

Alibaba Cloud



C-) Alibaba Cloud

www.alibabacloud.com

2022 Edition 8

Alibaba Cloud

with a Purpose



www.alibabacloud.com

CONTENTS

About Us

- Editor in Chief / Selina Yuan
- Executive Editors / Yangqing Jia, Feifei
 Li, David Chen, Shunmin Zhu
- Review Editors / WeiTong, Olivia Kang
- Technical Advisors / Sandy Zhang, Chenny Chen, Sue Zhou, Summer Yuan, Tianyi Chao, Crystal Bai
- PR Advisor / Crystal Liu

01

Keep on Innovating

02

Innovation with a Purpo Toward a More Sustaina Inclusive Future

Decrypting Alibaba Clo Self-Developed Server

03

Tech for Good: XuanTie Cores Open-Sourced

05

06

07

08

04

Hyperscale Data Center Green Technology

Lingjie: Integrating Dig with the Cloud for the N

Six Technical Directions Next-Generation Enterp Databases

Knocking on the Door to Database 2.0

09

Apsara Luoshen – A Clo Network Born from the

	05
ose: able and	06
ud's Chip - Yitian 710	10
e RISC-V IP	14
rs Powered by	20
ital Intelligence Next Wave of Al	22
of orise	26
o Cloud-Native	32
oud-Native Cloud	36



Selina Yuan

General Manager, Alibaba Cloud Intelligence International Business Unit

01 Keep on Innovating

We have seen 2021 as a year of transitional metamorphosis towards digitization. The year 2022 will continue toward a staggering technological evolution powering massive opportunities for enterprises to achieve digital transformation.

Alibaba Cloud is constantly innovating at the forefront of business intelligence that promotes resilience and sustainability, helping businesses navigate their digital transformation journey. We are increasing our efforts to build and foster an open, collaborative environment for innovation. Alibaba Cloud continues to develop at scale to design a new way ahead in the digital age and drive cloud-native transformation. We have been at the helm of several technological breakthroughs, from delivering unique solutions for the next-generation workforce to technologically supporting the Tokyo 2020 Summer Olympics.

Alibaba Cloud has introduced new solutions to continually improve the technological experience in the digital age. Alibaba Cloud and our practices aim to bring significant benefits such as cost savings, increased efficiency, and multi-dimensional data analysis to our own organization and empower many of our customers worldwide. This year we have unveiled a new in-house processor Yitian 710, and accelerated innovation through collaboration by making XuanTie RISC-V

series processors open-source. It will assist global software developers in developing the core ability, including chips, allowing the entire RISC-V ecosystem to benefit from more innovation. We have also been working on new ways to meet the demands of the next generation of workers, such as launching Tingwu, an Al-powered meeting assistant that translates virtual meeting audio into transcripts in real-time. In addition, we have launched Lingjie, our latest integrated AI platform, capable of providing real-time data analytics with millisecond delay and various out-of-the-box services.

We also witnessed a greener and more efficient 11.11 Global Shopping Festival this year, thanks to innovative green technologies such as liquid cooling and wind energy being applied at our hyper-scale data centers which can reduce the average PUE to as low as 1.09. By running 100% on the Alibaba Cloud for the first time, the 11.11 event saw computing resources reduced by half for every 10,000 transactions compared to last year.

As Socrates said, "The secret of change is not to fight the old, but to focus on building the new." We make joint efforts with our clients and partners to innovate constantly in this digital era.

Let us keep on experiencing innovation on this incredible digital journey with Alibaba Cloud.

reli

02 **Innovation** with a Purpose: **Toward** a More Sustainable and Inclusive Future

The staggering technological advancements have created massive opportunities for enterprises to achieve digital transformation, gain a competitive advantage, and accelerate business growth. Companies across the globe are now turning their focus toward leveraging the transformative power of technology to deal with challenging social issues of the world and make a global impact. At Alibaba Cloud, the digital technology and intelligence backbone of Alibaba Group, sustainability and inclusiveness lie at the heart of our innovation and business goals. We are continually stepping up the efforts to look beyond commercial business challenges, answer society's call, and cultivate an open, collaborative space for innovation.





From supporting the Tokyo 2020 Summer Olympics to releasing new solutions for the next-generation workforce, Alibaba Cloud continues to innovate at scale to shape a new way forward in the digital age and boost cloud-native transformation. Outlined below are some of the innovation initiatives by Alibaba Cloud:

Alibaba Cloud leveraged powerful cloud infrastructure and AI technology throughout the 2021 11.11 sales event. Our self-developed Al inference chip, Hanguang 800, improved the performance of Taobao's search algorithm by a whopping 200%.

We have also been innovating to address the needs of the next-gen workforce. Our Al-powered meeting assistant, Tingwu, converts the audio of virtual meetings into transcripts in real-time with close to 98% accuracy. Likewise, we recently launched an all-in-one PC and "Cloud Workplace" solution to enable people to work remotely.

Cloud-Native Technology

Cloud Computing

Alibaba Cloud's new in-house processor named Yitian 710 is one of the most advanced processors from China that has entered mass production. We also recently announced the development of cloud-native proprietary servers, called Panjiu, equipped with the self-developed server chip Yitian 710. Both innovations will deliver outstanding performance and boost overall computing efficiency. Another key technology by Alibaba Cloud is Apsara Luoshen, the high-performance network engine powering Alibaba Cloud. In this issue, we discuss the Apsara Luoshen Cloud Network Technology and its future outlook in detail.

2021 11.11 Global Shopping Festival

Hybrid Future

Digital Sports Experience

Alibaba Cloud has released new solutions to continue transforming the sports experience in the digital age and will be providing support during the upcoming Beijing 2022 Winter Olympics and the 2022 Asian Games in Hangzhou. We rolled out the Premium DingTalk for Sports Games solution to enable seamless collaboration and communication among organizations during large-scale sports events.

Sustainable Green Technology

At Alibaba Cloud, we have pulled out all the stops to reduce our carbon footprints and boost sustainability using green data centers. To create an environment-friendly experience for our customers, we used a rare 'liquid cooling technology' deployed at our Hangzhou hyperscale data center. The liquid cooling technology by Alibaba Cloud saves **70 million kilowatt-hours (kWh) electricity consumption every year** compared to traditional data centers. Similarly, our data center in Zhangbei County in Northwestern China helped to **reduce 26,000 tonnes of carbon dioxide emission** by leveraging clean, renewable energy. Moreover, our Inner Mongolia data center **increased the consumption of wind and photovoltaic power to 45% of its energy mix in 2021 vs. 38% in 2020**.

Open-Source Initiatives and Accessible Technology

To encourage more innovation through collaboration and accelerate technological development, we made our XuanTie RISC-V series processors open-source. It will help the global software developers build their own cost-efficient RISC-V-based chips and bring more innovation benefits to the entire RISC-V ecosystem. Alibaba Group has also been leveraging a slew of open-source technologies to develop its technical systems.





To name a few, Alibaba Cloud uses Flink, Hadoop, Kubernetes based technology, and open-source deep learning frameworks like TensorFlow.

In addition to our open-source initiatives, we have also doubled down on our efforts in promoting inclusivity through accessible technologies. To make shopping accessible for everyone during the 2021 11.11 sales event, we leveraged Alibaba DAMO's speech recognition technology to make it easier for the elderly population in China to get accustomed to the digital space. A smart voice assistant 'Taoxiaobao' avoids the need for older adults to type on mobile devices by enabling voice search. Other accessibility initiatives include an "image search for medicine" feature on the Taobao app.

Summing it up, modern technologies carry the vast potential to enable enterprises to conduct business to benefit all and create a big difference globally. Alibaba Cloud is continually exploring ways to serve our customers, optimize our existing technologies, and unlock massive value for society.





Alibaba Cloud recently unveiled a new processor design for its data centers, Yitian 710. The development of our proprietary cloud-native server series, called Panjiu, is also powered by this chip. With the advent of Yitian and Panjiu, Alibaba Cloud emphasized the fusion and optimization of software and hardware. The move helps improve the final process of the full-stack cloud infrastructure and realize technology and architecture innovation and self-development, from chips and components to the whole cabinet.



11

2022 Edition 8

Customizing our own server chips is consistent with our ongoing efforts toward boosting our computing capabilities with better performance and improved energy efficiency."

Jeff Zhang, President of Alibaba Cloud Intelligence and Head of Alibaba DAMO Academy during the 2021 Apsara Conference

Let's take a look at the five key questions on Yitian 710

What Are the Key Features of Yitian 710?

Unlike Hanguang 800, Yitian 710 is a generalpurpose server CPU. As the core component of a computer system, it empowers the receiving, processing, and computing of all data in the system.

Yitian 710 is built based on the state-of-theart 5 nm lithography process and integrates a staggering 60 billion transistors in the latest ARMv9 architecture. With 128 cores and a frequency up to 3.2 GHz, Yitian 710 effortlessly balances performance and power consumption. It supports the industry-leading DDR5 memory and PCle 5.0 interface to boost the transmission rate and adapt to various scenarios in the cloud. Additionally, Yitian 710 scores 440 on the SPECint 2017 benchmark suite. It also boasts of its top performance and an energy efficiency ratio of 20% and 50%, respectively, which is higher than industry standards. These effectively help data centers save energy and reduce carbon emissions.

How Does Yitian 710 Benefit Alibaba Cloud and Our Customers?

Performance, cost, and power consumption are among the top concerns of every cloud service provider and enterprise in the cloud. Yitian 710 is designed to satisfy the demanding requirements for high concurrency, performance, and energy efficiency of cloud applications. It combines leading CPU technologies with the specific needs of cloud applications to achieve a breakthrough in performance and energy efficiency ratio. Yitian 710 also works with the Apsara system to provide highly cost-effective services for customers in the cloud. Yitian 710 will be deployed in Alibaba Cloud data centers and gradually serve enterprises in the cloud.



What Are the Benefits of the 5 nm Process and its Implementation Challenges?

An integrated circuit (IC), more commonly known as a chip, contains complex circuits composed of a large number of transistors. The gate of a transistor is the narrowest feature on an IC, and its width represents the specification of the lithography process. The smaller width accommodates more transistors per unit area, making it a more powerful chip. However, a smaller width also indicates higher technical requirements.

As one of the first server CPUs built using the 5 nm process, Yitian 710 houses a total of 60 billion transistors, far more than its 7 nm predecessors. The 5 nm process poses extremely high challenges for the energy density and the internal layout of the IC. During the R&D process, the T-Head team flexibly utilized 30 different electronic design automation (EDA) software to deeply customize clock networks and instruction pointer (IP) technologies. In addition, the T-Head team also adopted cutting-edge multitier chip stacking technology to optimize chip performance and power consumption. 04

What Are the Advantages of Chips Developed by T-Head?

Unlike most chip manufacturers, T-Head does not rely on chip manufacturing as its primary business. Instead, the team focuses on developing technologies such as processor IPs, AI chips, and general-purpose server chips. In terms of processor IPs, the core team of T-Head has more than ten years of experience in CPU IP and chip development. The team is involved in the in-house R&D of instruction set architecture (ISA), CPU microarchitecture, and system on a chip (SoC) products. In terms of cloud chips, Alibaba Cloud, one of the world's top 3 cloud platforms, has enabled T-Head to gain a deep insight into the data center business, thereby developing industry-leading chips with higher productivity. At present, T-Head owns the Hanguang 800 Al inference chip and the Yitian 710 general-purpose chip, both of which have achieved record-high performance.





What Are the Benefits of in-house Chips for Alibaba Group?

Chips are the core element of computing systems and the foundation of all Internet applications and innovative technologies. Alibaba Group's ecosystem, which spans across a multitude of fields, including e-commerce, logistics, cloud computing, big data, and globalization, deals with some of the most challenging requirements on computing, networking, and machine learning every day. Successfully tackling these challenges requires a large number of chips, and the use of inhouse chips can reduce the overall computing costs within the Alibaba Group.

Furthermore, with in-house chips such as Yitian 710 and Hanguang 800, Alibaba Cloud can better empower enterprises to innovate by leveraging the power of the cloud. With the release of the Yitian processors and Panjiu servers, more businesses can now enjoy industry-leading computing power anytime and anywhere at a lower cost.

With the launch of "Yitian" and "Panjiu", Alibaba Cloud has perfected the last link of the full-stack cloud infrastructure, realizing technology and architecture innovation and self-research from chips, components to complete machines.

O Tech for Good: XuanTie RISC-V IP Cores Open-Sourced

100

Alibaba Cloud prioritizes cultivating an open, collaborative ecosystem for developing technologies and stretching the limits of imagination to optimize existing technologies and help shape a new digital future. Taking this belief forward, Jeff Zhang, President of Alibaba Cloud Intelligence speaking at the 2021 Apsara Conference, announced that T-Head has made the XuanTie RISC-V series processors open-source and opened a series of related software stack and development tools. Being one of the first global full-stack open-source series of processors, it will help to:

Alibaba Group APSARA

平头哥开源四款玄铁RISC-V处理器

-

Four XuanTie RISC-V IP Cores Open-sourced

25亿_{累计出货} 自研架构玄铁CPU

1986

2.5 billion cores shipped



- Promote the maturity of RISC-V architecture
- Accelerate the fusion and development of new solutions
- Drive innovation to faster implementation

Moreover, Alibaba Cloud plans to open-source more RISC-V processors and hopes to have more partners to develop valuable IP cores and essential software based on XuanTie to build an **open**, **transparent and inclusive RISC-V ecosystem**.



The Anatomy of Open-Source XuanTie RISC-V

With its open and flexible features, RISC-V architecture is anticipated to emerge as one of the most widely used CPU architecture in today's age of Artificial Intelligence of Things (AIoT). However, the current RISC-V architecture faces multiple challenges, including application fragmentation, low development efficiency, and complex adaptation to software and hardware since the software and hardware ecosystem is not mature yet.

The XuanTie RISC-V series processors adopt self-developed technology and are applicable across various scenarios ranging from low power consumption to high performance. Multiple applications include microcontrollers, intelligent household electricaappliances, smart grid, industrial control, image processing, artificial intelligence, multimedia, and automotive electronics. They also support an array of operating systems, such as AliOS, FreeRTOS, RT-Thread, Linux, and Android, broadening RISC-V architecture's vision for an open ecosystem.

It is critical to note that the full-stack open-source processors provide global developers with new architectural choices and also promote the maturity of the RISC-V technology and ecosystem. With years of research around the self-developed technology of chips, T-Head has been persistently demanding for opening it up to the public and making continuous efforts to promote the universalization of computing power.



Three Key Advantages of T-Head RISC-V

The T-Head RISC-V series is quickly gaining momentum worldwide as it offers some distinguished advantages.



10+ years of R&D Experience

With more than ten years of CPU and chip research and development experience, T-Head has been continuously working on self-developed instruction architecture, CPU microarchitecture, and system chip products.



Diverse Product Range & Large Market Occupancy

An extensive product family of T-Head RISC-V processor products ranging from low power consumption and low cost to medium and high performance is used in varied application scenarios, including voice, vision, Bluetooth, wireless, and MCU. T-Head actively promotes development and innovation across the RISC-V software and hardware ecosystem and has been successfully applied in more than 100 chips.

≗₽¤ ╢╢

Forerunner in Technical Strength

XuanTie 910 is capable of reaching 7.1 Coremark/MHz with a primary speed of 2.5GHz. It is compatible with Android and can run applications such as Chrome browser, making it one of the first chips to achieve this using RISC-V architecture

Alibaba Cloud e-Magazine

Open-Source XuanTie RISC-V Ecosystem and Its Significance



The open-source XuanTie series RISC-V processors currently include four mass production processor IPs: XuanTie E902, E906, C906, and C910, and full-stack software and tools based on OpenXuanTie's various operating systems (AliOS, RT-Thread, FreeRTOS, Linux, and Android). Alibaba Cloud opened the source code of the four high-performing RISC-V cores, now available as OpenE902, OpenE906, OpenC906, and OpenC910, respectively.

The open-source XuanTie series RISC-V processors allow developers to access the

XuanTie source code through GitHub and Open Chip Community to build their prototype chips required for IoT applications such as networking, gateway, and edge servers. With this strategic step, Alibaba Cloud envisions accelerating innovation across the thriving RISC-V software community and supporting global developers to build their own RISC-V-based chips cost-effectively. Therefore, now developers can leverage the benefits of open-source EDA collaboration, hardware architecture innovation, and the software application ecosystem enrichment.

T-Head's Open-Source Endeavors

Open-source is undoubtedly a force multiplier for developers worldwide, and T-Head has always been at the forefront of the opensource concept. The latest move to opensource XuanTie series RISC-V processors adds to their previous open-source achievements. The scale and impact of the current initiative across the global RISC-V community are unprecedented in today's connected digital era. T-Head has been contributing consistently to the open-source community over the past. It open-sourced XuanTie 910 as early as July 2019 to enable developers across the globe to download the FPGA code of the processor for free. Eventually, XuanTie 910 became a representative processor for RISC-V to advance to high-end chips. Moreover, with XuanTie 910 gaining Android compatibility recently for the first time, it augments RISC-V's vision for an open ecosystem.

Furthermore, T-Head has successively opened the source code of Wujian 100 Open, a low-

The Significance of Open-Source for Businesses

The rapidly evolving RISC-V architecture technology necessitates more collaboration among developers and enterprises worldwide. Open-source ecosystem promotes ecological progress, ensures faster technological maturity and stability across the software and hardware domains, and benefits every partner in the ecosystem.

Moreover, with the rising wave of AloT, RISC-V



power consumption, micro-control chip design platform, along with a complete set of related modules, such as processor, operating system, basic interface IP, software driver, and development tools.

T-Head continues to expand its open-source endeavors, and its research and efforts in the RISC-V domain are imperative to the promotion of the AloT strategy. Open-source XuanTie RISC-V series processors are set to drive the RISC-V ecosystem development, promote the Internet of Everything, and strengthen the global RISC-V community.

may be the key to the Internet of Everything (IoE) and a critical factor to broaden the Internet of Things ecosystem. The opensourced XuanTie RISC-V processor series effectively provides technology and service updates to enterprises moving from opensource to actual. It will enable the growth of the "cloud-terminal integration" of the AloT ecosystem with developers.

05 Hyperscale Data Centers **Powered by Green Technol**ogy

Our green data centers represent Alibaba Cloud's relentless effort to achieve a harmonious balance between technology, society, and the environment.

Let's look at the six interesting facts.

1. Increasing Adoption of Clean Energy in Alibaba Ecosystem

From 2015 through 2021, we have built more than 5 green data centers across China to support the development of the new digital infrastructure.



2021 11.11 Shopping Festival witnessed a further reduction in the carbon footprint. 26,000 tonnes of carbon dioxide emission will be reduced through the usage of clean energy.



5. Energy Saving "Soaking" Servers

6. Reduced Carbon Emissions by Harnessing Wind Energy

Alibaba Cloud's data center in Zhangbei County in Northwestern China leveraged wind power and reduced 8,000 tons of carbon dioxide during the 6.18 sales.

2. Combining Innovation and Energy Saving

Alibaba Cloud's Inner Mongolia data center also increased the usage of wind and photovoltaic power to 45% of its energy mix in 2021 vs 38% in 2020.

3. State-of-the-Art Hangzhou Data Center

Compared to traditional data centers, Alibaba Cloud Hangzhou data center saves 70 million kilowatt-hours of electricity each year. It is equivalent to the power consumption of all street lights around the West Lake continuously lit for 8 years.



2022 Edition 8

4. Propagating Clean Energy During the 11.11 Shopping Festival

Alibaba Cloud data centers save more than 70% of energy with the rare "soaking server" cooling technology and reduce the average annual PUE to 1.09.

06 Lingjie: Integrating Digital **Intelligence** with the Cloud **for the** Next Wave of AI



Yangqing Jia

Vice President of Alibaba Group, Senior Fellow of the Computing Platform BU, Alibaba Cloud Intelligence, Head of the System for Al Lab (SAIL), DAMO Academy

Yangging Jia leads the Computing Platform Business unit, responsible for Alibaba's big data platform (MaxCompute), Platform for AI (PAI), and other products. These products support the large-scale data storage, compute, analytics, and machine learning needs within Alibaba Group and Ant Group, as well as powering the digital transformation of customers on the Alibaba Cloud, spanning multiple industries. Before joining Alibaba, he led the deep learning project Caffe, which won the Mark Everingham Prize from the Computer Vision Foundation for its outstanding contribution to democratizing deep learning. He is a co-author of TensorFlow, co-leader of PyTorch 1.0, and the initiator of the ONNX standard. He received his bachelor's degree and master's degree in Automation from Tsinghua University and his Ph.D. in Computer Science from the University of California, Berkeley.



We believed that algorithms could change everything. However, recent statistics show that nearly half of Al projects struggle to go from the lab to the final production stage. The age of slash-and-burn artificial intelligence (Al) projects has already passed, and many scientists and Al engineers are now looking for more efficient ways to bring concepts to life. Today, we are increasingly seeing the integration of computing power, algorithms, data, and scenarios, especially the integration of Al and big data. But what exactly do we need to do to spark the next wave of artificial intelligence?

What Paradigm Should We Use to Explore Big Data and Al Systems?

Let's start with the end goal in mind and summarize our paradigm with four "S" - Scale, Speed, Simplicity, and Scenario.

With the ever-increasing volume of business data and the growing need for new agile business models, it is immediately obvious that scale plays an important role in the development of Al systems. At the same time, Al and big data research are still growing explosively; Moore's Law can hardly keep up with today's computing needs. Therefore, the second problem should focus on how we can speed up the development of systems.

Thirdly, to overcome the bottleneck of moving from the lab to the production stage, we need a standard and simple approach to achieve efficient deployment, effective operations and maintenance, and standard API and SDK. Lastly, we often see gaps between models developed in labs and actual business requirements. Thus, another critical problem is how to bridge these gaps to ensure intelligent, scenario-based decision-making. According to the concept of "4S", these are the problems that big data and AI must address.

Lingjie: Alibaba's Integrated **Big Data and Al Portfolio**



Alibaba Lingjie is our newly launched integrated AI platform that is aimed at addressing these issues. Lingjie integrates various components to provide users with an "out-of-the-box" solution to serve a wide variety of business applications. This includes infrastructure, big data and Al platform, Al open services, and AI applications.

Lingjie is the product of our engineering, algorithm, and business teams joining hands and working together to solve the fundamental issues of AI research and development. We hope it will open the doors to more possibilities.

Our Breakthrough Achievements in Al

Let's explore some exciting things we have done with Lingjie. The first was the superlarge-scale model obtained through largescale step-by-step training. Scientists at Alibaba's DAMO Academy worked with a model called M6. It is currently one of the largest multimodal models in the Chinese developer community. When AlexNet first came out in 2012, its model size was about 60 million parameters. But, M6 could reach 10 trillion parameters at its maximum level. Therefore, M6 can be applied to a wide range of scenarios, including designing customized products for intelligent manufacturing. Sure enough, AI may not be able to complete

all end-to-end implementations; it can however allow designers to spend more time innovating by reducing menial and repetitive tasks. These capabilities can be easily scaled to support commercial scenarios such as manufacturing. Furthermore, these capabilities can also be extended to non-commercial scenarios, including helping scientists run better, more accurate, and faster simulations for healthcare and environmental research.

Now, let's move on to another interesting example. As we mentioned above, the capabilities brought about by Lingjie empower researchers in the medical and healthcare

industry to achieve more breakthroughs. In this field, one of the biggest challenges is understanding or simulating the relationship between drugs and proteins. On the one hand, it is domain-based knowledge, and on the other hand, it requires massive amounts of data and calculation. It is hard to imagine biologists and chemists purchasing machines and running simulations all by themselves. Therefore, to solve such a problem, we developed a high-performance computing

The Next Wave of Artificial Intelligence

So, if everyone could become an Al developer, what would you do? Customers using big data and AI solutions on the cloud today may be able to provide a perspective to this question. Every industry has its own requirements, and through close collaboration with our customers, we understand that cost reduction and efficiency enhancement are among the top concerns in virtually all industries. But more importantly, customers are eager to solve a wide variety of problems effectively through standardized and replicable solutions.



cluster called Eflops. This platform is aimed at solving the problem of predicting drug activity using AI and has been tested in labs outside of Alibaba.

Of course, Al also brings a lot of changes to our lives in many spheres. We want to make the AI platform cloud-native in order to close the distance between a developer's idea, from the first line of code, and the final Al application.

We developed the Lingjie AI platform particularly to cater to these demands. Lingije was not built just for the sake of innovation, it was the result of us wanting to make Alsolutions easier to be realized. It is designed based on our experience accumulated over the years, focusing on Scale, Speed, Simplicity, and Scenario.

We hope that Lingjie will help developers from all walks of life be able to trigger the growth of new and innovative AI applications on the cloud. We hope we can inspire the next wave of artificial intelligence.

07 Six Technical Directions of Next-Generation Enterprise Databases



Feifei Li

Vice President of Alibaba Group, President of Database Products Business Unit, ACM Distinguished Scientist

Feifei Li is currently the Vice President of Alibaba Group, ACM Distinguished Scientist, President of the Database Products Business Unit of Alibaba Cloud Intelligence, and Director of the Database and Storage Lab of DAMO Academy. He has won multiple awards from NSF, ACM, IEEE, Visa, Google, HP, Microsoft, and IBM. He is a recipient of the ACM SoCC 2019 Best Paper Award Runner-Up, the IEEE ICDE 2014 10 Years Most Influential Paper Award, the ACM SIGMOD 2016 Best Paper Award, the ACM SIGMOD 2015 Best System Demonstration Award, and the IEEE ICDE 2004 Best Paper Award. He has been an associate editor, PC co-chairs, and a core committee member for many prestigious journals and conferences. He has led the R&D efforts of cloud-native database systems and products at Alibaba.

Gartner recognizes Alibaba Cloud as a Leader in the 2020 Magic Quadrant for Cloud Database Management Systems (DBMS). It implies that the Chinese database pioneer is at the forefront of cloud database development among the world's best.

Gartner has now combined Operational Database Management Systems (OPDBMS) and Data Management Solutions for Analytics (DMSA) into Cloud Database Management

Development Trends of Database Technologies in the Past Ten Years

The evolution from traditional database architectures to cloud-native architectures was the most prevalent trend of database technologies in the last decade.

The rise of cloud service providers was a typical feature during the past decade. The growth and success of many cloud service providers resulted from the opportunities brought by the emerging technologies of cloud computing.



Systems (DBMS) because the global research giant believes that "There is Only One DBMS Market." The move implies that the Cloud DBMS Magic Quadrant is more competitive and has a higher value, and lays the ground for a trend that calls for databases and data warehouses integration.

As the database industry evolves rapidly, let's explore six technical trends of the nextgeneration enterprise databases.

The system of cloud-native technologies derived from cloud computing had given birth to cloud-native databases and cloudnative data warehouses, such as Alibaba Cloud PolarDB and Alibaba Cloud AnalyticDB.

Another dominant trend was the growth of distributed technologies, which have evolved from their rudimentary forms over the past decade, making distributed databases and data warehouses possible.

Six Trends of Enterprise Database Technologies for the Next Decade

First, cloud-native and distributed technologies will be further integrated into a seamless architecture to provide better elasticity and high availability.

Second, the profound application of intelligent technologies. This involves using Al and related technologies to achieve intelligent O&M of databases, such as index recommendations, MySQL governance, and exception detection.

Third, integration of database and big data technologies, including hybrid transaction/ analytical processing (HTAP) and integration of online/offline resources. In the past decade, databases and big data services were separate. Databases were all about offline processing, whereas big data supported online services. However, customers tend to use one system to manage the entire data application process, covering data generation, processing, storage, and final consumption. It's an increasingly urgent need to reduce the costs of data migration and storage. Customers hate synchronizing data every day. Integration of online and offline resources, HTAP, or online query and offline computing capabilities can help resolve all these problems in no time. This will be an essential trend in the next decade. That's why Gartner has combined the initially separated **Operational Database Management Systems** (OPDBMS) and Data Management Solutions for Analytics (DMSA).

Fourth, multi-model database. How to process unstructured or semi-structured data such as text and images along with structured data, is critical. Integrated multimodel databases can help to deal with unstructured or semi-structured data.

Fifth, hardware and software integration. The development of hardware, such as non-volatile memory (NVM) and high-speed networks, plays a critical role in designing database systems.

Last but not least, security and reliability. Such issues are not an emerging trend but will continue to evolve. For example, it's crucial to figure out how to build a blockchain into a database system to make it tamper-proof or protect databases by using encryption technology.



Industry Trends: Key Technologies for Next-Gen **Enterprise Databases**

HTAP: Big Data and **Database Integration** Hybrid Row/Column-Stores + Mixed Workload + Distributed **Computing and Analytics** Intelligent **Technologies** 00 + Self-Healing + Self-Tuning Software and Hardware Integration ARM CPU & Alpha CPU

Alibaba Cloud Database Services: The Way Ahead

Alibaba Cloud will continue concentrating on varied fronts, including:

- Two core services for online transaction processing (OLTP): PolarDB and its distributed variant PolarDB-X.
- Core services for online analytical processing (OLAP): AnalyticDB, a cloudnative data warehouse.

Two core services in the NoSQL field:

Lindorm, a cloud-native multi-model database, and Tair, a cloud-native inmemory database.

In addition, the cloud pioneer will also provide managed RDS and NoSQL databases and tools. Regarding managed services, we tend to build a management platform,

Cloud-Native and Distributed Architecture

Decoupling of Compute, Memory Processing

Multi-Model Databases

Graph/Time-Series/...



ረዝቲ

Security and Reliability

SQL Auditing + End-To-End

which combines the strong points of cloudnative and intelligent technologies to provide premium management services, such as automatic instance management and high availability.

We also want to focus on our core competencies, increase investment in key in-house services, provide a management platform, and cooperate with ecosystem partners. For example, instead of investing in the development of MongoDB, we have signed an agreement with MongoDB to use its latest version. From the development of database kernels, provision of O&M services to application development as an Internet service provider (ISV), our core business strategy has always been to leverage and increase the power of an ecosystem.

FERTINET. Fortinet on Compute Nest

Advanced Threat Protection for Your Peace of Mind

Why Fortinet on Compute Nest

Automate comprehensive security architecture with simplified management, consistent visibility, and industry-leading threat intelligence. Get the best of both Alibaba Cloud and Fortinet, and enjoy our limited-time offer with FREE POC!



The Convergence of Network and Security

Fortinet is the only vendor recognized as a Leader for Network Firewalls and WAN Edge Infrastructure Using Single Platform and is committed to being a trusted advisor and partner for network and security.

Cloud-Native Integration for Efficiency

Fortinet security products have been integrated with Alibaba Cloud platform and services natively, allowing you to manage network and security policies on FortiGate in a cloud-native manner.





Industry Validated and Recognized Solutions

Fortinet Security Fabric, the industry's highest-performing cybersecurity platform, was named as a representative solution of Gartner's Cybersecurity Mesh Architecture.



Featured Products by Fortinet

FortiGate	FortiWeb	
	Our Cus	
	sodexo	

Contact our sales representative for migration support, FREE POC, and limited-time special offer! https://www.alibabacloud.com/campaign/fortinet

FortiManager

FortiAnalyzer

tomers

FUNG GROUP



008 Knocking on the Door to Cloud-Native Database 2.0



David Chen Senior Technical Expert of Database Tools and Utilities at Alibaba Cloud We are witnessing a new era of technological and industrial revolutions that is restructuring the global development scenario. As represented by cloud computing, nextgeneration information technologies are being interpenetrated as the new engine that drives the intelligent and innovative development of the digital economy.

On the other hand, enterprises need to modify cloud databases for the post-pandemic era. COVID-19 has accelerated the migration of businesses across industries to the Internet for digital operations. Along with this rapid transition, the demand for data storage and value mining is on the rise, resulting in three distinct trends:

- Rapid data volume increase
- Faster data migration to the cloud
- The use of a cloud-native and distributed architecture

What Disruptive Changes Will Cloud-Native Database 2.0 Concepts and Services Bring to Enterprises?

An evident change is that enterprises can manage a copious volume of data in databases, such as standalone databases and cloud-native distributed databases like PolarDB, and import the data to data warehouses for analysis, reporting, archiving, and offline data value mining. Enterprises can develop storage and computing capabilities that integrate databases, data warehouses, and online and offline resources.

Another change is the collective management of online data assets, and enterprises can build data warehouses for agile analysis at lower costs. This is because cloud-native technologies enable the seamless real-time



Do Database Engineers Need to Learn New Technologies to Use These Tools Well?

In the past, application developers worked with databases, and data engineers or scientists worked with data warehouses. Alibaba Cloud Data Management Service (DMS) enables developers to process data like data engineers by connecting databases and data warehouses. The platform has lowered the requirements for skills. connection of databases and data warehouses in a controlled environment, facilitating O&M activities such as DDL-based changes and capacity scaling without interruption to business. Such a feat is impossible for a traditional data integration solution.

In addition, the system includes an extract, transform, and load (ETL) procedure that substantially speeds up the traditional data integration process, and the source database acts as the data warehouse's operational data store (ODS) for computing. Physical data migration is no longer necessary. Enterprises can build data warehouses as needed, like in no time.

In addition, small and micro-businesses, or medium and large enterprises, can all manage their data and development in DMS. Also, because the storage resources of a cloud-native data warehouse are infinitely expandable, the model for computing and processing user data barely changes with the swelling data volume.

Is Open-Source the Future of Cloud Databases?

Alibaba Cloud thrives on open-source databases and has been an active participant and promoter of open-source communities. We constantly share practical problems, bugs fixed, and original techniques. This time, we have disclosed some of the core, value-added technologies to return the favor.

Based on the long-term practice of providing

cloud database services, Alibaba Cloud has developed an insightful understanding of database technologies, and we have mastered our knowledge in terms of high availability, distributed architectures, and storagecomputing separation. We are looking for more co-builders to jointly provide more valuable solutions for users by launching open-source projects.



A Promising Future for Cloud Databases with Alibaba Cloud

First, apply the cloud-native distributed database technology. The cloud-native distributed technology can support online processing and analysis of large volumes of data, ensure high service availability, and provide the foundation for multi-model data processing engines. As a result, cloudnative distributed databases will become the preferred option for businesses.

Second, use DMS to manage enterprises' data assets. DMS enables collective data masking and lineage analysis to help bring out the value of data. It provides enterprises with services that cover the entire lifecycle of data, from data production and storage, transmission and processing, to computing

and analysis. Enterprises can lower the barrier for data processing and analysis and harvest the benefits only by seamlessly connecting their various databases, allowing free data flows, and developing some comprehensive solutions.

Due to the rapid increase of business data, different data types, more complex business characteristics, and the demand for real-time data insights, traditional database solutions are required to be more scalable, with better availability and real-time flexibility. We have seen too much of this, just like a calling from the digital era. With world-class proven products and services, Alibaba Cloud can be your right choice to achieve such transition.

09

Apsara Luoshen – A **Cloud-Native Network Born from the Cloud**



Zhu Shunmin Head of Alibaba Cloud Network Services

According to its narrow definition, cloud-native entails microservices, containers, continuous delivery, and DevOps. However, I believe that the wide application of cloud computing is redefining cloud-native, which may become more inclusive. In a broad sense, cloud-native simply refers to any software, hardware, or architecture intended for cloud use. As a typical and core Infrastructure as a Service (laaS) product of cloud computing, the cloud network was born from the cloud and for the cloud. In fact, it is the real cloudnative network.

Cloud-Native Apsara Cloud Network System

Apsara Cloud Network System is an in-house cloud-native system built for the cloud through the support of other cloud capabilities. It represents the unique way we, at Alibaba Cloud, develop cloud-native technologies. Apsara Cloud Network System encompasses capabilities such as large-scale software-defined network (SDN) control, Sailfish software-hardware integrated forwarding and CyberStar elastic



Alibaba Cloud Network Architecture



In-house SDN Controllers: Removing Concerns Over Size and Complexity

SDN controllers play a key role in the efficient management of a large data system. In the trend of cloud-native applications, a great number of users are using Docker and Kubernetes in the cloud. It poses a huge challenge to the size and density of cloud networks. SDN controllers enable users to accommodate large amounts of instances in a virtual private cloud (VPC) and meet their



In-house Sailfish Software-Hardware Integrated Forwarding Platform: Removing Concerns Over Performance

The Sailfish hardware forwarding platform is built on a hardware-software integrated architecture based on high-performance forwarding components such as dedicated programmable chips and intelligent network interface controllers (NICs). The platform needs for highly densified elastic network interfaces (ENIs) on an Elastic Compute Service (ECS) instance. Also, given the scale of public cloud instances, management efficiency is critical. Cloud-native SDN controllers are built to take care of these issues and relieve users from worries about their data systems' massive size and complexity.

is designed to enable high-performance data forwarding between hosts and primary gateways. The platform has been widely deployed in many cloud network applications in combination with core modules such as XGW, MoC, and ALI-LB.

XGW: In-house High-performance Gateway with Integrated Software and Hardware

The development of network devices often involves both software and hardware because they work together in a complementary way toward higher flexibility and performance. Before the birth of XGW, specific customers requested a hybrid cloud by using connections over Express Connect circuits with a jaw-dropping bandwidth of 10 Tbit/s. Considering reasonable redundancy, nearly a hundred high-performance x86 servers were required, leading to unacceptable costs and operational efficiency. It encouraged us at Alibaba Cloud to develop XGW, a high-performance gateway with integrated hardware and software.



CyberStar: Enabling Flexible Customization and Making Investment More Effective

As more large enterprises migrate their business to the cloud, they often require higher elasticity to support scenarios such as live streaming and large promotional events. These scenarios frequently require a drastic capacity increase that traditional architecture built on x86 servers cannot meet.

Therefore, the elasticity based on the cloudnative ECS service becomes a natural choice. The most appealing feature of CyberStar, an elastic network element platform, is that the network elements are not deployed on bare metal servers or in resource pools dedicated to computing. Instead, they get deployed on Alibaba Cloud ECS instances, which means "infinite" resources and "infinite" scalability. As the business network elements no longer get deployed on traditional x86 physical servers, scalability problems such as a timeconsuming deployment and complex scaling procedure have been solved. The CyberStar platform is now widely applied to core services such as Application Load Balancer (ALB), NAT Gateway, PrivateLink, CEN-TR, Global Accelerator (GA), and VPN Gateway.

Future Outlook

On the 2020 IDC FutureScape webinar, IDC predicted a significant trend toward clouddedicated devices. By 2024, more than 25% of new workloads on public clouds will get processed by dedicated infrastructure components from service providers. I believe this day will arrive much earlier because cloud-native dedicated devices built on integrated hardware and software have already become a popular choice in cloud



network-related business scenarios. Likewise, the CyberStar elastic network element platform built on the cloud-native ECS service has also been deployed to cope with a wide range of scenarios.

In the future, cloud network technologies will continue to evolve in these two directions, and the core advantages of cloud networks will rely on cloud-native architectures.

METAVERSE

A PRIMER ON THE NEXT GENERATION OF THE INTERNET

Connect people and objects in the virtual world to reality by integrating digital and physical worlds, digital and real identities, and digital and real-world economies.



KEYTECHNOLOGY PILLARS OFTHE METAVERSE



Holographic Construction

Constructs a geometric model of people, maps, and objects in the virtual world to create an immersive environment.





People and objects in the virtual world simulate the physical world dynamically and follow the physical rules of our reality to construct their world.





Blends the information from the virtual and physical worlds, builds high precision 3D maps using AR and mixed reality to achieve accuracy.



Virtuality-Reality Linkage

Transforms the physical world using second-generation robots, like AlibabaTianxun, a complete solution to unattended monitoring, inspection, and data security.

Construct Your Own Metaverse Leveraging Alibaba Cloud Technology

Leverage Alibaba Cloud's advanced technological solutions, built using Artificial Intelligence (AI), and Machine Learning, to create your own virtually-unique Metaverse.



Supports fast and high-quality construction of metaverse material and digital content with Alibaba Cloud Elastic High-Performance Computing (E-HPC) power.



nternet of Thinas (loT)

Offers reliable wireless connectivity and signal transmission services to ensure seamless virtualityreality integration of Metaverse.



Provides enterprise-level data privacy, high-quality networking, and smooth-user experience for the Metaverse.



Contact us to take your smart technology platform to the next level with Metaverse!

2022 Edition 8



Provides an integrated view of virtual and physical worlds, enhances the performance of the Metaverse via intelligent speech, Machine learning, OSS, and NLP.



chain as a

Helps users create a secure, stable, and valuebased environment on the Internet with efficient collaboration and mutual trust.



Alibaba Cloud Contair Service for Kubernete

Integrates virtualization, rich scheduling, deployment, and security capabilities to help users build metaverse applications quickly.





42