

## USE CASE

## SONiC Alibaba Data Center Network Use Case

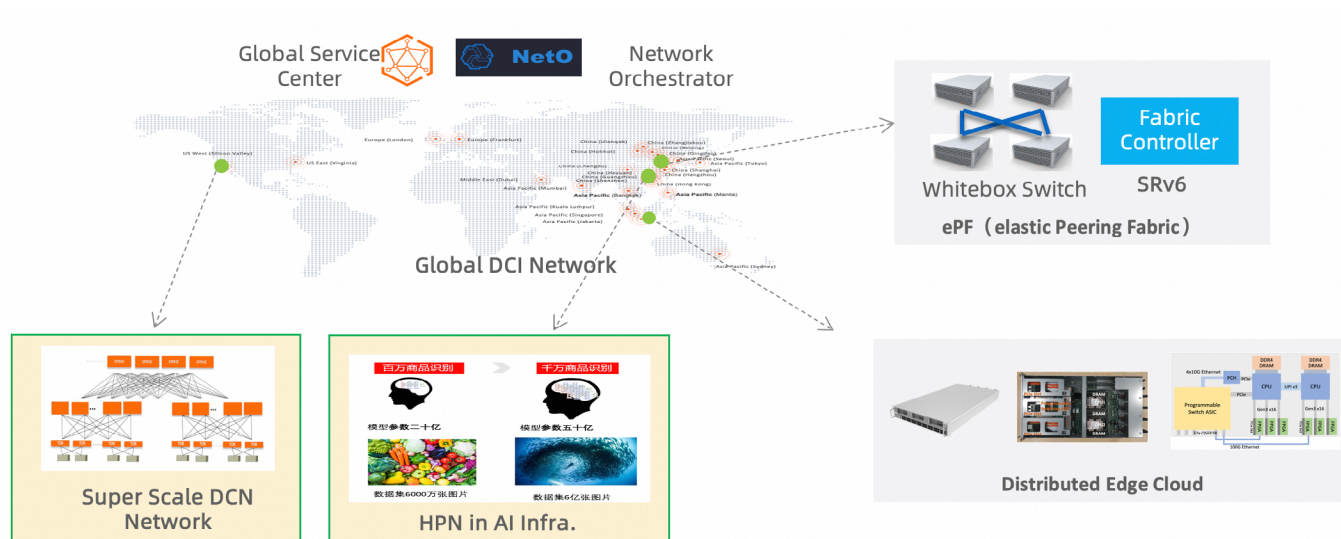
## Overview

Alibaba Cloud's worldwide infrastructure spans 28 regions, providing a combined total of 86 available zones and an extensive network featuring over 3,200 CDN nodes.

The cornerstone of this cloud infrastructure is built on Alibaba predictable network infrastructure. The overarching strategy behind the development of the Alibaba network revolves around the utilization of a unified protocol stack, IPv6/SRv6, alongside single-chip

white boxes which the control plane is based on SONiC and the data plane is programmable. To establish a predictable network, it is important to strike a harmonious balance between optimizing both endpoints and the network infrastructure. Moreover, it's crucial to foster a collaborative design that bridges applications and the network, resulting in an endpoint-network synergy system.

Figure 1: The big picture — AliCloud network



SONiC **NPL** High performance network for AI/ML based on merchant silicon white

Due to SONiC's openness, maturity, strong community support and other various benefits, Alibaba chose SONiC as the data center devices' network operating system, which allows us building up an endpoint-network synergy system easily. In this context, we demonstrate how Alibaba harnesses SONiC to enhance its data center infrastructure and align it with the demands of AI infrastructure.

## Deployment Details

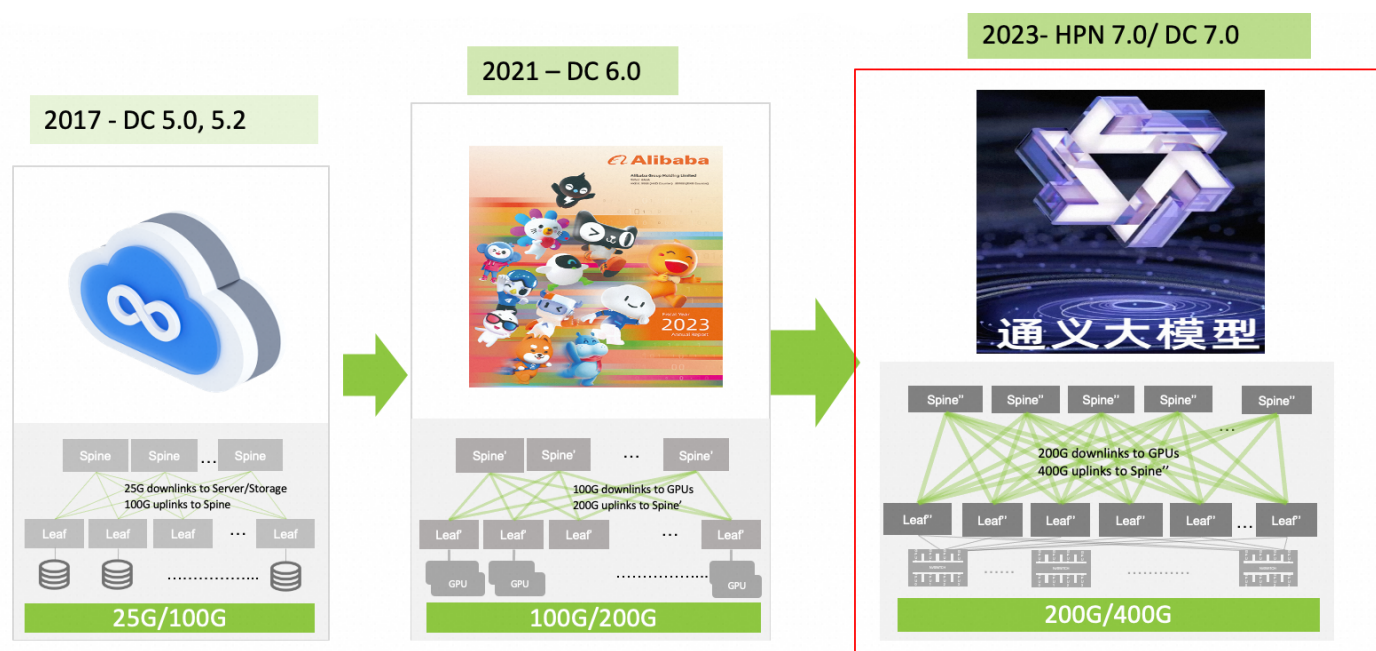
Alibaba embarked on its SONiC journey in 2017. The subsequent diagram illustrates the evolution of Alibaba's data center infrastructure as it progressed along with its SONiC journey.

At present, Alibaba Cloud has deployed over 100,000 white box network devices, all operating on the SONiC-based software stack as their network operating system. Beginning in 2017, Alibaba's data center network architecture has been optimized for compute and storage

use cases. Each PoD follows a CLOS architecture, with each TOR switch utilizing 25G downlinks to connect to servers and storage devices, and 100G uplinks to connect with spine switches. This architecture has evolved through two generations, known as DC5.0 and DC5.2. Alibaba introduced the Solar RDMA technology, which was presented in a paper at Sigcomm 2022, integrating storage logic into the network pipeline to enhance storage performance.

As the demand for computing increased, Alibaba initiated the development of its own AI infrastructure using SONiC-based white box devices. The DC6.0 architecture was defined for this purpose, catering to both AI infrastructure and standard cloud Elastic Compute Services (ECS) requirements. In the case of AI infrastructure, each TOR employs 100G downlinks to connect with GPUs and 200G uplinks to connect with spines, and this AI infrastructure played a pivotal role in supporting the 2023 Asia Olympic Games.

Figure 2: The hotspot — AI infrastructure with AliCloud



With the growing popularity of ChatGPT and the varying requirements of AI infrastructure and Data Center Networking, Alibaba's DCN has evolved into two distinct architectures. The HPN7.0 cluster is dedicated to large-scale AI infrastructure, utilizing 51.2T switches for both TORs and spines. Meanwhile, DC7.0 represents the next generation of cloud compute / storage infrastructure. The upcoming slide will provide an overview of the HPN7.0 architecture.

HPN7.0 deviates from the conventional CLOS architecture. At the TOR level, it features a configuration where 16 TOR switches form a segment. This configuration allows for the connection of up to 1000 GPUs in one segment. Each TOR segment connects with two independent spine planes, with each plane accommodating as many as 64 spine switches. This innovative architecture empowers Alibaba's AI infrastructure to efficiently address the needs of extensive AI deployments in just two tiers, resulting in reduced network latency and streamlined traffic management within the AI fabric.

## Benefits

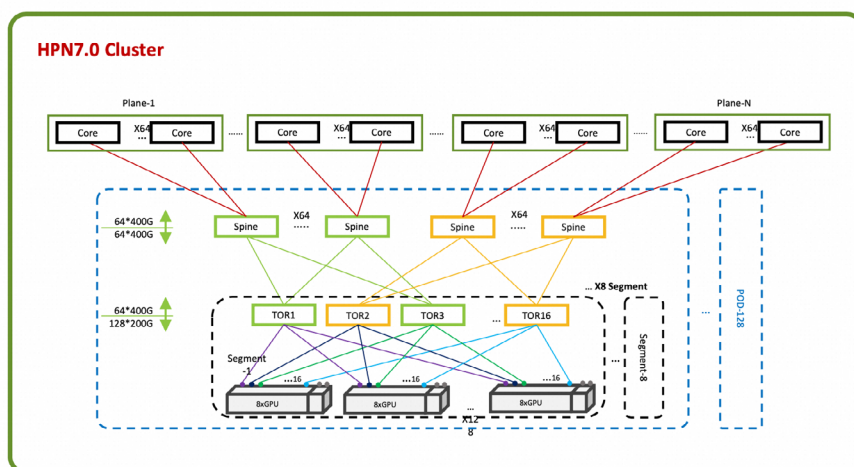
Alibaba has complete autonomy over our SONiC-based network operating system, allowing us to customize it to meet our precise needs. This enables us to concentrate our resources and efforts on features that are specifically critical to us, eliminating the necessity to invest in or develop extraneous components.

This approach offers several notable advantages:

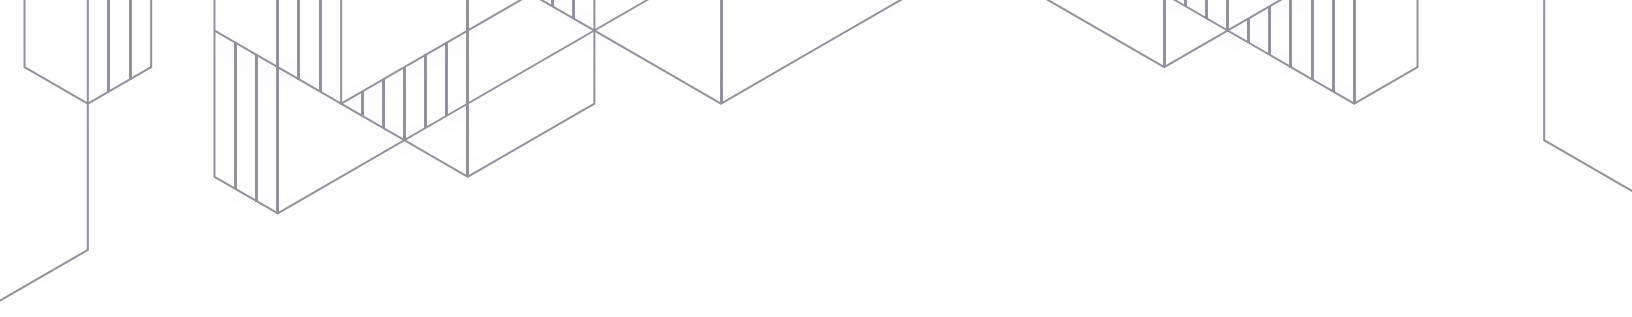
- **Substantial Cost Savings and Seamless Infrastructure Management:** Alibaba achieved significant cost savings, including both initial capital expenses (CAPEX) and ongoing operational expenses (OPEX), by deploying SONiC-based network operating systems across a diverse range of white box devices, including traditional data center switches, optical switches, DCI routers, and gateway devices. By consolidating under a single network OS, Alibaba's cloud network infrastructure becomes seamlessly monitorable, manageable, and extensible. SONiC's versatility allows rapid detection

Figure 3: High performance network of Alibaba cloud

- High performance RDMA network
- Large Radix scale out
- Multi-rail communication
- Adaptive Routing ready
- Seamless integration between SONiC and in-house automation systems



Our approach to addressing **AI emergency!**



and resolution of network outages, typically within just 60 seconds, effectively mitigating and preventing potential disruptions or impacts within a maximum timeframe of 10 minutes.

- **Enhanced Network Quality Through Efficient Testing:** SONiC empowers Alibaba to achieve enhanced network quality through targeted testing efforts. With a deep understanding of how our devices are deployed and utilized, Alibaba focuses development and testing efforts exclusively on what is truly necessary for our unique requirements. This means that we may require only 10-20% of the effort that normal vendors do to achieve the same objectives within our specific area of interest.
- **Agile Update Cycles for Up-to-Date Infrastructure:** Alibaba maintains rapid update cycles to keep its network infrastructure current. Even with focused testing efforts, software updates are sometimes unavoidable during the device's life cycle due to factors such as enhanced requirements, shifts in deployment strategy, or issues discovered in the field. With full control over our SONiC-based network operating system, we can implement necessary changes at our discretion without external dependencies.
- **Accelerated Introduction of New Features and Devices:** SONiC's open and mature architecture enables Alibaba to accelerate the introduction of new features and devices, fostering agility and innovation in our network capabilities. Beginning in April, Alibaba dedicated all development resources to the creation of the 51.2T switch for HPN7.0. Despite starting from scratch, we successfully deployed new devices to handle our AI traffic in less than four months. We attribute this nimble development to SONiC's adaptability and openness.

**“The future of SONiC is highly promising, owing to its adaptability, open-source collaborations, cost-effectiveness, and the dynamic community that supports it.”**

- **Thriving Community and Collaborative Innovation:** Alibaba benefits from the thriving SONiC community, fostering a rich environment for sharing ideas, resolving issues, and pushing the boundaries of network technology. As more organizations engage with SONiC, the community's collective knowledge and impact will grow. Alibaba not only reaps the benefits of SONiC's versatility but also contributes enhancements to push the boundaries of SONiC. Currently, Alibaba leads the SONiC routing working group, OTN work group to introduce innovations in the SONiC routing domain and optical switch domain.

## Conclusion

The realm of networking is undergoing swift transformation. This year has marked the onset of AI infrastructure, catalyzed by ChatGPT, which presents a prime opportunity for the advancement of network technologies. SONiC emerges as a robust platform for swiftly incorporating these innovations. The future of SONiC is highly promising, owing to its adaptability, open-source collaborations, cost-effectiveness, and the dynamic community that supports it. Alibaba Cloud remains committed to the ongoing evolution of our network infrastructure to ensure predictability. We will collaborate closely with the SONiC community to drive continuous enhancements in SONiC.



# Join SONiC

Become a SONiC member to collaborate, learn and shape the future of the Open Network Operating System.

[sonicfoundation.dev/join-sonic](https://sonicfoundation.dev/join-sonic)