FROST & SULLIVAN
沙利文

头豹 LeadLeo

# 2023 China Cloud-Native Database Market Report

Database Paradigm Shift、AI4DB、DB4AI、Observability

2023. 12

LeadLeo Research Institute
Frost & Sullivan (China)

# Table of Content



**Note: This research framework represents the entire content of the full version, while the abstract version has only selected some parts for overview reading.**

# 1.1 Cloud-Native Database Definition

## Key Insight

The essence of cloud-native databases lies in the architectural transformation and evolution of cloud service levels brought about by cloud-native principles. By closely aligning with cloud computing infrastructure, database architecture is designed to fully leverage the advantages of cloud computing, resulting in expanded database performance and functionality.

❑ **The core of a cloud-native database is the architectural transformation and service efficiency improvement brought about by cloud-native principles**

The cloud-native database defined in this report is designed based on the characteristics of cloud computing infrastructure, making full use of cloud resources such as computing, storage, and networking. This design results in a database with enhanced performance and expanded functionality, characterized by high scalability, high elasticity, high security, observability, and a high degree of automation. Cloud-native databases primarily represent the adaptation and evolution of database deployment patterns and architectures to cloud environments, independent of data models. Both SQL and NoSQL databases can evolve into cloud-native databases. Frost & Sullivan believes that the core of a cloud-native database lies in the architectural transformation and service efficiency improvement brought about by cloud-native principles. Specifically:

- Architectural Transformation: Cloud-native databases adopt a compute-storage separation architecture, closely designed to the characteristics of cloud computing infrastructure. Under this design philosophy, the efficiency and utilization of cloud-based computing and storage resources are improved, reducing network overhead. This optimization enhances data transfer performance and storage capabilities, resulting in enhanced database performance. Furthermore, by breaking free from the constraints of traditional database architectures in a cloud environment, cloud-native databases have greater control over cloud resources, enabling on-demand resource usage and elasticity. This significantly increases agility in responding to fluctuations in user business traffic and reduces the cost of database services.
- Service Efficiency Improvement: Cloud-native databases deeply integrate with technologies and tools in the cloud computing environment, including AI and observability tools. By applying these technologies and tools, database vendors can offer services such as elastic load balancing, intelligent monitoring, and smart operations to users, thereby expanding the functional boundaries of database services. Additionally, with the help of these tools, database vendors can build database DevOps teams, enabling rapid iteration of database capabilities and feature evolution.

Based on these changes, cloud-native databases can more comprehensively meet users' demands for cost reduction and efficiency improvement.

**The Two Core Elements of Cloud-Native Databases**

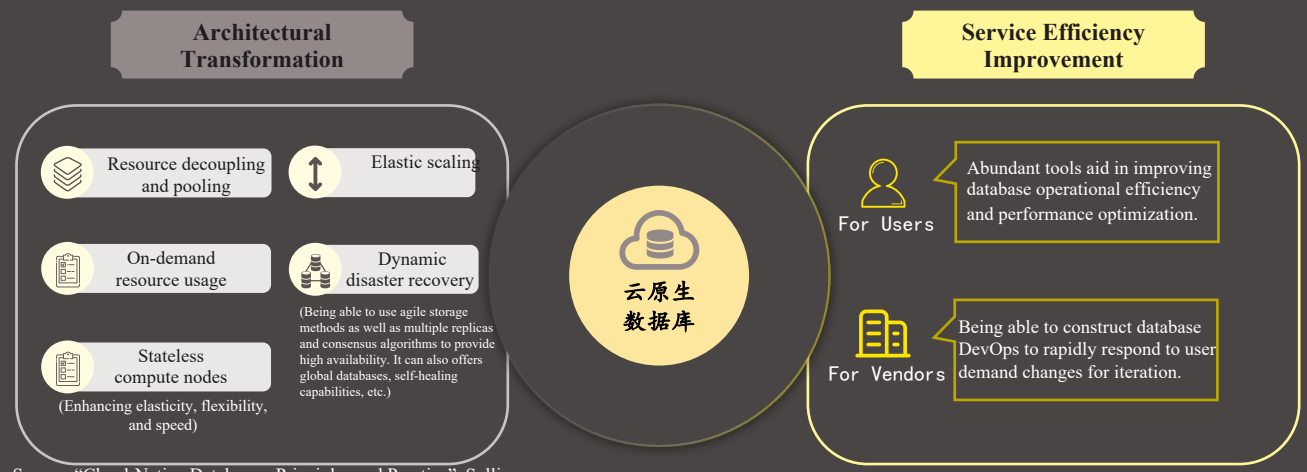


Source: "Cloud-Native Databases: Principles and Practice", Sullivan

# 1.2 Development Phases of Cloud-Native Databases

**Key Findings:**

The deployment of databases in the cloud has evolved from cloud-hosted to cloud-native, essentially shifting from 'resource-centric' to 'application-centric'. As technologies and tools continue to improve, users will be able to more easily implement database migration and understand its benefits. This will enhance their knowledge and acceptance of cloud-native databases, effectively driving the development of the cloud-native database market.

## ❑ The deployment of databases in the cloud has gone through three major phases

The deployment of databases in the cloud has gone through three major phases: Cloud Hosting, Cloud Services, and Cloud-Native. Demand changes, technological innovations, and architectural advancements are the core influencing factors that have shaped the development of these three phases.

Initially, Cloud Hosting primarily addressed the maintenance pressure of physical equipment brought about by IDC data center deployments. Users would deploy traditional database software onto cloud-hosted servers, with database management still being the responsibility of the users. As cloud service providers improved their technical capabilities and expanded their teams, they began to take on the operational aspects of database management. They optimized database performance and functionality, allowing users to access high-performance databases through interfaces, thus reducing the cost and burden of user database operations. This enabled users to benefit more effectively from cloud computing services.

In the first two phases, vendors primarily focused on relieving users of the burden of resource and database operations management and reducing costs by leveraging shared resource pools. They also improved database performance and functionality to some extent. However, due to the limitations of traditional architectures, databases were still unable to fully harness the advantages of cloud computing, such as high scalability, elasticity, and availability.

With the increasing significance of data and continuous growth in enterprise data capacity, high concurrency demands have risen. Additionally, cloud computing has become the primary environment for handling workflows. In these scenarios, Cloud Service stage databases faced challenges such as low scalability and low resource utilization efficiency, making it difficult to meet changing demands. Therefore, through technological innovation and architectural advancements, databases have fully integrated with cloud computing systems, transitioning into the Cloud-Native phase, which has become an inevitable evolution in the deployment of databases in the cloud.

**The evolution of cloud-deployed databases is a shift from "resource-centric" to "application-centric."**

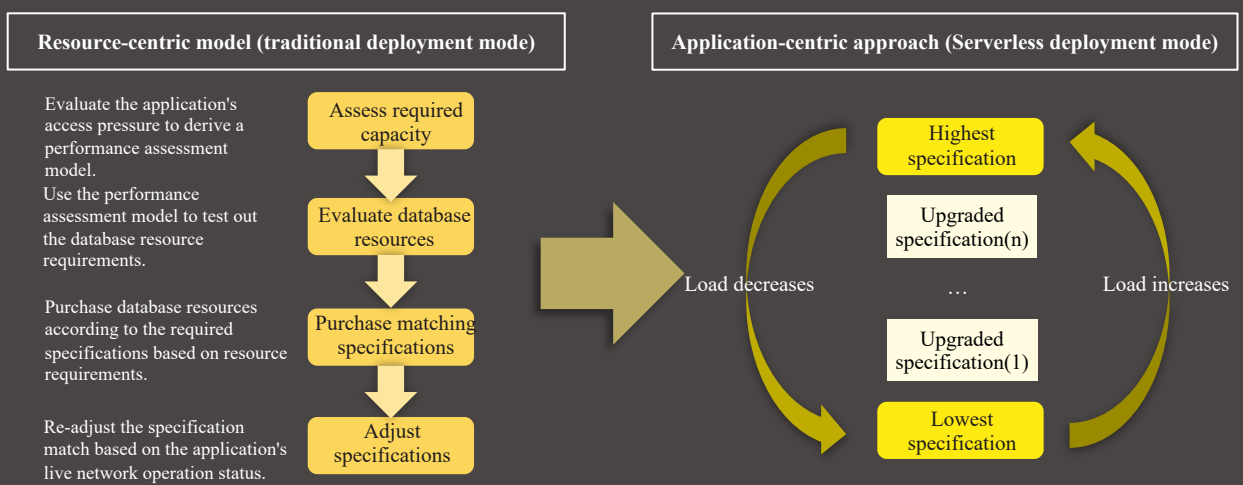


Source: 51CTO,Sullivan

## ❑ Cloud-native databases are entering a rapid development phase, with Serverless deployment mode as the core direction

Over the past decade since the first cloud-native database was introduced, cloud computing and cloud-native technologies have continuously improved, supporting the iteration and enhancement of cloud-native databases. These databases have significantly improved in performance, and the industry's product offerings have become more diverse. This includes databases with various data models and deeply integrated intelligent operations tools, allowing for greater flexibility and comprehensiveness in meeting the needs of different scenarios. Cloud-native databases have completed their technological foundation and commercial validation, and they are now entering a rapid development phase.

In this rapid development phase, the key to the growth of cloud-native databases lies in their ability to fully leverage the benefits of cloud-native principles, such as elastic resource allocation, precise billing, and ease of database management. This enables them to efficiently and flexibly address the needs of enterprise applications. Consequently, finer-grained resource deployment and the Serverless deployment model, which eliminates the need for users to manage the underlying infrastructure, have become the central focus of development in this current stage.

**Enterprises have a high willingness to adopt cloud-native databases**

**Survey of enterprise willingness to use cloud-native databases**

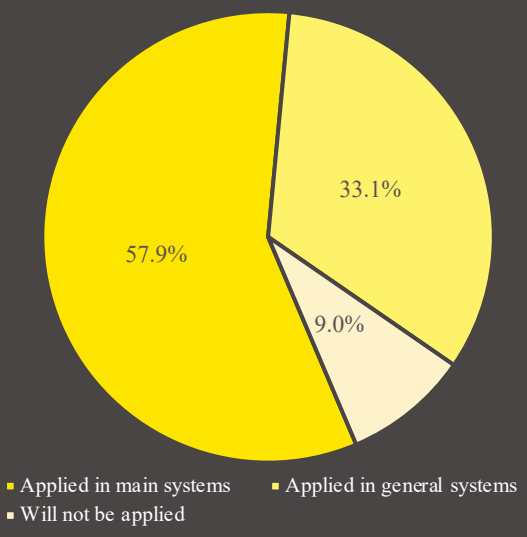


**Main factors affecting the choice of cloud-native databases**

**Insufficient Awareness / Habit Conflict:**
**Issue:** Traditional DBAs have a limited understanding and mastery of cloud computing technologies.
**Solution:** Database vendors can raise awareness by providing technical sharing, case demonstrations, etc. They can also facilitate the transition by offering tools for extension and adaptation, and by providing Serverless deployment to make it easier for users to get started and change habits.

**Data security**
**Issue:** Cloud-native databases are primarily deployed in public cloud environments, which increases potential data security threats.
**Solution:** Database vendors need to continuously improve database security features, such as multi-layer access control, auditing and data encryption, data backup and recovery, and service security technologies.

**Complex Architecture or High Coupling:**
**Issue:** Existing databases have a high degree of coupling with applications, and the application business logic is complex. The workload of conversion is too large, and the benefits are unclear.
**Solution:** Leading database vendors possess methodological expertise, offering migration assessment, migration strategy support, and ecological tools like database migration tools, which can help users migrate databases smoothly. Furthermore, by providing capabilities like FinOps and Serverless, they can assist users in effectively perceiving benefits.

## ❑ User awareness of cloud-native databases is gradually increasing, further strengthening user demand

According to statistics from the China Academy of Information and Communications Technology, 91% of surveyed enterprises consider using cloud-native databases, indicating that the value of cloud-native databases has been perceived by users.

According to interviews and research conducted by LeadLeo Research Institute, the acceptance of cloud-native concepts by enterprises has significantly improved. However, due to the fact that traditional database deployment and application methods have been used for many years and are more aligned with user knowledge and habits, some enterprises hesitate when considering a switch to cloud-native databases. In reality, the transition to cloud-native databases does not bring significant changes to existing applications. The hesitation of these enterprises primarily stems from a lack of understanding of cloud-native databases. As Serverless database technology becomes more mature and database vendors enhance their ecosystem tool capabilities (such as FinOps and database migration tools), it will facilitate enterprises in conducting testing and usage more flexibly and conveniently. This will enable them to effectively perceive cloud-native databases, gradually accumulate cases for promotional purposes, and ultimately enhance the overall market awareness of cloud-native databases.

Source: CAICT, Sullivan, LeadLeo

# 1.3 Analysis of the Necessity of Cloud-Native Databases

**Key Finding:**

**Cloud-native databases are leading a shift in the philosophy of database usage towards "data-centric." Database vendors are achieving this transformation by offering more comprehensive database services, reducing the enterprise's focus on database operations and management, and allowing them to concentrate more on extracting the value from data and implementing business logic. This approach aligns better with the data-driven business and the increasing trend of enterprises migrating to the cloud, meeting their evolving database usage needs.**

## ❑ Cloud-native database services transform the concept of database usage to be "data-centric."

The traditional database usage concept revolves around being "database-centric," focusing primarily on the management and maintenance of databases, including tasks such as configuring hardware resources, optimizing data storage, and tuning query performance. When dealing with different business workloads, users need to manage and maintain multiple databases with similar or dissimilar data models and handle data interactions between these databases.

Cloud-native databases are delivered in the form of cloud computing services, and their core concept shifts towards being "data-centric." On the operational and management front, thanks to the seamless integration of cloud-native databases with cloud resources and deep integration with complementary tools, data can be accessed seamlessly, processed efficiently, and analyzed intelligently. In this mode, cloud-native databases can provide dynamic resource management, scalability, and rapid fault recovery capabilities. Moreover, database vendors can establish data links between different database services with various data models, matching them to different business workloads and enabling automatic data synchronization between databases. As a result, enterprises can focus more on extracting the value of data and implementing their business logic.

Furthermore, cloud-native databases can offer global cross-region data synchronization, allowing applications to access databases with low latency from nearby locations. This capability can help enterprises expand their business nationally and globally, providing global cross-region database disaster recovery capabilities.

## ❑ Cloud-native databases will be the core foundation supporting the development of enterprise data-driven business strategies

As fine-grained operations continue to develop, the interconnectivity and complexity of data between multiple business units within enterprises are increasing. Faced with diverse and large-scale data, enterprises need database solutions that can maintain stable performance and low latency. Meanwhile, with the ongoing trend of cloud computing strategies in various countries, the migration of enterprise operations to the cloud will continue to deepen. Currently, China has introduced a series of policies to promote the widespread adoption and deep application of cloud computing technologies, establishing the foundation for the demand for cloud-based services. In this environment, cloud-native databases alleviate the complexity of using one or multiple databases by providing comprehensive services and leveraging resource elasticity to help enterprises adapt to rapidly changing business demands. They form the core foundation for driving the development of data-driven business strategies within enterprises.

**"Functions that can be realized under the 'data-centric' concept of cloud-native databases**

**Improvement to observability**

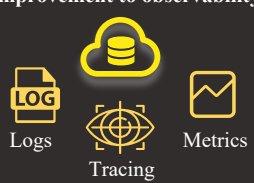


- Cloud-native databases are fully combined with cloud resources, which can better collect internal and external database data for problem analysis and system diagnosis
- With ecosystem tools, users' stress in discovering and understanding issues can be significantly reduced, and intelligent tools can be provided to assist in performance tuning.

**Global Database**

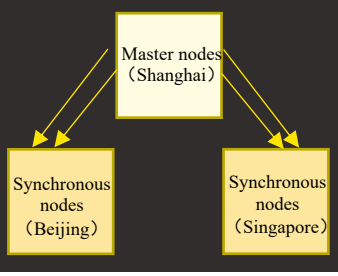


- Providing nearby reading, physical replication-based, low-latency cross-regional synchronization capabilities. Use database resources in different regions with the same experience and improve disaster recovery capabilities.

**Deeply integrated database services**

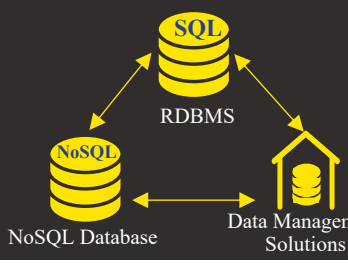


- Improving data processing efficiency through deep integration
- Alleviating users' stress in operating and managing multiple types of databases, and focus on data analysis and the application itself.

Source: Sullivan

# 2.2 Collaborating Databases with Artificial Intelligence Continuous

Key Findings

The deep integration of AI and databases comprises two directions: AI4DB and DB4AI. AI4DB can significantly improve the efficiency of database operations and management as well as performance optimization, with tremendous application potential. Meanwhile, DB4AI can address the challenges of AI implementation, achieving increased AI efficiency and cost reduction. These two approaches complement each other, jointly driving advancements and innovations in AI and database technologies.

❑ **The application of AI4DB has enormous potential, and database vendors should begin laying the groundwork for technological development and capability enhancement.**

AI4DB involves using AI to optimize databases, enabling more precise, efficient, and agile database operations and management. It enhances efficiency in handling complex tasks, reduces daily workloads, and frees up human cognitive bandwidth for more critical tasks. The implementation of AI4DB can be divided into four development stages:

- Suggesting Stage: In this stage, the focus is on reducing tasks that consume human cognitive resources through plug-and-play intelligent tools. Examples include load management optimization, database performance monitoring and alerts, database auditing, and SQL optimization. Currently, most database vendors have deployed intelligent operations and maintenance tools to accomplish these functions. Some tools are also incorporating Natural Language Processing to simplify SQL writing and optimization.
- Assisting Stage: AI assistance can be further integrated into database components. For instance, combining AI tuning models with query optimizers allows for user-level parameter tuning for query optimization, adapting to different query characteristics. This stage concentrates on enhancing specific capabilities within the database. Some vendors have started using AI and machine learning to continuously evolve diagnostic logic and inspire new approaches to root cause analysis.
- Enhancing Stage: Building upon the previous stage, this phase aims to reduce the need for manual intervention in AI model optimization. It utilizes deep reinforcement learning to provide dynamic tuning and recommendations for optimal indexing schemes, adapting to varying environments. At this stage, the database incorporates AI services internally, using AI algorithms for self-enhancement. This can further reduce reliance on DBA expertise and offer optimization methods that are difficult to deduce through experience alone.
- Self-Building Stage: With a rich accumulation of data and algorithms, AI gains the ability to autonomously design, assemble, evaluate, monitor, and maintain databases.

Databases are complex systems and serve as critical data backbones for numerous business operations. In this context, the application of AI4DB holds significant potential, though it faces challenges in terms of practicality and reliability. However, the enhancement of AI capabilities relies on continuous data accumulation. Therefore, database vendors should initiate efforts to build and apply relevant capabilities to achieve technological development and capability enhancement in AI4DB.

**Evolution stages of AI4DB**

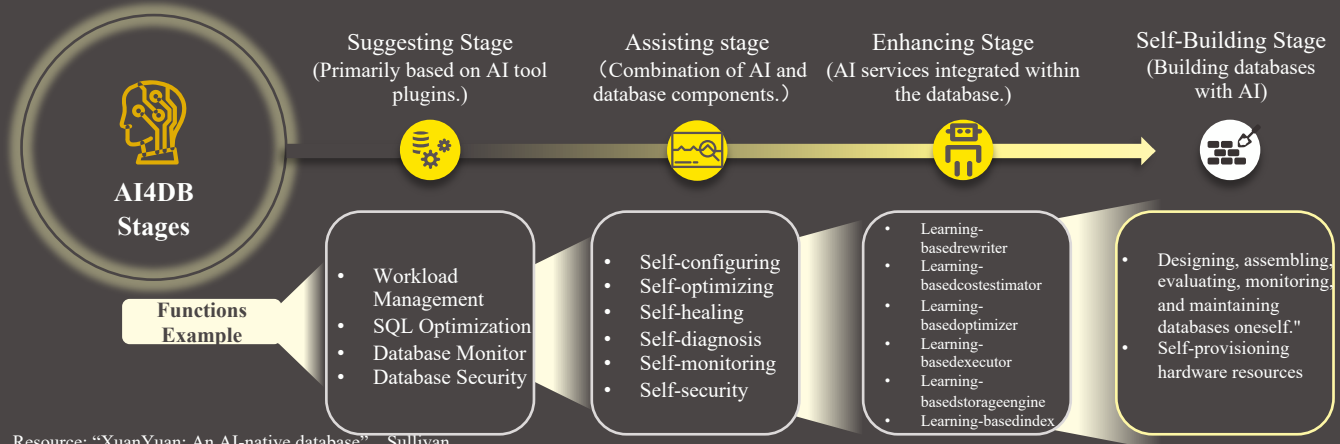


Resource: "XuanYuan: An AI-native database", Sullivan

## ❑ Establishing the capabilities of DB4AI can address the challenges of AI implementation, achieving increased efficiency and cost reduction

Sullivan argues that two significant challenges in current AI implementations are increasing hardware requirements and efficiency issues. On one hand, as models become increasingly complex with more parameters, hardware requirements and costs also escalate. On the other hand, the interaction between data and models is quite intricate, and in the face of evolving business needs, continuous algorithm improvement, and expanding data volumes, this complexity is bound to grow. Traditional development processes can lead to low efficiency in model deployment, including issues related to data flow and workflow efficiency.

DB4AI refers to optimizing the efficiency of data storage, management, and operations within the database to meet the demands of AI model training, making the training process more efficient. Specific capabilities that can be achieved include:

- Simplified Modeling: Through unified management of data, features, and AI models within the database, and by providing a unified SQL interface, users can perform model development, training, and inference using SQL without the need for multiple languages (such as Python or R) for AI development.
- Assisted Computing: Adding AI operators and support for vector calculations within the database enables the use of database computing capabilities for model training and inference tasks, reducing the computational burden of AI training.
- Model Reusability: By persisting AI models through mechanisms like materialized views and query tables, models become easier for users to manage and reuse.

By utilizing the above capabilities, a deep integration of data and AI can avoid data fragmentation across models, reducing data management complexity and latency. This leads to improved data flow efficiency and streamlined workflows. Moreover, by unifying the management of data, features, and models and supporting AI training and inference tasks, the hardware requirements for AI-related tasks can be reduced, subsequently lowering hardware costs.

**The value of DB4AI**

The processes and elements involved in AI model training

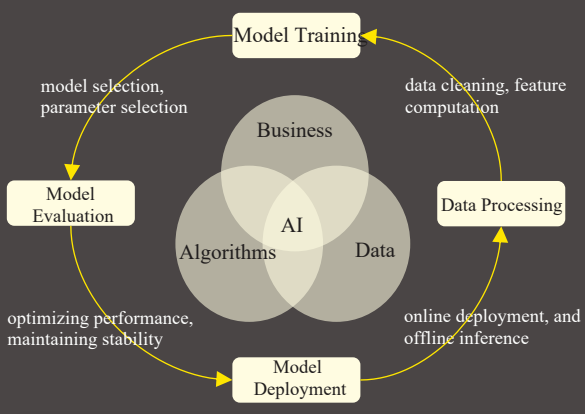


- Deploying AI requires comprehensive consideration of three main elements: business, algorithms, and data. Over time, these elements are continuously iterating and evolving, such as business requirements that change continuously; algorithms that need to be regularly updated while maintaining quality; and data volume that is continuously expanding, presenting increasing challenges to the practical implementation of AI.
- The AI model training process includes multiple steps, where the interaction between the model and data involves multiple systems, with complexity increasing during the iteration process.

The comparison of traditional algorithm development processes evolving to the use of DB4AI for algorithm development.

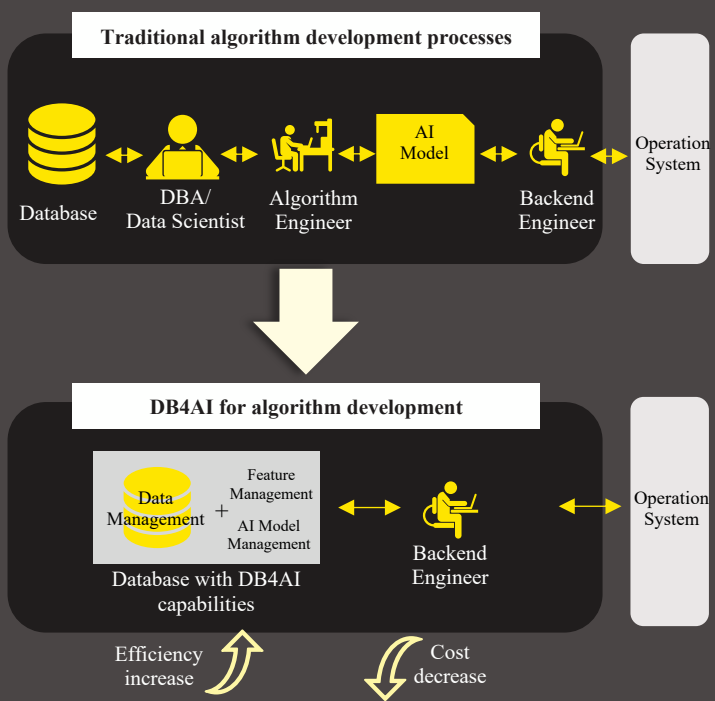


Resource: "XuanYuan: An AI-native database", Sullivan

# 2.3 Resource Utilization Optimization

Key Findings

To fully harness the value of cloud-native databases in maximizing resource utilization and reducing costs, users are confronted with two significant challenges: resource configuration optimization and cost management difficulties. Database vendors should offer comprehensive resource deployment strategies, ongoing optimization of Serverless performance, and provide FinOps tools to assist users in optimizing resource efficiency.

## ❑ Providing efficient deployment strategies and FinOps tools will be crucial in helping users optimize resource utilization efficiency

Improving resource utilization to save costs is a core value proposition of cloud-native databases, but realizing this value fully still faces two major challenges: resource configuration optimization and cost management difficulties.

The challenge of resource configuration primarily arises from the complexity of the available configuration options and the dynamic nature of business workloads. Relying on manual efforts to carefully match these options makes it challenging to maximize the value of resource configuration in complex and ever-changing environments. Moreover, most users still manage resources with a traditional mindset, leading to significant waste.

The difficulty in cost management stems from the introduction of cloud-native technologies, which has brought changes to traditional centralized financial budgeting and IT management models. Traditional database financial budgeting involves pre-assessing the quantity of hardware and software resources to be purchased, resulting in a clear budget. In the cloud-native environment, resource elasticity leads to cost uncertainty and variability. Furthermore, cloud-native databases may be deployed across multiple business units, increasing the complexity of service tracking and cost allocation.

The key to addressing these challenges lies in tackling resource waste, often caused by overprovisioning of resources to meet business demands. To address this, cloud-native databases should be able to provide deployment strategies that combine the preprovisioning mode for read/write nodes within a single cluster with Serverless, achieving resource guarantees and elasticity to handle spikes. Serverless database technologies are continually evolving, with ongoing performance optimizations to minimize resource consumption. Additionally, these deployments should offer the ability to automatically adjust based on workload changes, reducing the need for manual configuration and optimizing resource allocation capabilities.

Simultaneously, providing FinOps tools is essential to guide users in transitioning from traditional financial budgeting and IT management models to adapt to the cloud-native environment. FinOps tools should offer visual guidance for specification recommendations, budgeting, and quota management. They should also check for unreasonable configurations and provide capabilities such as load-aware scheduling and topology-aware scheduling to support ongoing rate and elasticity optimization. By modeling and quantifying business computing resource usage and establishing complementary metrics such as watermarks and redundancy levels, enterprises can receive recommendations and implement various strategies, such as load balancing, resource consolidation in offline modes, and automated scaling, to effectively optimize resource allocation and truly achieve cost reduction and efficiency improvement.

**Cloud-native database deployment mode strategies and key points for implementation of enterprise FinOps**

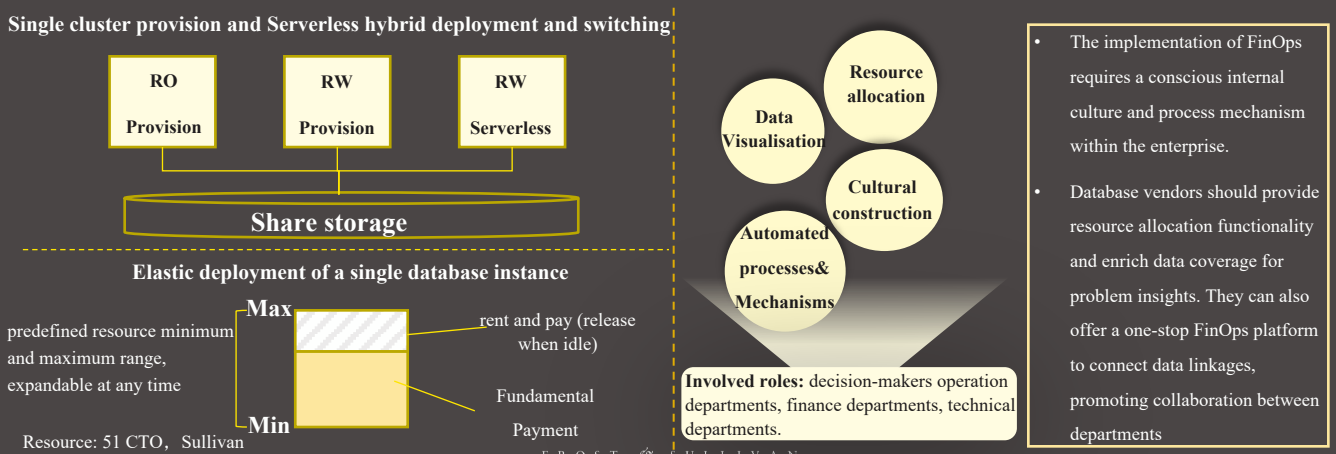


Resource: 51 CTO，Sullivan

# Evaluation Metric System for Cloud-Native Database Vendors

This report establishes six major dimensions for evaluating the performance of competitors: Cloud Native and Innovative Features and Technologies, Fundamental Features and Technologies, Operations and Management, Security Assurance, Services and Market, and Ecosystem Construction. Based on the scoring, the capabilities of the competitors are categorized into different tiers, and the comprehensive capabilities of the competitors are showcased based on the total score obtained in all these areas.

| Tier 1 Indicator | Tier 2 Indicator | Key Points |
|---|---|---|
| Cloud-Native&Innovative Functions and Technologies | Product Classification | Data models, load scenarios, distributed architecture, deployment scenarios. |
| | Cloud Computing Service Capabilities | Resource isolation, resource utilization optimization, hybrid multi-cloud. |
| | Cloud-Native Capabilities | Scalability, large-scale performance and stability, Serverless, HTAP. |
| | Distributed Architecture Innovation | Transaction throughput optimization, compute-storage separation, HTAP performance. |
| | Observability Capabilities | Performance indicators, logs, link tracking. |
| | DB4AI Capabilities | SQL calls for machine learning, in-database training and inference, automatic feature engineering. |
| Fundamental Functions and Technologies | Transaction Scenario Capabilities | Transaction architecture, consistency, concurrency control, transaction throughput optimization. |
| | Analysis Scenario Capabilities | Sharding rules, smooth expansion, performance expansion loss, cross-domain partitioning. |
| | Distributed Architecture Basic Capabilities | Physical resource layer, scalability, cross-regional deployment capabilities. |
| Operations and Management | Operational Optimization Capabilities | Intelligent operational tools, load balancing, query optimization, SQL optimization. |
| | AI4DB Operational Technology Reserves | Database configuration, performance optimization, database monitoring, AIGC applications. |
| | Database Management Capabilities | Resource isolation level, network topology detection, host load shifting. |
| Security capabilities | Data Backup Management Capabilities | Backup recovery, backup management, node management, high availability features. |
| | Database Disaster Recovery Construction | Data center failure switch-over plans, off-site disaster recovery switch-over duration, various fault RTOs. |
| | Database Access Security | Permission management, identity authentication, anti-malicious intrusion, access control. |
| | Data Security | Grading and classification, data desensitization, data encryption, data masking. |
| | Audit Capabilities | Audit record scope and protection. |
| | Security Protection Capabilities | Vulnerability scanning, anti-SQL injection, anti-ransomware, DDoS protectio |
| Services and Market capabilities | Development and Open Compatibility | Development interfaces, SQL standards, syntax compatibility, open engines, new hardware. |
| | Migration Adaptation and Transformation | Highly Portable Tools and Services, Migration Plans, Zero-Downtime Upgrades Technological Sustainability |
| | Technological Sustainability | Research and Development Pathways, Standard Writing, Academia-Industry Research, Talent, Evangelism, Partnerships |
| | Localization and Adaptation Construction | Servers, Chips, Middleware, Operating Systems, Optimal Combinations |
| Ecological construction | Service Support | Implementation Services, Value-Added Services, Expert Teams, Product Documentation |
| | Commercial Maturity | Industry Application Breadth - Product Coverage Across Various Industries<br>Industry Application Depth - Granularity and Depth of Product Implementation<br>Advantageous Service Models - Continuous Optimization of Service Capabilities to Meet Needs |

# 2023 Comprehensive Competitive Performance in the Chinese Cloud-Native Database Solutions Market——Frost Vendor ScoreCask ™

This report establishes six major dimensions for evaluating the performance of competitors: Cloud Native and Innovative Features and Technologies, Fundamental Features and Technologies, Operations and Management, Security Assurance, Services and Market, and Ecosystem Construction. Based on the scoring, the capabilities of the competitors are categorized into different tiers, and the comprehensive capabilities of the competitors are showcased based on the total score obtained in all these areas.

## ❑ 2023 Comprehensive Competitive Performance in the Chinese Cloud-Native Database Solutions Market—Frost Vendor ScoreCask ™

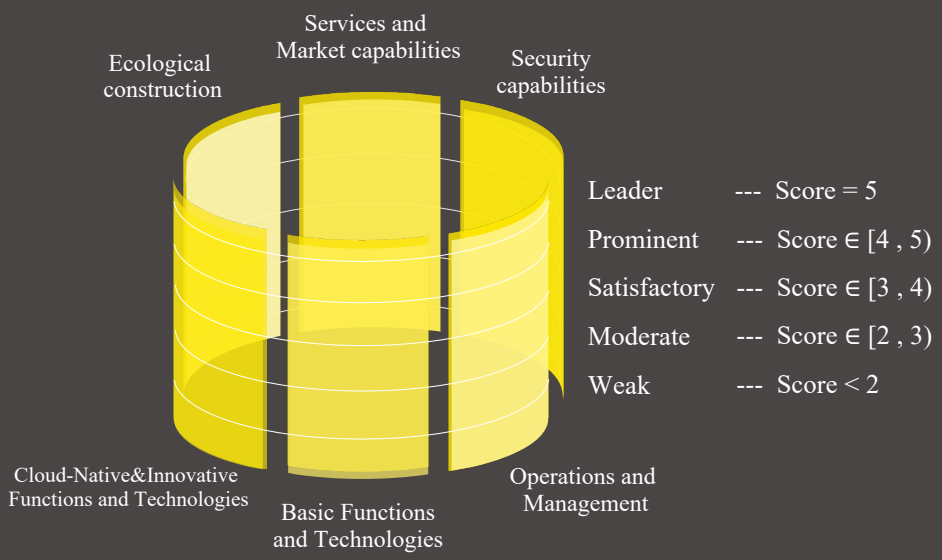


*Note: The scoreboard follows a logic of ascending from bottom to top, corresponding to a comprehensive score from low to high, with the highest score being 5 points. There are six dimensions for evaluation, forming a total score with a maximum of 30 points as the highest score.*

### ❑ Based on assessment results, the report recommends the top ten manufacturers selected for inclusion in this report.

The report employs the Frost Vendor ScoreCask ™ model for a comprehensive assessment of cloud-native database manufacturers in China. The evaluation system covers six dimensions: Cloud Native and Innovative Features and Technologies, Fundamental Features and Technologies, Operations and Management, Security Assurance, Services and Market, and Ecosystem Construction.

Based on the assessment results, Frost & Sullivan has shortlisted 10 competitors from Chinese cloud-native database manufacturers. This selection process is not merely based on performance in a single dimension but considers the capabilities of each manufacturer across all six dimensions. The report summarizes the core competencies of each manufacturer in product orientation, user orientation, and market development, providing a reference for readers in product research and selection decisions.

The report's conclusions on the comprehensive competitive strength of the competitors' cloud-native database products and services are only applicable to the cloud-native database market development at this stage. Frost & Sullivan will continue to monitor the cloud-native database market to capture competitive trends.

Sullivan Market Insight

# 2023 Comprehensive Competitive Performance of China's Cloud Native Database Market Frost Vendor ScoreCask ™ —— Amazon Web Services（AWS）

## Key Insights

Amazon Web Services, with its extensive technical expertise and experience in the fields of databases and cloud computing, has developed a cloud-native database solution that seamlessly integrates databases with big data and AI tools. This solution is designed to fulfill the core performance requirements of databases while also catering to a diverse range of application scenarios and user needs, establishing Amazon Web Services as a leader in China's cloud-native database solution market.

### 2023 Frost Vendor ScoreCask for Amazon Web Services （AWS）

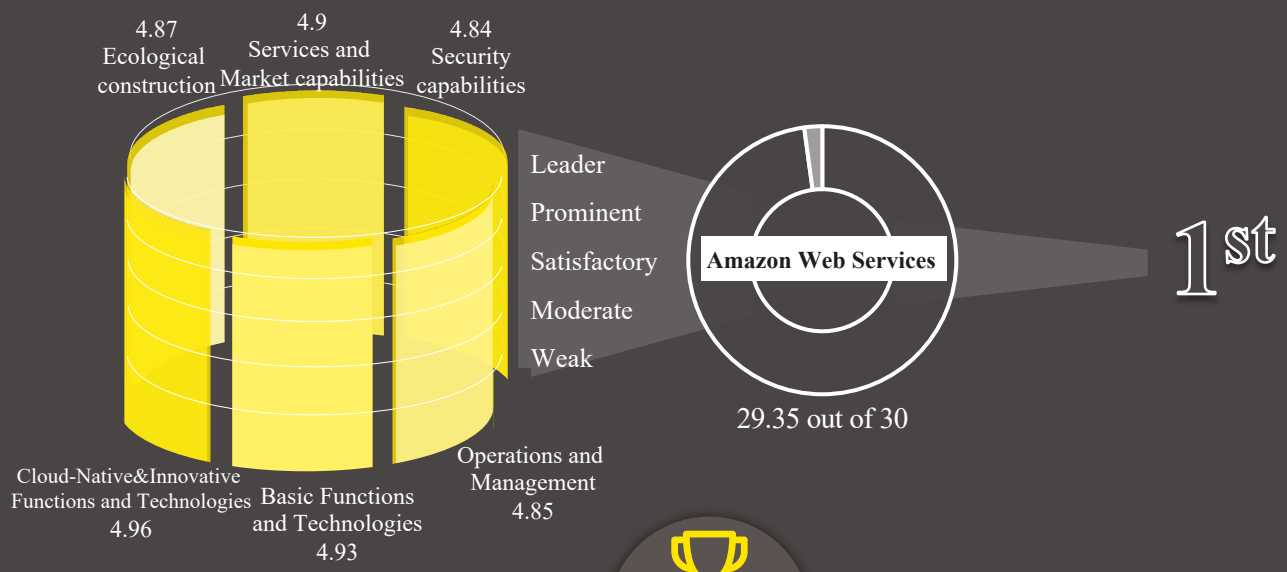


Comments

**Amazon Web Services leads in all six indicators and ranks first in overall score**

- **Product-Oriented Capabilities:** Amazon Web Services deeply integrates its database development capabilities with foundational cloud computing technologies to create a range of cloud-native databases, achieving market-leading levels in terms of cloud resource utilization and user experience. The company continually leads database innovation, introducing various serverless features for databases, enabling zero-ETL integration among multiple data services, and enhancing several databases with vector search capabilities to accelerate generative AI application innovation.
- **User-Oriented Capabilities:** Amazon Web Services offers a variety of tools and services to help database users reduce costs and increase efficiency. In the cloud, the company provides a suite of real-time monitoring and intelligent operation tools to alleviate user operational burdens, allowing users to focus on business innovation. Significant efforts are invested in developing and localizing technical documentation systems, providing users with a solid foundation for learning and continuous optimization of their architectures.
- **Market Development-Oriented Capabilities:** Amazon Web Services actively disseminates cloud-native database technology practices and papers, positively influencing global industry development. Moreover, the company collaborates closely with Amazon Web Services partners to offer users a more comprehensive and enriched array of products and services.

Source: Frost Sullivan

# Amazon Web Services（AWS）

## Key Insights

Amazon Web Services, with its extensive technical expertise and experience in the fields of databases and cloud computing, has developed a cloud-native database solution that seamlessly integrates databases with big data and AI tools. This solution is designed to fulfill the core performance requirements of databases while also catering to a diverse range of application scenarios and user needs, establishing Amazon Web Services as a leader in China's cloud-native database solution market.

### ❑ AWS continues to drive end-to-end data strategies, focusing on enhancing the user experience

Leveraging its profound expertise in database and cloud computing technologies, Amazon Web Services has launched a series of industry-leading innovative services and features that focus on actual user needs, aiding users in reducing costs and enhancing efficiency while enjoying a high-quality experience.

The company has developed several fully-managed, high-performance cloud-native databases, offering peak performance for various scenarios. It provides numerous database features supporting global distributed application systems, enabling cross-regional database disaster recovery to ensure the continuity of customer businesses worldwide. Additionally, it helps reduce access latency and enhance customer experience by positioning services closer to end-users globally.

Amazon Web Services is committed to the full-stack Serverless transformation of computing, storage, databases, and associated tools. Currently, most database services support a Serverless deployment model, enhancing their elastic scaling capabilities and improving performance in large-scale scenarios. This allows users to enjoy greater agility and optimized costs without the need to manage or configure servers.

The company is actively investing in a zero-ETL future, reducing the complexity of tool interoperability, and enhancing the efficiency of data collaboration. It deeply integrates database services with big data analytics, machine learning, and generative AI services, supporting users in implementing end-to-end data strategies and focusing on business innovation to maintain competitiveness in a rapidly changing market.

**Amazon Web Services provides a comprehensive range of cloud database services.**

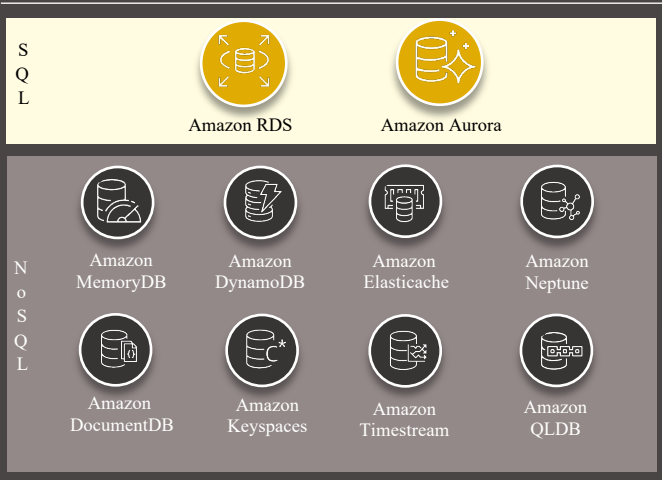


**Purpose-built databases with ultimate performance**
Modern applications require databases that are purpose-built. The purpose-built databases provide exceptional performance and cost-effectiveness.

**Serverless deployment enabling agile innovation**
Serverless architecture eliminates the need for maintenance and offers dynamic scalability, ensuring support for unpredictable workloads and business growth.

**Global architecture via one-click deployment**
Supporting global business expansion, enabling one-click database expansion worldwide to enhance the access experience for users in multiple regions, and enabling global-level disaster recovery to improve global business continuity.

**Data movement and extensive integration**
Cloud-native databases are extensively integrated with other cloud services, such as big data analytics and artificial intelligence tools, enabling the building of intelligent applications, facilitating data-driven agile decision-making, accelerating innovation, and shortening innovation cycles.

**Progress in Cloud-Native Database Deployment**

**Resource Utilization and Performance Optimization:**

- Most of AWS database services support Serverless deployment. Amazon Aurora Serverless v2 provides supreme scalability and pay-as-you-go pricing to avoid waste and reduce costs.
- The introduction of the Amazon Aurora Limitless Database feature allows databases to automatically expand beyond the write limitations of a single database and achieve high performance through database sharding.
- Supporting mixed use of Serverless and provisioned deployment within the same cluster allows flexible adaptation to various workloads and reduces migration costs.

**Enhancement of DB4AI Capabilities**

- The company has provided vector retrieval capabilities to several databases, bringing the vector retrieval function closer to the location of the source data storage and aligning more closely with users' original database usage habits, thereby accelerating generative AI innovation.

**Progress in Zero-ETL**

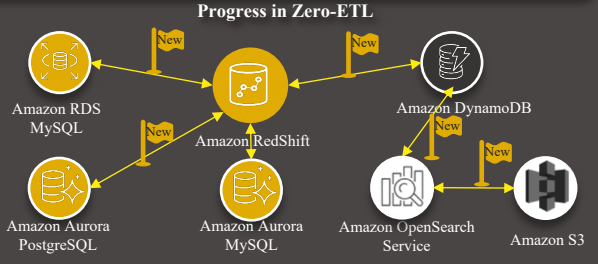


- The company successively launched zero-ETL integration of Amazon Aurora MySQL, Amazon Aurora PostgreSQL, Amazon DynamoDB, and Amazon RDS for MySQL with Amazon Redshift. Data written to these databases can be almost instantaneously and automatically replicated to the Amazon Redshift data warehouse, eliminating the need to build and maintain data pipelines and greatly simplifying the process of analysing transactional data.
- The zero-ETL integration between Amazon DynamoDB and Amazon OpenSearch Service enables easier, near-real-time powerful full-text and vector search queries on Amazon DynamoDB data.
- The zero-ETL integration between Amazon S3 and Amazon OpenSearch Service allows for direct use of the rich interactive features of OpenSearch Service on infrequently queried data stored in Amazon S3 data lakes.

Source: Amazon Web Services, Frost Sullivan

# Methodology

- Frost & Sullivan has conducted in-depth research on the market changes of 19 major industries and 532 vertical industries in China with more than 1,000,000 industry research samples accumulated and more than 10,000 independent research and consulting projects completed.
- Rooted on the active economic environment in China, the research institute, starting from data management and big data fields, covers the development of the industry cycle, follows from the enterprises' establishment, development, expansion, IPO and maturation. Research analysts of the institute continuously explore and evaluate the vagaries of the industrial development model, enterprise business and operation model, Interpret the evolution of the industry from a professional perspective.
- Research institute integrates the traditional and new research methods, adopts the use of self-developed algorithms, excavates the logic behind the quantitative data with the big data across industries and diversified research methods, analyses the views behind the qualitative content, describes the present situation of the industry objectively and authentically, predicts the trend of the development of industry prospectively. Every research report includes a complete presentation of the past, present and future of the industry.
- Research institute pays close attention to the latest trends of industry development. The report content and data will be updated and optimized continuously with the development of the industry, technological innovation, changes in the competitive landscape, promulgations of policies and regulations, and in-depth market research.
- Adhering to the purpose of research with originality and tenacity, the research institute analyses the industry from the perspective of strategy and reads the industry from the perspective of execution, so as to provide worthy research reports for the report readers of each industry.

# Legal Disclaimer

- The copyright of this report belongs to LeadLeo. Without written permission, no organization or individual may reproduce, reproduce, publish or quote this report in any form. If the report is to be quoted or published with the permission of LeadLeo, it should be used within the permitted scope, and the source should be given as "LeadLeo Research Institute", also the report should not be quoted, deleted or modified in any way contrary to the original intention.

- The analysts in this report are of professional research capabilities and ensure that the data in the report are from legal and compliance channels. The opinions and data analysis are based on the analysts' objective understanding of the industry. This report is not subject to any third party's instruction or influence.

- The views or information contained in this report are for reference only and do not constitute any investment recommendations. This report is issued only as permitted by the relevant laws and is issued only for information purposes and does not constitute any advertisement. If permitted by law, LeadLeo may provide or seek to provide relevant services such as investment, financing or consulting for the enterprises mentioned in the report. The value, price and investment income of the company or investment subject referred to in this report will vary from time to time.

- Some of the information in this report is derived from publicly available sources, and LeadLeo makes no warranties as to the accuracy, completeness or reliability of such information. The information, opinions and speculations contained herein only reflect the judgment of the analysts of leopard at the first date of publication of this report. The descriptions in previous reports should not be taken as the basis for future performance. At different times, the LeadLeo may issue reports and articles that are inconsistent with the information, opinions and conjectures contained herein. LeadLeo does not guarantee that the information contained in this report is kept up to date. At the same time, the information contained in this report may be modified by LeadLeo without notice, and readers should pay their own attention to the corresponding updates or modifications. Any organization or individual shall be responsible for all activities carried out by it using the data, analysis, research, part or all of the contents of this report and shall be liable for any loss or injury caused by such activities.

Research Analyst

Tira Huo

13660690528

tira.huo@frostchina.com

Jackey Hu

18576027961

jackey.hu@frostchina.com

Research Project Team Director

Livia Li

13149946576

livia.li@frostchina.com

www.frostchina.com ; www.leadleo.com

https://space.bilibili.com/647223552

https://weibo.com/u/7303360042