



WHITEPAPER

PhoenixAI Cloud

Security Whitepaper

ABOUT THIS WHITEPAPER

This whitepaper aims to provide you with a comprehensive understanding of the security measures implemented within the PhoenixAI Cloud platform and equip you with the necessary knowledge to ensure the confidentiality, integrity, and availability of your data.

WHAT'S INSIDE

1. Introduction
2. Architecture
3. Accounts
4. Authentication
5. Authorization
6. Auditing
7. Data Security

Introduction

PhoenixAI Cloud, a powerful and secure cloud data platform, provides enhanced security capabilities. This whitepaper aims to provide you with a comprehensive understanding of the security measures implemented within the PhoenixAI Cloud platform and equip you with the necessary knowledge to ensure the confidentiality, integrity, and availability of your data.

PhoenixAI Cloud is a Bring Your Own Cloud (BYOC) service that enables organizations to leverage the power and flexibility of PhoenixAI, a high-performance analytics platform, in their preferred cloud environment. With PhoenixAI, you can harness these capabilities while maintaining control over your cloud infrastructure.

These principles were key in building the PhoenixAI Cloud service, and continue to be adhered to as new features are added and in how the service is managed:

1. **Data Sovereignty:** The data plane resides within the customer's Virtual Private Cloud (VPC). PhoenixAI's control plane handles cluster management, while user data remains under the customer's control.
2. **Principle of Least Privilege(PoLP):** This principle is consistently applied:
 - a. Minimal IAM permissions are set during PhoenixAI deployment
 - b. Access is controlled by the customer during troubleshooting
 - c. RBAC is used for all operations to minimize unnecessary exposure.
3. **Data Encryption and Compliance:** PhoenixAI Cloud prioritizes data security with encryption for transmission and storage. Holding SOC 2 and GDPR certifications showcases our commitment to rigorous security and compliance standards, ensuring data protection and regulatory adherence.

In the following sections, we will delve into the architecture of PhoenixAI Cloud, highlighting the underlying components and their roles in ensuring the security of your data. We will then explore account management, authentication methods such as LDAP and Okta, privilege systems such as the RBAC model, and external privilege system integration.


Architecture

PhoenixAI Cloud consists of two main components: the Control plane and the Data plane.

- **Control plane:** The Control plane serves as the central management hub for PhoenixAI clusters. It encompasses cluster deployment, monitoring, scaling, billing, optimization, WebGUI, and more. The Control plane is deployed within PhoenixAI's VPC.
- **Data plane:** The Data plane includes the BE (Backend Engine) and FE (Frontend Engine) of the PhoenixAI cluster. The BE handles data storage and SQL execution, while the FE manages metadata storage and SQL plan generation and optimization. The Data plane is deployed within the customer's

VPC. PhoenixAI clusters support two modes: Classic cluster and Elastic cluster, corresponding to the shared-nothing and shared-data models.

- **Communication** between the two planes occurs via a secure network connection (e.g., AWS Private Link). PhoenixAI uses gRPC over TLS connections for both public access and private link service.



An example is worth one thousand words

Reading through this short example will provide you with a better understanding of where your data resides, who has access to it, and how that access is controlled. Cindy is a data engineer working on a new dataset. This is her workflow:

Create a PhoenixAI cluster

This is a new project with a new dataset, so Cindy decides to create a new PhoenixAI cluster to work in. Cluster creation is initiated in the PhoenixAI Cloud Control plane.

- Cindy logs in to PhoenixAI Cloud with her company-provided Okta SSO credentials.
- Cindy creates a new cluster. After selecting the instance types, there is a choice to deploy manually or use a quick deployment, and whether to deploy in an existing VPC or a new one.

Whichever choice is made, the destination for the cluster is in an AWS account that she has administrative rights to. The cluster is in the Data plane. Some important distinctions to call out here that are directly related to security and privacy are:

- The Control plane and Data plane are separate
- The Control plane is in a PhoenixAI VPC and is an administrative UI
- The Data plane is in your VPC and is where your data is stored; only your users and people authorized by you have access to your Data plane
- You can create a VPC that meets your company security standards and deploy the Data plane there
- Authentication to the Control plane and Data plane are separate. Separate accounts, and separate authentication methods are used.
- Cindy hands the Cloudformation details to her colleague who has the AWS privileges to create the cluster and the cluster is created. The initial password for the database is set during this process, and ongoing authentication will be set up afterward to comply with the rules set by Cindy's company.
- Once the Cloudformation job finishes Cindy clicks the Finish button in the Control plane and the cluster status is available in the Control plane.

Work with the PhoenixAI cluster

The cluster connection details, other than the password as that is set by you during the Cloudformation step, are available in the Control plane on the specific cluster's overview page.

After collecting the hostname and port from the overview page in the Control plane, Cindy leaves the Control plane and from now on is working in the Data plane, which is hosted in a VPC in her company AWS account. Her first priority is to integrate the PhoenixAI cluster with her company LDAP server so that she can collaborate with her team.

Once logged in as the database admin Cindy configures the authentication policy for the cluster to use her company LDAP service.

Work with your data

Now that the cluster is created and LDAP is configured Cindy and her colleagues can get to work. The data is stored in a VPC owned by Cindy's company, and configured to the company data access standards.

Collaborating with PhoenixAI support

Sometimes you need a hand. PhoenixAI support engineers do not have access to your Data plane, so they cannot see your queries, data, logs, etc. At some point, Cindy and her team may need help. They can request that PhoenixAI support connect to their database by contacting support and inviting the support engineer to join the organization. After the session, the support engineer is removed from the organization.

Data access

Data can be categorized into three levels:

- **Level 0:** User's data and metadata. All this information is stored within the customer's Data plane, and PhoenixAI cannot access it unless granted explicit permission by users. Stored data and metadata are encrypted.
- **Level 1:** User query history, profiles, and audit logs. This data might contain sensitive information. PhoenixAI can access this data in the Data plane with user authorization to assist in diagnosing slow queries and optimizing performance.
- **Level 2:** PhoenixAI logs and machine monitoring data unrelated to user-specific information. PhoenixAI has default access to this data for monitoring and ensuring customer SLAs.

PhoenixAI's Control plane does not store any of the above data.

For specific data storage and access methods, refer to the diagram below:

Data Flow Diagram - PhoenixAI Cloud 7.2023

last updated: 07.06.2023

PhoenixAI VPC

monitoring metric (e.g. qps) - 14 days persisted EBS (AWS), NodeStorage(Azure)

service logs (like request ID) for PhoenixAI Cloud mysql RDS and web audit log like "scale out/in"
no data in these logs

query request (not persisted anywhere)

Query Answer - if GUI used to request. Not persisted anywhere, just shown to GUI

sent through TLS to GUI and not persisted




"save query request from browser" saves result data to file

Control plane

Data plane

cross-domain request "cors", 10-minute expiration time, directly to user-browser display; no storage

Customer VPC



Accounts

PhoenixAI Cloud includes both a data plane and a control plane, resulting in two separate account systems.

Control plane


The control plane is accessed with a web browser. Within an account, the following actions can be performed:

- Cluster Management
 - View Cluster details (including node inspection/cluster status monitoring)
 - Creation/deletion of clusters
 - Configuration changes
 - View alerts

- Manage AWS Authorization Information
- API Management
- User Management (addition, deletion, modification, and retrieval of the PhoenixAI cloud user under the account)
- View Usage Consumption Statistics
- Billing-related
 - Historical bill viewing
 - Current bill viewing
 - Settlement

Data plane

The data plane, on the other hand, is within the user's VPC and access is through database user accounts using the MySQL protocol supporting LDAP as a third-party authentication solution.



Authentication

To ensure the security isolation of data, the user management system of PhoenixAI Cloud is divided into two sets—one for the control plane and one for the data plane. Therefore, authentication also consists of two independent systems. In the Control plane, users can log in using their PhoenixAI accounts or integrate with third-party Single Sign-On (SSO) solutions like Okta. In the Data plane, login occurs through the MySQL protocol, supporting LDAP as a third-party authentication solution.

Web manager authentication

OKTA

SSO (Single Sign-On) enhances user experience and security by allowing access to multiple applications with a single set of credentials. Okta offers robust SSO solutions for seamless and secure authentication. PhoenixAI SSO support enables unified access control through integration with Okta, streamlining user management and bolstering data security.

Cluster authentication

MySQL authentication

PhoenixAI utilizes the MySQL protocol for compatibility, allowing direct connections through MySQL clients. PhoenixAI ensures secure connectivity via HTTPS. Alternatively, you can establish a direct connection using the internal FE 9030 port, ensuring data bypasses PhoenixAI's network entirely.

LDAP

LDAP (Lightweight Directory Access Protocol) is a protocol used for managing and accessing distributed directory services. It finds extensive use in scenarios requiring centralized user authentication and information storage, such as corporate internal email systems and other databases. Its mechanism involves hierarchical data organization, allowing efficient data retrieval and modification across various networked devices and platforms.


PhoenixAI integrates LDAP authentication to enhance security and user management. LDAP (Lightweight Directory Access Protocol) is utilized as an external authentication source. Users' credentials and roles are stored in the LDAP directory. When a user tries to access PhoenixAI, the system communicates with the LDAP server to verify the provided credentials. Upon successful authentication, the user's roles and permissions are fetched from LDAP. This integration streamlines user management, promotes consistency across systems, and enables centralized control over access rights in PhoenixAI, bolstering overall data security.

PhoenixAI Cloud enhances the LDAP functionality within PhoenixAI through the following feature updates:

1. **Direct LDAP Authentication:** Users can now directly log in using their LDAP credentials, eliminating the need for manual user creation in PhoenixAI.

2. **Automated Role Assignment:** Users' roles are automatically determined based on their LDAP groups, alleviating the need for manual role assignments.


The following diagram shows the flow of interaction between PhoenixAI and LDAP:



By integrating this feature, organizations that wish to centrally manage user information via LDAP and match roles with LDAP groups can seamlessly manage user authentication and authorization within the PhoenixAI Cloud platform. The implementation involves configuring connectivity, specifying authentication order, creating role mappings, and refreshing the cache to ensure synchronization between LDAP and PhoenixAI.

This advancement streamlines user management, optimizes security, and enhances the overall experience of LDAP integration within the PhoenixAI Cloud data ecosystem.

Authorization



Web manager access control


In the PhoenixAI Cloud web manager, you can create new roles and manage existing roles:

Role management				Create role
Name	Description	Create at	Actions	
Account admin	Account admin is a built-in role that has all privileges enabled in this Phoeni...	12/5/2022, 11:39:51 PM		
CloudNative		12/28/2022, 11:25:50 PM	Remove	
cluster_manager	Only member with this role can create/edit/drop cluster. Just for saving more cost!	3/18/2023, 7:35:34 PM	Remove	
DEV		12/6/2022, 5:08:30 AM	Remove	
DLA		1/31/2023, 3:43:59 AM	Remove	
DW		12/28/2022, 11:26:53 PM	Remove	
PM		12/8/2022, 7:18:33 PM	Remove	
Public	Public is a built-in role that has no privileges enabled in this PhoenixAI Cloud ac...	12/5/2022, 11:39:51 PM	Remove	

Cluster access control

PhoenixAI employs a role-based approach to permission management. Instead of assigning privileges directly to users, privileges are assigned to roles. Users are then associated with one or more roles, simplifying user management and privilege assignment. However, to ensure compatibility with the old MySQL-style privilege system, we have retained the option to assign privileges directly to users.

The privilege system allows administrators to grant or revoke specific privileges at different levels of granularity, such as databases, tables, columns, and functions. This ensures that users have precisely the necessary permissions without unnecessary access. Here is a diagram to describe the PhoenixAI privilege item tree:



1. **Predefined Roles:** PhoenixAI comes equipped with various built-in roles for new users. These roles include root, cluster_admin, db_admin, and user_admin.
2. **Role Hierarchy and Privilege Inheritance:** Custom roles can be crafted to match specific business needs. Roles can also be assigned to other roles, establishing a hierarchy for efficient privilege inheritance and streamlining management. The maximum inheritance depth for a role is 16 levels.
3. **Principle of Least Privilege via Role Switching:** Users can switch to different roles using the "set role" command, adhering to the principle of least privilege. Active roles provide users access to role-specific privileges during their session. Default roles, activated upon user login, enhance privilege protection. Users without a default role still have the automatically activated public role.

Row access policy and column masking policy

In scenarios involving business data, there's a requirement to dynamically determine whether a role/user has permission to access specific rows, sensitive data columns, or obfuscated data. This dynamic capability ensures that data access is based on session information rather than storing multiple copies of the data. While views can address some of these needs, they have limitations when applied across multiple tables. Views lack flexibility in permission delegation and can lead to management complexities.

Column Masking and Row Access policies are employed to safeguard sensitive data from unauthorized access while allowing authorized users to query sensitive data during runtime. These policies do not modify the actual data within existing tables but rather apply data transformations based on user-defined conditions. This enables column-level data masking or row-level data filtering. For instance:

- A Column Masking Policy can dynamically return either encrypted or unencrypted data columns based on a user's role.

- A Row Access Policy filters rows based on a user's role and specific attributes, for example, region.

The benefits of these policies include flexibility for multiple use cases, simplified data modification, avoidance of excessive view creation, and more granular access control. They can be applied to TABLEs, VIEWs, MaterializedViews, and External Table objects. The policies are effective in queries, insertions, CTAS, updates, and deletions.

Auditing

Audit logs play a crucial role in monitoring and maintaining the security and compliance of systems. They provide a detailed record of various actions and activities performed within a system, such as database operations, configuration changes, and user interactions. Audit logs enable organizations to track and analyze these events, ensuring accountability, detecting anomalies, and aiding in incident investigations.

You can find all the database-related audit logs in PhoenixAI's `fe.audit.log`, which encompass activities like DDL (Data Definition Language) and DML (Data Manipulation Language) operations. These audit logs are maintained at the cluster level, allowing you to access comprehensive information through the query history page on the Cluster GUI. Admins (users who have `user_admin` and `db_admin` roles) have the privilege to check all of the query history, and users without `user_admin` and `db_admin` roles can only access history query records run by themselves. No cluster audit logs are allowed for these users.

Start time	Query ID	Status	SQL	Warehouse	Duration	Memory	Client IP	Query user	Queue Time	Cpu Cost	Scanned Bytes
06/30/2026, 16:03:24	e5cb69dc-74d7-11f1-9634-0630363a3d13	Success	SELECT i.campaign_id, c.campaign_name, COUNT(*) AS impressions, COALESCE(SUM(c.ost),0) AS spend ...	default_warehouse	373ms	8.57MB	35.92.74.227	admin	0ms	46ms	37.57MB
06/30/2026, 16:03:24	e5cb42c7-74d7-11f1-9634-0630363a3d13	Success	SET enable_materialized_view_rewrite = true	default_warehouse	0ms	-	35.92.74.227	admin	0ms	0ms	-
06/30/2026, 16:03:24	e5cb1bb3-74d7-11f1-9634-0630363a3d13	Success	ROLLBACK	default_warehouse	0ms	-	35.92.74.227	admin	0ms	0ms	-
06/30/2026, 16:03:24	e5caf4a2-74d7-11f1-9634-0630363a3d13	Success	SELECT i.campaign_id, c.campaign_name, COUNT(*) AS impressions, COALESCE(SUM(c.ost),0) AS spend ...	default_warehouse	374ms	9.53MB	35.92.74.227	admin	0ms	58ms	38.54MB
06/30/2026, 16:03:24	e5cad91-74d7-11f1-9634-0630363a3d13	Success	SET enable_materialized_view_rewrite = true	default_warehouse	0ms	-	35.92.74.227	admin	0ms	0ms	-
06/30/2026, 16:03:24	e5caa680-74d7-11f1-9634-0630363a3d13	Success	ROLLBACK	default_warehouse	0ms	-	35.92.74.227	admin	0ms	0ms	-
06/30/2026, 16:03:24	e5ca0a31-74d7-11f1-9634-0630363a3d13	Success	SELECT i.campaign_id, c.campaign_name, COUNT(*) AS impressions, COALESCE(SUM(c.ost),0) AS spend ...	default_warehouse	376ms	8.33MB	35.92.74.227	admin	0ms	65ms	37.36MB
06/30/2026, 16:03:24	e5c9e31f-74d7-11f1-9634-0630363a3d13	Success	SET enable_materialized_view_rewrite = true	default_warehouse	0ms	-	35.92.74.227	admin	0ms	0ms	-
06/30/2026, 16:03:24	e5c9e31e-74d7-11f1-9634-0630363a3d13	Success	SET enable_materialized_view_rewrite = true	default_warehouse	0ms	-	35.92.74.227	admin	0ms	0ms	-
			SELECT i.campaign_id, c.campaign_name, C								

Additionally, within the PhoenixAI web manager page, you can access a range of information related to cluster actions such as purchases, deletions, scaling, suspensions, and resumptions. The diagram below illustrates this functionality.

Cluster event of Demo-Webinar						
Action	Details	Started At	Ended At	Result		
Resume	Manual resume by Sida Shen	6/30/2026, 3:59:52 PM	6/30/2026, 4:00:36 PM	Finished		
Suspend	Manual suspend by Sida Shen	6/29/2026, 2:44:30 PM	6/29/2026, 2:45:20 PM	Finished		
Resume	Manual resume by Sida Shen	6/29/2026, 1:55:57 PM	6/29/2026, 1:56:45 PM	Finished		
Suspend	Manual suspend by Sida Shen	6/29/2026, 11:44:20 AM	6/29/2026, 11:45:44 AM	Finished		
Resume	Manual resume by Sida Shen	6/29/2026, 9:51:40 AM	6/29/2026, 9:52:25 AM	Finished		
Suspend	Manual suspend by Sida Shen	6/26/2026, 4:22:21 PM	6/26/2026, 4:23:12 PM	Finished		
Resume	Manual resume by Sida Shen	6/26/2026, 4:05:21 PM	6/26/2026, 4:06:06 PM	Finished		
Suspend	Manual suspend by Sida Shen	6/26/2026, 2:26:16 PM	6/26/2026, 2:27:36 PM	Finished		
Warehouse scaling	Decommission Warehouse[default_warehouse] Node size to 3 by Sida Shen	6/26/2026, 1:49:01 PM	6/26/2026, 1:50:26 PM	Finished		
Resume	Manual resume by Sida Shen	6/26/2026, 1:48:05 PM	6/26/2026, 1:48:51 PM	Finished		
					Total 117	
					10/page	
					< 1 2 3 4 5 6 ... 12 >	

Data Security

Data encryption in transit

PhoenixAI Cloud uses SSL during transit.

Data encryption at rest

PhoenixAI persists metadata on the FE node's disk. This metadata includes, but is not limited to, all of the table schema and data distribution information. In PhoenixAI Cloud, we provide optional disk encryption capability. This ability ensures all the data in EBS cannot be directly parsed in plain text after the data is obtained, which guarantees data security and minimizes potential safety hazards caused by data leakage.

Also, in an elastic cluster, PhoenixAI Cloud can persist the table data in the customer's cloud storage, for example, in an S3 bucket. With the SSE encryption capability enabled by default in AWS S3, PhoenixAI Cloud ensures that the data in the elastic cluster will not be accessed in plain text by the AWS S3 server. On this basis, users can choose to encrypt the S3 bucket with their own AWS KMS key to SSE encrypt the data in an elastic cluster in a customized way.

User identity information such as username/password is stored in the database with field-level encryption (Bcrypt algorithm) to ensure that the user identity information cannot be parsed in plain text on the server side.

Backup service

In the event of hardware failure, human error, malicious attacks, software bugs, and natural disasters, and to meet compliance requirements, backup and recovery serve as an important safeguard for database systems to protect data security.

PhoenixAI backup refers to the process of creating copies of data to prevent data loss, while restore is the process of utilizing backup files to reconstruct the database back to the state at the time the backup was taken.

PhoenixAI supports backing up to remote object storage and HDFS, supports incremental backups based on partitions, and supports backup and recovery at the database level. The destination for a backup is known as a repository.

Here are some scenarios of backup and restore:

1. **Regular Backup**

To prevent data loss, databases need to be backed up regularly, usually by setting a backup policy, e.g. backing up the database every night at midnight. Regular backups ensure that the database can be restored to a previous state in case of any problems.

2. **Backup Before Upgrade**

When upgrading the database version, a backup could be performed before the upgrade to prevent any accidents during the upgrade process. The backup can be used to restore the database.

3. **Data Migration**

When migrating a cluster to a new cluster, a backup should be done first, then the backup data can be restored in the new cluster to ensure data integrity during migration.

4. **Emergency Recovery**

If the database fails or data corruption occurs due to an attack during operation, the backup can be used to perform emergency recovery and restore the database back to the state before the failure.

5. **Test Environment Restore**

Developers often need to use real data in test environments for testing. Migrating a production backup into a development environment can be used to provide the necessary real data.

Disclaimer: Information provided in this Security Whitepaper is aligned with current laws and regulations as of July 2026. Content may change from time to time due to changes in local or international laws and regulations. A new version of this document will be published at the time revisions have been made. Any questions related to PhoenixAI's security process and procedures can be directed to privacy@phoenixdata.ai