded template.
- 6. In the **Mapping** tab, the system defaults to selecting the same name database from the target data source. If it does not exist, you need to manually select the target database, and finally click **Save and Pre-Check**.
- 7. On the **Pre-check** tab, wait for the system to complete the pre-check. After the pre-check passes, click **Launch**.
 - You can check Enable data consistency comparison. After the synchronization task is completed, this will automatically start a data consistency comparison based on the source data source to ensure data consistency on both ends. Depending on the Type you selected, the timing for starting Enable data consistency comparison is as follows:

- **Full**: Starts after the full data replication is completed.
- **Full** + **Incremental**, **Incremental**: Starts when the incremental data first matches the source data source and **Delay** is 0 seconds. You can click **View Details** to view the synchronization delay on the **Details** page.



- If the pre-check fails, click **Details** in the **Actions** column on the right side of the failed check item to troubleshoot the cause. After manually fixing the issue, click **Check Again** to re-run the pre-check until it passes.
- For check items where **Result** is **Warning**, you can choose to fix or ignore them based on the specific situation.
- 8. On the **Launch** page, you will see a prompt **Launch Successfully**, indicating that the synchronization task has started. At this point, you can perform the following actions:
 - Click View Details to view the execution details of each stage of the synchronization task.
 - Click **Back to list** to return to the **Replication** task list page.



Introduction to database comparison

The NineData database comparison function supports the consistency comparison between two data sources. In case of inconsistency, it also generates SQL job to correct the inconsistency. User can execute the SQL job on the target side to achieve data and schema consistency.

Features

- Schema Comparison: Schema comparison is the comparison of metadata (MetaData), which is used to compare whether the definitions of objects in the two databases are consistent. For example, compare whether the columns, indexes, and constraints (primary keys, foreign keys, and constraints) in the two tables are the same.
- Data Comparison: Data comparison is the comparison of user data (UserData), which is used
 to compare whether the data in the two databases are consistent. The feature is usually for
 purpose of data consistency, in case such as, backup and recovery, data migration, and data
 synchronization.

Supported Datasources

Source Datasource	Target Datasource	Data Comparison	Schema Comparison
MySQL	MySQL	V	✓
	PostgreSQL	V	-
	Oracle	V	-
	Doris	V	-
	SelectDB	V	-
	ClickHouse	V	-

Source Datasource	Target Datasource	Data Comparison	Schema Comparison
	Elasticsearch	-	-
	Kafka	-	-
	Greenplum	√	-
	Redshift	-	-
	StarRocks	√	-
	SingleStore	√	-
	TiDB	✓	-
	Hive	√	-
	ADB PostgreSQL	-	-
PostgreSQL	MySQL	✓	-
	PostgreSQL	√	✓
	Oracle	✓	-
	Doris	√	-
	SelectDB	√	-
	ClickHouse	✓	-
	Greenplum	√	-
	StarRocks	√	-

Source Datasource	Target Datasource	Data Comparison	Schema Comparison
	SingleStore	✓	-
	TiDB	✓	-
	Hive	✓	-
	ADB PostgreSQL	✓	-
	Sybase	✓	-
Oracle	MySQL	✓	-
	PostgreSQL	✓	-
	Oracle	✓	✓
	OB-Oracle	✓	✓
	Doris	✓	-
	SelectDB	✓	-
	ClickHouse	✓	-
	Greenplum	✓	-
	StarRocks	✓	-
	SingleStore	✓	-
	TiDB	✓	-
	Hive	V	-

Source Datasource	Target Datasource	Data Comparison	Schema Comparison
	ADB PostgreSQL	✓	-
	SQL Server	✓	✓
SQL Server	Doris	√	-
	StarRocks	√	-
MongoDB	MongoDB	✓	-
Redis	Redis	✓	-
	MySQL	-	-
Kafka	ClickHouse	-	-
	Kafka	-	-
SingleStore	MySQL	✓	-
	PostgreSQL	✓	-
	Oracle	✓	-
	Doris	✓	-
	SelectDB	✓	-
	ClickHouse	√	-
	Greenplum	✓	-
	StarRocks	V	-

Source Datasource	Target Datasource	Data Comparison	Schema Comparison
	SingleStore	√	-
	MySQL	✓	-
	PostgreSQL	✓	-
	Oracle	√	-
	Doris	√	-
StarRocks	SelectDB	√	-
	ClickHouse	✓	-
	Greenplum	✓	-
	StarRocks	√	-
	SingleStore	✓	-
ClickHouse	MySQL	√	-
	PostgreSQL	✓	-
	Oracle	√	-
	Doris	√	-
	SelectDB	✓	-
	ClickHouse	√	-
	Greenplum	V	-

Source Datasource	Target Datasource	Data Comparison	Schema Comparison
	StarRocks	√	-
	MySQL	√	-
	PostgreSQL	√	-
	Oracle	V	-
	Doris	√	-
SelectDB	SelectDB	V	-
	ClickHouse	√	-
	Greenplum	√	-
	StarRocks	V	-
	SingleStore	√	-
Greenplum	MySQL	V	-
	PostgreSQL	√	-
	Oracle	√	-
	Doris	√	-
	SelectDB	✓	-
	ClickHouse	√	-
	Greenplum	V	-

Source Datasource	Target Datasource	Data Comparison	Schema Comparison
	StarRocks	√	-
	SingleStore	✓	-
	TiDB	✓	-
	Hive	√	-
	MySQL	√	-
	PostgreSQL	√	-
ADB PostgreSQL	Oracle	✓	-
	TiDB	✓	-
	ADB PostgreSQL	√	-
Sybase	PostgreSQL	✓	-
TiDB	MySQL	√	-
	PostgreSQL	✓	-
	Oracle	✓	-
	ClickHouse	✓	-
	Greenplum	✓	-
	TiDB	✓	-
	Hive	V	-

Source Datasource	Target Datasource	Data Comparison	Schema Comparison
	ADB PostgreSQL	V	-
	MySQL	V	-
	PostgreSQL	V	-
Hive	Oracle	V	-
Hive	Greenplum	V	-
	TiDB	V	-
	Hive	-	-
DWS	DWS	V	-
DaMeng	DaMeng	V	-
DB2	DB2	√	-
OB-Oracle	OB-Oracle	V	-

Use case

Use case	Description	
Database	When enterprise users manage multi-projects or multiple databases of	
schema	the same type, due to the dispersion or independence of organizations or	
consistency	teams, changes in database structures and objects are not synchronized	
across	to all data centers or sub-projects in time, and application errors are	
departments	reported only after they are discovered. NineData Schema Comparison	
and regions	service regularly compares the metadata between the source and target	
	instances to ensure the consistency of the data structure of all	

Use case	Description	
	databases, and provides repair suggestions to avoid business interruption.	
Cross-region, cross-cloud data verification	In order to achieve low-latency reading and writing at the nearest location, and the ability of databases to be active and reused, enterprises usually build data replication and synchronization between different regions or cloud platforms. However, due to the limitations of the system architecture and other factors, data inconsistency may occur, so periodic data verification is essential. In the case of a huge amount of data, full comparison is expensive and inefficient. NineData Data and schema comparison service is based on group mixed inspection and fast hashing , which provides efficient data consistency check. At the same time, for inconsistencies, NineData provides SQL change scripts, which can be executed on the target side to correct the inconsistency.	
Integrity of data during ELT/ETL process	Both classic ETL and ELT in the cloud-native era require data from multiple data sources to be imported into a data warehouse to support OLAP and BI analysis. However, the source and destination are usually heterogeneous, and data consistency may be affected during data aggregation and schema transformation. NineData Data and schema comparison service supports data comparison between homogeneous and heterogeneous databases, which greatly reduces the risk of data inconsistency in the process.	



Data comparison

The NineData database comparison function supports the comparison of user data (UserData), for purpose of data consistency, in case such as, backup and recovery, data migration, and data synchronization. This article describes how to perform data comparison.

Prerequisite

The data sources that need to be compared have been added to NineData.

Steps

- 1. Log in to NineData console.
- 2. In the left navigation bar, click **Comparison > Data**.
- 3. In the upper right corner of the **Data Comparison** page, click **Create Data Comparison**.
- 4. On the **Source & Target** tab, configure according to the following table, and click **Next**.

Parameters	Desciption
Name	Enter the name of the comparison task. For the convenience of subsequent search and management, please try to use a meaningful name. Up to 64 characters are supported.
Source	The source of the task, and the system will compare based on the structure of the data source.
Target	The target for the comparison task.
Frequency	Select how often the comparison is performed. • One-time: Perform the comparison task only once. User can manually perform the comparison task multiple times.

Parameters	Desciption	
	Periodic : Periodically perform a compare task.	
Periodic (Periodic comparison only)	Choose how often to compare, user can choose any one or more days from Monday to Sunday, or check Every day to compare daily.	
Time Zone (Periodic comparison only)	Select the time zone according to the actual business situation. This parameter affects the time displayed in the comparison task.	
Launch Time (periodic comparison only)	Select the time to launch the comparison task, from 00:00 to 23:59 . User can also click Now to select the current time as the start time.	
Method	Choose how the data will be compared. • Full Data: Compare all selected data.	
	Quick Comparison : Choose one or more of the following	
	comparisons.Compare the total number of records in the table.	
	 Find all columns of numeric type and bool type in the table, and compare the maximum value, minimum value, and average value. 	
	 According to the Sample Proportion configuration you set, a random sample of corresponding proportions of records is compared in each table. 	
	Note: For Redis data sources, options related to BigKey are also	
	provided. If the "Only Compare BigKey Value Length, not compare	
	value" option is selected and a BigKey is encountered, only the	
	length of the value is compared, not its content. You can customize	
	the length of the BigKey value by configuring Threshold , which	
	defaults to 10000 . If the length of the value corresponding to a key	

Parameters	Desciption	
	exceeds 10000, the key is considered a BigKey. The unit of the	
	threshold varies depending on the data type, as follows:	
	 For string types: the length of the value corresponding to a key 	
	exceeds 10,000 characters.	
	 For list, set, sorted set, stream types: the length of the value 	
	corresponding to a key exceeds 10,000 elements.	

- 5. On the **Objects** tab, select the content to be compared and click **Next**. Both **Full Instance** and **Customized Object** are supported.
 - **Full Instance**: Select all databases and objects in the source instance.
 - Customized Object: Customize the database and objects to be compared. In the
 Source list on the left, select the databases and objects to be compared, and click
 > add to **Target** list.
- 6. On the **Mapping** tab, configure the mapping relationship between the source and target databases and object names, and click **Save and Pre-Check**.



If the database or object name in the target is different from that in the source, configuration is required. For MySQL-to-MySQL and SQL Server-to-SQL Server comparisons, you can also click **Mapping & Filtering** on the right side of the table list and configure the mapping between the target field name and the source field name.

7. On the **Pre-check** tab, wait for the system to complete the pre-check. After the pre-check passes, click **Launch**.

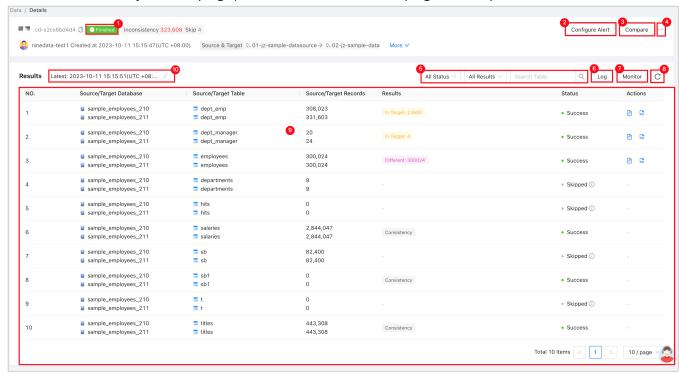


 If the pre-check fails, user need to click the **Details** in the **Actions** column to the right of the target check item to correct the failure, and then perform the precheck again.

- Warnings can be repaired or ignored according to specific circumstances.
- 8. On the **Launch** tab, a message is displayed indicating that **Launch Successfully**. User can do the following:
 - If it is a **Periodic** comparison task, click **Compare** to execute a comparison task immediately.
 - Click **View Details** to view the execution of the comparison task.
 - Click **Back to List** to return to the **Data Comparison** list.

View comparison results

- 1. Log in to NineData console.
- 2. In the left navigation bar, click **Comparison > Data**.
- 3. On the **Data Comparison** page, click the **Task ID**. The page description is as follows.



Serial number	Function	Description
1	Task status	Compare the current state of the task.
2	Configure alerts	After configuring alerts, the system notifies the user in case of a task failure.
3	Compare now	Only in-progress tasks are displayed, and a comparison task begins immediately. Option Description: • Method • Full Data: Perform a full data comparison. • Quick Comparison: Perform a quick comparison. • Inconsistent Recheck: Compare inconsistent data from the previous comparison. • Ignore Datatype: Ignore differences in data type and format between source and target, and only compare the valid numerical values.
4	More	 Pause: Pause the task, only applicable to the task whose status is Running. Edit: Edit the configuration information of the task. Terminate: End the task. Delete: Delete the task. The action is irreversible, please operate with caution.
5	Filter items	Filter the comparison results by task status , task results , and table names .
6	Log	Records all logs during the execution of the comparison task, and support quick filtering and positioning of target logs by log type, log generation time, keywords and other items.

Serial number	Function	Description
7	Monitor	Monitoring information of the task, showing the number of comparison records per second. For MySQL and SQL Server, a flow control function is also provided. After the flow control is enabled, when the thread_running parameter of the source database reaches the threshold you configured, the comparison will stop until the value of this parameter drops below the threshold before resuming the comparison, ensuring the stability of the data source.
8	Refresh	Get the latest status of comparison tasks.
9	List of comparison results	 Show the results of the data comparison between the source and target. Click the icon in the Actions column on the right side of the comparison list (displayed in case of inconsistency): View the comparison details of the table definition DDL statements on the source and target ends. Click the icon in the Actions column on the right side of the comparison list (displayed in case of inconsistency): generate change SQL, user can apply SQL on target instance for execution, and correct the inconsistent content.
10	History of comparison task results	Click to view a list of all comparison results. Click on a target list item to switch to the details of that comparison result.



Schema comparison

NineData database comparison function supports schema comparison, that is, metadata comparison, which is used to compare whether the definitions of objects in two databases are consistent. This article describes how to perform schema comparisons.

Prerequisite

- The data sources that need to be compared have been added to NineData.
- The two data sources that need to be compared for schema compatibility are of the same type.

Steps

- 1. Log in to NineData console.
- 2. In the left navigation bar, click **Comparison** > **Schema**.
- 3. In the upper right corner of the **Schema Comparison** page, click **Create Schema Comparison** .
- 4. On the Source & Target tab, configure according to the following table, and click Next.

Parameters	Description
Name	Enter the name of the comparison task. For the convenience of subsequent search and management, please try to use a meaningful name. Up to 64 characters are supported.
Source	Compare the source data source of the task, and the system will compare based on the structure of the data source.
Target	The object data source for the comparison task.
Гисанован	Colort how often the commonican is newformed

90 / 127

Parameters	Description	
	 One-time: Perform the comparison task only once. User can manually perform the comparison task multiple times. Periodic: Periodically perform a compare task. 	
Periodic (Periodic comparison only)	Choose how often to compare, user can choose any one or more days from Monday to Sunday, or check Every day to compare daily.	
Time Zone (Periodic comparison only)	Select the time zone according to the actual business situation. This parameter affects the time displayed in the comparison task.	
Launch Time (Periodic comparison only)	Select the time to launch the comparison task, from 00:00 to 23:59. User can also click Now to select the current time as the start time.	

- 5. On the **Objects** tab, select the content to be compared and click **Next** . Both **Full Instance** and Customized Object are supported.
- Full Instance: Select all databases and objects in the source instance.
- Customized Object : Customize the database and objects to be compared. In the Source list on the left, select the databases and objects to be compared, and click > add to Target list.
- 6. On the **Mapping** tab, configure the mapping relationship between the source and target databases and object names, and click Save and Pre-Check.



If the database and object names in the target end are inconsistent with those in the source end, user should configure the mapping logic.

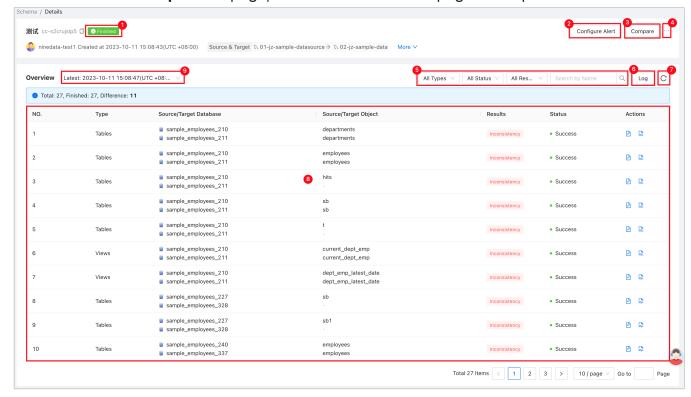
7. On the **Pre-check** tab, wait for the system to complete the pre-check. After the pre-check passes, click Launch.



- If the pre-check fails, user need to click the **Details** in the **Actions** column to the right of the target check item to correct the failure, and then perform the precheck again.
- Warnings can be repaired or ignored according to specific circumstances.
- 8. On the **Launch** tab, a message is displayed indicating that **Launch Successfully** . You can do the following:
 - If it is a **Periodic** comparison, click **Compare** to execute a comparison task immediately.
 - Click **View Details** to view the execution of the comparison task.
 - Click Back to List to return to the Schema Comparison list.

View comparison results

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click **Comparison > Schema**.
- 3. On the **Schema Comparison** page, click the **Task ID**. The page description is as follows.



Serial number	Function	Description
1	Task status	Compare the current state of the task.
2	Configure alerts	After configuring alerts, the system notifies the user in case of a task failure.
3	Compare now	Only in-progress tasks are displayed, and a comparison task begins immediately.
4	More	 Pause: Pause the task, only applicable to the task whose status is Running. Edit: Edit the configuration information of the task. Terminate: End the task. Delete: Delete the task. The action is irreversible, please operate with caution.
5	Filter items	Filter the comparison results by task status, task results, and table names.
6	Log	Records all logs during the execution of the comparison task, and supports quick filtering and positioning of target logs by log type, log generation time, keywords and other items.
7	Refresh	Get the latest status of comparison tasks.
8	List of comparison results	Show the results of the structural comparison between the source and destination. Click the in the Actions column on the right side of the comparison list: view the comparison details of the table definition DDL statements of the source and target.

Serial number	Function	Description
		 Click the in the Actions column on the right side of the comparison list (displayed in case of inconsistency): generate change SQL, user can apply SQL on target instance for execution, and correct the inconsistent content.
9	History of comparison task results	Click to view a list of all comparison results. Click on a target list item to switch to the details of that comparison result.



Adding Sybase Data Source (Standalone Architecture)

This section describes how to add a Sybase data source with a standalone architecture to NineData, for data replication from Sybase to PostgreSQL.

Prerequisites

The Sybase data source is deployed in a standalone architecture.

Steps

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click **Datasource** > **Datasource**.
- 3. Click the **Datasource**tab, and then click **Create Datasource** on the page. In the pop-up window for data source type, select **Database** > **Sybase**. Configure the parameters on the **Create Datasource** page according to the table below.



TIP

If you make a mistake, you can click the !arrow_down icon at the top of the **Create Datasource** page to reselect.

Parameter	Description
Name	Enter the name of the data source. To facilitate subsequent search and management, please use meaningful names as much as possible.
Connection	Select the access method for the data source. Here, select Private Network.

Parameter	Description
Architecture Type	Select the architecture type of the data source. Here, select Standalone .
Host	Enter the internal network connection address and port for the Sybase data source.
DB Account	Enter the login username for the database.
DB Password	Enter the login password for the database.
Character	Enter the target character set. NineData will automatically convert the characters in the database to this character set for reading, such as big5hk, utf8, to avoid garbled text.
Access Region	Fixed as NineData Dedicated Cluster .
Environment	Choose based on the actual business use of the data source, as an environmental identifier for the data source. The default environments provided are PROD and DEV , and you are also supported to create custom environments.

4. After completing all configurations, click on **Connection Test** to test if the data source can be accessed normally. If prompted with **Connection Successfully**, click on **Create Datasource** to finish adding the data source. Otherwise, please recheck the connection settings until the connection test is successful.

Appendix I: Creating Environments

To meet different enterprise needs, NineData not only provides default environments such as **PROD** and **DEV**, but also allows you to manually create custom environments.

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click on **Datasource** > **Datasource**.

- 3. Click on the **Environment** tab, which displays all environments under the current organization. Here, you can **Edit** or **Delete** the custom-created environments.
- 4. Click on **Create Envir.** in the upper right corner of the page.
- 5. Complete the configuration according to the table below and click **OK**.

Parameter	Description
Name	Enter the name of the environment, which can contain 1 to 16 characters. This name serves as the identifier of the environment and needs to be distinctive. Please use a meaningful name if possible.
Label Color	Select the color in which the environment tag is displayed on the console.

Appendix II: Adding Multi-active Labels to Data Sources

If your data source needs to be used to create multi-node replication tasks, to prevent circular replication, you need to set multi-active labels for all data sources participating in the replication, and each data source's multi-active label must be globally unique.

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click on **Datasource** > **Datasource**.
- 3. Click on the target data source ID to open the **Details** page.
- 4. In the data source details area (which includes information such as data source name, ID, creator, creation time, etc.), click on **More**.
- 5. Find **Multi-Active Tag**, and click on the \square icon to the right of it.
- 6. Enter the multi-active label and click **OK**.
 - Multi-active labels can contain 1 to 64 characters.
 - Multi-active labels must be globally unique and cannot be repeated with other multiactive labels.

Basic service

Adding Sybase Data Source (Master-Slave **Architecture**)

This section describes how to add a Sybase data source with a master-slave architecture to NineData for data replication from Sybase to PostgreSQL.

Prerequisites

The Sybase data source is deployed in a master-slave architecture.

Procedure

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click on **Datasource** > **Datasource**.
- 3. Click on the **Datasource** tab, and then click on **Create Datasource** on the page. In the popup dialog for data source types, select **Database** > **Sybase**, and configure the parameters on the Create Datasource page according to the table below.



()TIP

If you make a mistake, you can click the \vee icon at the top of the **Create Datasource** page to make a new selection.

Parameter	Description
Name	Enter the name of the data source. For ease of future reference and management, try to use a meaningful name.
Connection	Select the access mode of the data source. Here, select Private Network.

Parameter	Description
Architecture Type	Select the architecture type of the data source. Here, select Master Slave .
Host	Enter the internal connection address and port of the Sybase primary server.
Standby Node	Enter the internal connection address and port of the Sybase MSA replication node.
DB Account	Enter the login username of the database.
DB Password	Enter the login password of the database.
Character	Enter the target character set. NineData will automatically convert the characters in the database to this character set for reading, such as big5hk, utf8, to avoid garbled text.
Access Region	Fixed as NineData Dedicated Cluster .
Environment	Select based on the actual business purpose of this data source as an environment identifier for the data source. Default options include PROD and DEV , and you can also create custom environments.

4. After completing all configurations, click on **Connection Test** to test if the data source can be accessed normally. If prompted with **Connection Successfully**, click on **Create Datasource** to finish adding the data source. Otherwise, please recheck the connection settings until the connection test is successful.

Appendix I: Creating Environments

To meet different enterprise needs, NineData not only provides default environments such as **PROD** and **DEV**, but also allows you to manually create custom environments.

1. Log in to the NineData console.

- 2. In the left navigation bar, click on **Datasource** > **Datasource**.
- 3. Click on the **Environment** tab, which displays all environments under the current organization. Here, you can **Edit** or **Delete** the custom-created environments.
- 4. Click on **Create Envir.** in the upper right corner of the page.
- 5. Complete the configuration according to the table below and click **OK**.

Parameter	Description
Name	Enter the name of the environment, which can contain 1 to 16 characters. This name serves as the identifier of the environment and needs to be distinctive. Please use a meaningful name if possible.
Label Color	Select the color in which the environment tag is displayed on the console.

Appendix II: Adding Multi-active Labels to Data Sources

If your data source needs to be used to create multi-node replication tasks, to prevent circular replication, you need to set multi-active labels for all data sources participating in the replication, and each data source's multi-active label must be globally unique.

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click on **Datasource** > **Datasource**.
- 3. Click on the target data source ID to open the **Details** page.
- 4. In the data source details area (which includes information such as data source name, ID, creator, creation time, etc.), click on **More**.
- 5. Find **Multi-Active Tag**, and click on the \square icon to the right of it.
- 6. Enter the multi-active label and click **OK**.
 - Multi-active labels can contain 1 to 64 characters.
 - Multi-active labels must be globally unique and cannot be repeated with other multiactive labels.



Basic service

Adding PostgreSQL Data Source (Primary-Secondary Architecture)

This section describes how to add a PostgreSQL data source with a primary-secondary architecture to NineData.

Prerequisites

The PostgreSQL data source is deployed in a primary-secondary architecture.

Procedure

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click on **Datasource** > **Datasource**.
- 3. Click on the **Datasource** tab, and then click on **Create Datasource** on the page. In the popup dialog for data source types, select **Database** > **PostgreSQL**, and configure the parameters on the **Create Datasource** page according to the table below.



TIP

If you make a mistake, you can click the \lor icon at the top of the **Create Datasource** page to make a new selection.

Parameter	Description
Name	Enter the name of the data source. For ease of future reference and management, try to use a meaningful name.
Connection	Select the access mode of the data source. Here, select Private Network.

Parameter	Description
Host	Enter the private IP (Virtual IP) address and port of PostgreSQL to support high-availability automatic switching between primary and secondary nodes.
DB Account	Enter the login username of the database.
DB Password	Enter the login password of the database.
Access Region	Fixed as NineData Dedicated Cluster .
Environment	Select based on the actual business purpose of this data source as an environment identifier for the data source. Default options include PROD and DEV , and you can also create custom environments.

4. After completing all configurations, click on **Connection Test** to test if the data source can be accessed normally. If prompted with **Connection Successfully**, click on **Create Datasource** to finish adding the data source. Otherwise, please recheck the connection settings until the connection test is successful.

Appendix I: Creating Environments

To meet different enterprise needs, NineData not only provides default environments such as **PROD** and **DEV**, but also allows you to manually create custom environments.

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click on **Datasource** > **Datasource**.
- 3. Click on the **Environment** tab, which displays all environments under the current organization. Here, you can **Edit** or **Delete** the custom-created environments.
- 4. Click on **Create Envir.** in the upper right corner of the page.
- 5. Complete the configuration according to the table below and click \mathbf{OK} .

Parameter	Description
Name	Enter the name of the environment, which can contain 1 to 16 characters. This name serves as the identifier of the environment and needs to be distinctive. Please use a meaningful name if possible.
Label Color	Select the color in which the environment tag is displayed on the console.

Appendix II: Adding Multi-active Labels to Data Sources

If your data source needs to be used to create multi-node replication tasks, to prevent circular replication, you need to set multi-active labels for all data sources participating in the replication, and each data source's multi-active label must be globally unique.

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click on **Datasource** > **Datasource**.
- 3. Click on the target data source ID to open the **Details** page.
- 4. In the data source details area (which includes information such as data source name, ID, creator, creation time, etc.), click on **More**.
- 5. Find **Multi-Active Tag**, and click on the \square icon to the right of it.
- 6. Enter the multi-active label and click **OK**.
 - Multi-active labels can contain 1 to 64 characters.
 - Multi-active labels must be globally unique and cannot be repeated with other multiactive labels.

Basic service

Adding PostgreSQL Data Source (Single-node Architecture)

This section describes how to add a PostgreSQL data source with a single-node architecture to NineData.

Prerequisites

The PostgreSQL data source is deployed in a single-node architecture.

Procedure

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click on **Datasource** > **Datasource**.
- 3. Click on the Datasource tab, and then click on Create Datasource on the page. In the popup dialog for data source types, select Database > PostgreSQL, and configure the parameters on the Create Datasource page according to the table below.



TIP

If you make a mistake, you can click the \lor icon at the top of the **Create Datasource** page to make a new selection.

Parameter	Description
Name	Enter the name of the data source. For ease of future reference and management, try to use a meaningful name.
Connection	Select the access mode of the data source. Here, select Private Network.

Parameter	Description
Host	Enter the internal connection address and port of the PostgreSQL data source.
DB Account	Enter the login username of the database.
DB Password	Enter the login password of the database.
Access Region	Fixed as NineData Dedicated Cluster .
Environment	Select based on the actual business purpose of this data source as an environment identifier for the data source. Default options include PROD and DEV , and you can also create custom environments.

4. After completing all configurations, click on Connection Test to test if the data source can be accessed normally. If prompted with Connection Successfully, click on Create Datasource to finish adding the data source. Otherwise, please recheck the connection settings until the connection test is successful.

Appendix I: Creating Environments

To meet different enterprise needs, NineData not only provides default environments such as **PROD** and **DEV**, but also allows you to manually create custom environments.

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click on **Datasource** > **Datasource**.
- 3. Click on the **Environment** tab, which displays all environments under the current organization. Here, you can **Edit** or **Delete** the custom-created environments.
- 4. Click on **Create Envir.** in the upper right corner of the page.
- 5. Complete the configuration according to the table below and click **OK**.

Parameter	Description
Name	Enter the name of the environment, which can contain 1 to 16 characters. This name serves as the identifier of the environment and needs to be distinctive. Please use a meaningful name if possible.
Label Color	Select the color in which the environment tag is displayed on the console.

Appendix II: Adding Multi-active Labels to Data Sources

If your data source needs to be used to create multi-node replication tasks, to prevent circular replication, you need to set multi-active labels for all data sources participating in the replication, and each data source's multi-active label must be globally unique.

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click on **Datasource** > **Datasource**.
- 3. Click on the target data source ID to open the **Details** page.
- 4. In the data source details area (which includes information such as data source name, ID, creator, creation time, etc.), click on **More**.
- 5. Find **Multi-Active Tag**, and click on the \square icon to the right of it.
- 6. Enter the multi-active label and click **OK**.
 - Multi-active labels can contain 1 to 64 characters.
 - Multi-active labels must be globally unique and cannot be repeated with other multiactive labels.

Account management

Introduction to account management

NineData provides multi-person collaboration capabilities, which can be enabled after creating or joining an organization. This article introduces concepts such as organizations, users, roles, etc.

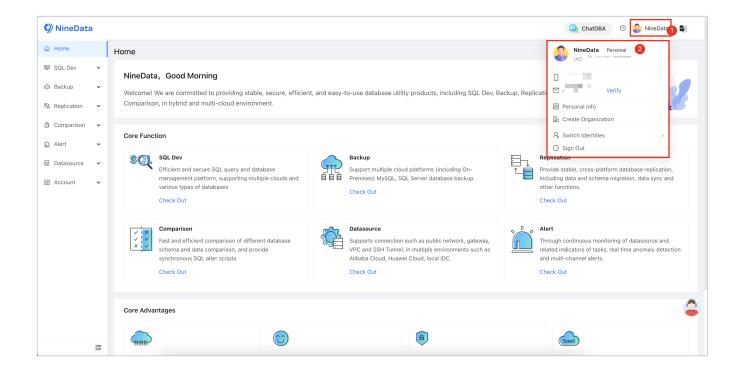
Organization

Typically, the system provides you with a default organization and automatically switches to it. You can manually switch to **Personal Mode**, which is suitable for individual developers, for tasks such as backing up, restoring, synchronizing, comparing their own databases, and performing daily development tasks using the SQL window.

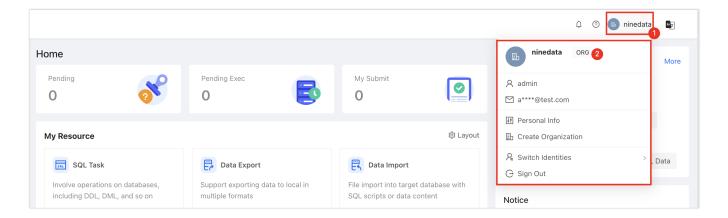
However, for enterprise users, there is a need for multi-person collaborative development. In this scenario, various roles are responsible for different businesses, so each role requires different permissions to meet data security needs. NineData's **Organization Mode** allows different roles and resource management permissions to be assigned to each user under the organization, facilitating multi-user collaborative development within the same organization. This ensures data security while enhancing overall productivity.

You can hover over the icon in the upper right corner of the page to see the current mode.

Personal Mode



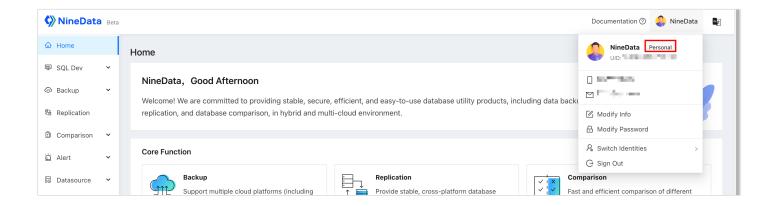
Organization Mode



User

A user is an independent entity in NineData. In Personal Mode, the user is yourself, with full permissions to all resources. In Organization Mode, the user is part of an organization, with the role determined by the system administrator and corresponding resource permissions assigned.

You can hover over the icon in the upper right corner of the page to see the current mode.



Role

Roles determine a user's permissions to use resources. NineData provides the following 5 preset roles:

- System Administrator: The user who creates the organization is by default the system
 administrator, with full permissions to data sources within the organization, management
 permissions for organization users and roles, access to audit logs (user behavior records),
 and permissions for all functional modules.
- Production Environment Administrator: By default, has full permissions to production environment data sources and functional modules.
- Production Environment Read-Only: By default, has read-only permissions to production environment data sources and full permissions to all functional modules.
- Development Environment Administrator: By default, has full permissions to development environment data sources and functional modules.
- Regular Member: Has full permissions to all functional modules.

System administrators can delete or edit all preset roles except for **System Administrator**, and can also add custom roles.

Permissions

Permissions in NineData are divided into four categories:

- Management Permissions: Manage users and roles within the current organization, accessible only to Administrator.
- Service Permission: Determines whether users can see the six major modules on the left side of the console, including DevOps, Backup, Replication, Comparison, Datasource,

Alert, subdivided into sub-functional entries within each module, such as SQL Console.

- **Datasource Permission**: Determines whether users can perform operations such as editing, SQL querying or modification, backup and recovery, data replication, and database comparison on target data sources. **Datasource Permission** needs to be used in conjunction with **Service Permission**. For example, a user has permission to replicate data for a certain data source, but has not been granted permission for the replication module, they still cannot use the data replication feature.
- Other Permission: Includes Alert Permission and Sensitive Data Management.
 - **Alert Permission**: Grants permission to configure alarms.
 - **Sensitive Data Management**: Grants permission to configure sensitive data sources.



Managing Organizations

For enterprise users, NineData's organization feature facilitates effective collaboration among multiple users by managing each user's permissions through roles. This article explains how to manage organizations.

Limitation

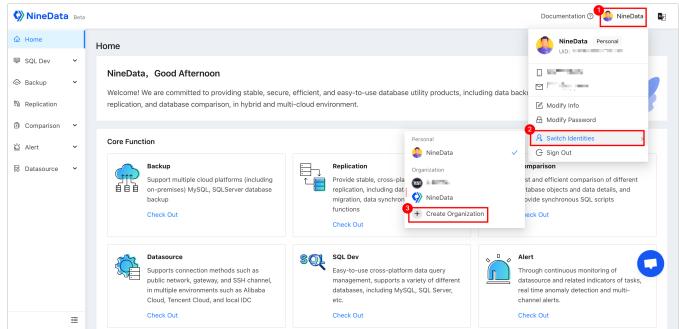
You can create or join a maximum of 2 organizations.

Creating an Organization

The system has already created a default organization for you, which you can use directly. If you need to create a new organization, follow this process.

Steps

- 1. Log in to the NineData console.
- 2. In the upper right corner of the NineData console page, hover over the user avatar, click **Switch Identities** and then **Create Organization**.



3. On the **Create Organization** page, configure the parameters according to the table below.

Parameter	Description
Name	Enter the organization's name, typically using the company name for easy identification. Maximum of 24 characters.

4. Click **OK** to complete the organization creation.

Switching to an Organization

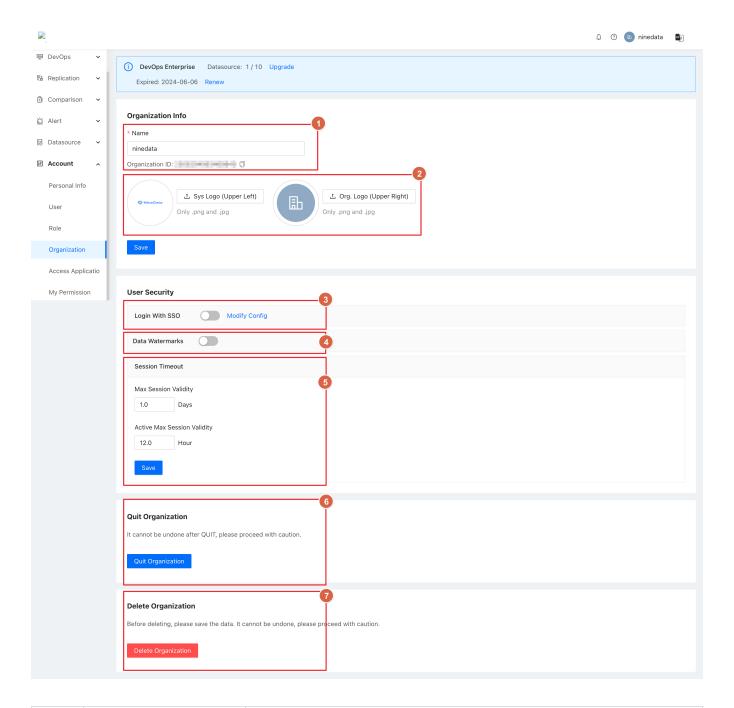
- 1. Log in to the NineData console.
- 2. In the upper right corner of the NineData console page, hover over the user avatar, click **Switch Identities**, and select the organization you created under **Organization**.

Viewing or Editing Organization Information

Prerequisites

To edit organization information, the user role must be **Administrator**.

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click **Account Management** > **Organization Information** to view the current organization's information.
- 3. On the **Organization Information** page, you can perform the following operations.



No.	Parameter	Description
1	Modify Organization Name (admin only)	Modify the organization name. Maximum of 24 characters. After modification, click Save Changes .
2	Change Logo (admin only)	Change the System Logo in the upper left corner of the page and the Organization Logo in the upper right corner

No.	Parameter	Description
		of the page. Click the corresponding button to upload the logo file, and then click Save Changes .
3	Configure SSO Login (admin only)	Enable or disable SSO login for the organization.
4	Manage Data Watermark (admin only)	Enable or disable data watermark for the organization.
5	Configure Session Timeout (admin only)	Configure the expiration time after logging in to the console. After expiration, users need to log in to the console again. • Max Session Validity: Default is 24 hours. After this time, users must log in again. • Active Max Session Validity: Default is 12 hours. If users do not perform any operations within this time, they must log in again.
6	Exit Organization	Exit the current organization. If you are the only system administrator in the organization, you need to assign the system administrator role to other members of the organization before exiting.
7	Delete Organization (admin only)	Delete the current organization. Warning : After deletion, all data sources, members, etc., in the organization will be lost, and the operation cannot be undone. Please proceed with caution.



Creating Users

Users are a fundamental part of an organization. If you are a system administrator, you can create new users within your organization.

Prerequisites

Your role is **Administrator**.

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click **Account > User**.
- 3. On the **User** page, click the **Invite** button in the upper right corner.
- 4. In the **Invite** dialog box, configure according to the table below.

Parameter	Description
User Type	 Supports new user and existing user: New User: A brand-new user who has never been added to the NineData system. Existing User: A user who has been added before but has been removed.
Login Name	The login name of the user, supports Chinese and English characters, and special characters: o For new user : Enter a username that has never been added before. o For existing user : Enter a username that has been added previously.
Role	Select the user's role. The role determines the user's permissions. NineData has built-in 5 types of roles to choose from, please refer to the

Parameter	Description	
	console for details.	
Add	Click Add to create users in bulk.	

5. Click **OK** to complete the creation of the user. If you are creating a **new user**, the page will display **Login name**, **Password**, and **Login Link** information. If it is an **existing user**, only the **Login Link** will be displayed. Click to copy this information to the clipboard for easy sharing with the user. Users can log in to NineData using this information.



Editing Users

If you are a system administrator, you can edit target users. Editable items include **Login Name**, **Username**, **Query per day Max Times**, **Query per day Max Rows**.

- Login Name: Change the login account name of the target user when logging into NineData.
- Username: Change the display name of the user within the NineData organization.
- Query per day Max Times (For database DevOps use only): Configure the number of queries a current user can execute per day, default is 2000 times. The date on the right is the deadline, meaning before 0 o'clock on that date, the user can perform 2000 queries every day, if left blank, it is effective indefinitely.
- **Query per day Max Rows** (For database DevOps use only): Configure the total number of rows a current user can query per day, default is 20000 rows. The date on the right is the deadline, meaning before 0 o'clock on that date, the user can query 20000 rows of data every day, if left blank, it is effective indefinitely.

Prerequisites

Your role is **Administrator**.

Steps

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click **Account** > **User** to view all users in the current organization.
- 3. On the **User** page's user list, find the user you need to edit and click **Actions** in the **Edit** column to the right of the user.



TIP

You can filter by role category in the upper right corner of the page, or directly enter a username to search and quickly locate the target user.

4. Configure the user in the pop-up window, then click **OK**.



Removing Users

To prevent data security incidents, users who are no longer part of a project can be removed. Once removed, the user will no longer be able to log in to the NineData platform.

Prerequisites

Your role is **Administrator**.

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click **Account > User**.
- 3. On the **User** page, find the user you wish to remove and click **Actions** under the **Remove** column to the right.
- 4. In the confirmation window, click **Remove** again to proceed. After this, the user will no longer be able to log in to NineData with that username. If you need to reverse this action, you can re-add the user by adding an **existing user**. For specific methods, please refer to the **Creating Users** section.

Account management

Configuring User Permissions

Prerequisites

Your role must be **Administrator**.

Steps

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click **Account** > **User** to view all users in the current organization.
- 3. In the user list on the **User** page, click the **Authority** in the **Actions** column on the right side of the target user. In the popup window of **Authority**, configure the user permissions.

The **Authority** consists of three sections: **Application**, **Role Grant**, and **User Grant**.

- Application: Displays all permissions applied by the user through submitting
 permission application forms, divided into three sections: Datasource, Database, and
 Sensitive Column. You can revoke the permissions applied by the user by clicking the
 Revoke in the Actions column on the right side of the target permission.
- Role Grant: Displays the roles that the user is currently bound to and the corresponding permissions. You can click Edit Grant to bind new roles to the user or unbind existing roles.
- User Grant: Displays the personal permissions granted to the user separately. Personal
 permissions complement other permissions the user has (Application, Role Grant) and
 there is no priority among them. You can grant or revoke personal permissions for the
 user by toggling the slider on the right side of the target permission.
- 4. Click Save.



Managing Roles

Roles determine the permissions users have for accessing resources. NineData provides 5 predefined roles. You can edit any role except for **Administrator** and also create custom roles. Roles that are no longer needed can be manually deleted.

Prerequisites

You must have the role of **Administrator**.

Notes

- The role named **Administrator** in the predefined roles cannot be edited or deleted.
- When creating or editing role names, you cannot use role names already existing under your current NineData account.

Viewing or Adding Roles

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click **Account** > **Role**. You can choose to rename role names or edit role permissions.
- 3. Click the **Create Role** button at the bottom left of the page.
- 4. In the **Create Role** dialog, enter the name of the new role. This name is used to identify the role's purpose, so use a meaningful name, with a maximum of 16 characters.
- 5. Click **OK** to complete the creation of the new role. The page will automatically redirect to the permission configuration page for the newly created role, where you can configure the role's permissions according to your actual business needs.

Editing Existing Roles

You can edit any role except for **Administrator**, including the role name and specific permissions.

1. Log in to the NineData console.

- 2. In the left navigation bar, click **Account** > **Role**. You can choose to rename role names or edit role permissions.
 - Renaming a Role: In the list of roles on the left side of the page, hover over the role you want to rename, click the ☑ icon on the right. In the **Edit Role** dialog that appears, enter the new role name and click OK. Role names can have a maximum of 16 characters.
 - The \square icon will not appear next to the **Administrator** role because it cannot be modified.
 - Editing Role Permissions: Click the name of the target role in the list of roles on the left side of the page. In the permissions configuration page on the right side, edit the permissions for that role. After editing, click **Save** at the bottom left of the page.

Deleting Roles

You can delete any role except for **Administrator**.

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click **Account** > **Role**. You can choose to rename role names or edit role permissions.
- 3. In the list of roles on the left side of the page, hover over the role you want to delete, and click the \square icon on the right.
 - The \square icon will not appear next to the **Administrator** role because it cannot be deleted.
- 4. In the **Edit Role** dialog that appears, click the **Delete** button at the bottom left.
 - After deleting a role, users assigned to that role will lose the role and all associated permissions. This action cannot be undone, so proceed with caution.
- 5. In the confirmation dialog that appears, enter the name of the role and click the **Delete** button at the bottom right.

Binding Roles

Roles only take effect after being bound to target users. Before binding, ensure that the target users have joined the organization.

- 1. Log in to the NineData console.
- 2. In the left navigation bar, click **Account** > **Role**. You can choose to rename role names or edit role permissions.
- 3. Click the name of the target role in the list of roles on the left side of the page. In the **Members** on the right side of the page, click **Add**.

You can view users already bound to the current role in the **Members**. If there are too many users, you can enter keywords in the search box to search.

4. In the **Add** dialog that appears, click the checkbox below **Members** and select the target users from the list. You can select multiple users for batch binding.

Users already bound to the current role will not appear in the list. If you cannot find the target user, make sure that the user has joined the organization and is not already bound to the current role.

5. Click **OK**.



View audit logs

NineData supports the audit log feature. By viewing the audit log, system administrators can track the operation history that has occurred in the console, in order to analyze the changes made by users and the occurrence of security incidents.

Function Description

The audit log feature of NineData mainly records who did what to which object at when. Specifically, the audit log usually includes the following information:

- Who: User ID, organization, IP address, etc.
- When: Timestamp, records the time when the operation occurred.
- **Object**: The object of the operation, usually a data resource (data source, database, table, column, etc.), user, task (backup, replication, comparison, etc.).
- What: The operation performed, which basically covers all operations performed on the NineData console.

The audit log records the user's operations based on the above information and displays the details. Depending on the operation object, the audit log will display specific content.

Precautions

The following operations will not be recorded:

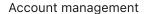
- Viewing operations for non-sensitive information, such as: a user viewed the list of replication tasks.
- Personal-related operations, such as: a user saved SQL, viewed SQL operations in the SQL window.
- Database operations triggered by specific tasks, such as: a data replication task initiated a
 write operation to the target data source.

Prerequisite

You must have the role of **Administrator**.

Procedure

- 1. Log in to the NineData Console.
- 2. In the left navigation pane, click **Account**, then click **Audit Log**.
- 3. On the audit log page, you can switch between tabs to view Operation Logs or SQL Logs.
 - Operation Logs: Records the actions performed by users in the NineData Console. You can use the filter function to quickly find the target operation record. The filterable items include User, Module, Type, Time, and Event Object. After configuring the filter items, click Search to quickly locate the required operation record. If the recorded Event Object is a specific instance ID, you can directly click to jump to it.
 - SQL Logs: Records the modifications made by users to data sources in the form of SQL statements. You can use the filter function to quickly find the target SQL execution record. The filterable items include User, Module, Type, Datasource, Database, Time, and Query (SQL statement). After configuring the filter items, click Search to quickly locate the required SQL execution record.
- 4. Click **Details** on the right side of the target record to view all information of that record.



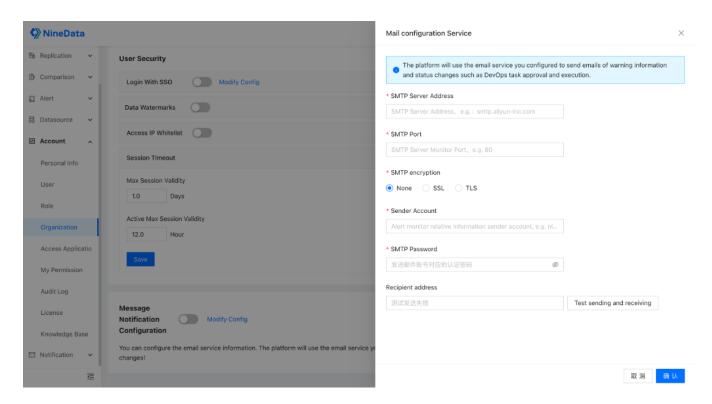
Configure Message Notifications

NineData supports message notification functionality. By configuring an internal SMTP email service, users can receive NineData alert messages, approval workflow messages, and more.

Prerequisites

Your role must be **System Administrator**.

- 1. Log in to the NineData console with an administrator account.
- 2. In the left navigation bar, click **Account > Organization**.
- 3. At the bottom of this page, you'll find a Message Notification Configuration section. If you need email notifications, enable the feature toggle and click Modify Config. The following page will appear:



Enter your SMTP email service details, click **Test Sending and Receiving**. If you receive the test email, the SMTP information is correct. Click **OK** to save.

4. Once configured, you can update your email address in **Account > Personal Info**.