Aspirations to Configure a High-speed Network Connecting GPU Server Clusters that Support the AI Era Using Open Technology/Systems

Highlights

Introduction/Architecture Points

- Realize a simple multi-tenant environment via switch architecture
- Evenly share bandwidth among tenants
- Provide consistent latency with the required GPU server bandwidth and no bottle necks

The demands for GPU computing are growing with the rapid evolution of generative and other AI technologies. GPU technology is also advancing at a rapid pace to keep up with the demand for high-performance. Building out large-scale computing processes to leverage numerous GPUs is difficult however, without setting up suitable facilities that can cope with the greater power consumption and heat that high-performance computing requires. GPU computing services through cloud/data center providers are the solution to satisfy these needs. SAKURA internet provides GPU servers under the name Koukaryoku. The organization plays a vital role in delivering the infrastructure necessary to compete in the fierce AI development happening worldwide. SAKURA internet chose Arista Networks' DCS-7800-series switches for the network infrastructure that connects a large number of GPU servers in a cluster.

Customer Data

SAKURA internet Inc.

Address: 1-12-12 Umeda, Kita-ku, Osaka-shi, Osaka-fu

Established: December 1996

SAKURA internet

SAKURA internet came into fruition as a server rental service for private users in 1996 during the dawn of the internet age, before ‘cloud’ computing had come to mean what it does today. The Company began with a vision to turn ‘What you want to do’ into ‘What you can do’ with the aim to create the internet and a society where the dreams of passionate people willing to persevere can be realized. SAKURA internet is proud of its track record of being engaged with pioneering initiatives. In January 2016, the Company announced the start of efforts to provide specialized services delivering computing power under the “Koukaryoku Computing” concept, which responded to the demand for the tremendous computational resources necessary for deep and other machine learning. In anticipation of higher computational demand thereafter driven by AI, SAKURA internet aimed to further expand its GPU services according to the market needs. In June 2023, the Japan Ministry of Economy, Trade and Industry approved plans to secure cloud programs as specific critical products under the Economic Security Promotion Act. In light of this announcement, SAKURA internet decided to develop Koukaryoku cloud services for generative AI recognizing the provision of computational resources for AI as essential infrastructure for fostering a digital society in Japan.

On January 31, 2024, the Company started offering Koukaryoku PHY bare metal servers as a cloud service for regenerative AI provided through bare metal servers equipped with NVIDIA H100 GPUs.

SAKURA internet Koukaryoku Venture

We talked with Takefumi Sudo, Assistant Director of the Cloud Service Department under the Cloud Service Business at SAKURA internet Inc. who oversees the Koukaryoku services,about the launch of these comprehensive Koukaryoku services. He explains, “I integrated GPUs into our services after Founder and President Kunihiro Tanaka asked whether or not we could incorporate GPUs as a service or business.” The biggest drive for these services was the GPU power consumption. “I knew customers would need to take advantage of data centers if GPU servers consumed more power than typical processes due to the difficultly of setting facilities like that up on site. The next question became what form and type of interface would we need to provide these computational resources through our data center. I thought about how we could package these services as a product, which happened to fit very well with advancements in machine learning technologies, the birth of generative AI, and other such technologies.” Takefumi Sudo reflects.

The Koukaryoku PHY servers provided by SAKURA internet today are located at the Ishikari Data Center.

As a suburban data center that began operations in November 2011, the Ishikari Data Center is a high-efficiency facility that takes advantage of the cold Hokkaido climate by broadly adopting a method to use outside air for cooling. The full workload of regenerative AI and other processes currently carried out on these high-performance GPUs takes a tremendous amount of data to repeatedly run an enormous number of calculations. Unlike interactive applications that demand a meticulous exchange of data with good responsiveness, the long distance to the data center causes almost no lag. The use of ambient air cooling can also tackle the large amount of heat produced by the GPUs at a low cost. These are two of the major benefits. A broadband network is indispensable for group communication between GPU servers without any lag necessary to take full advantage of the latest GPU computing performance. SAKURA internet decided to use the Arista Networks DCS-7800 series for the network infrastructure of the large-scale GPU cluster that connects numerous GPU servers.

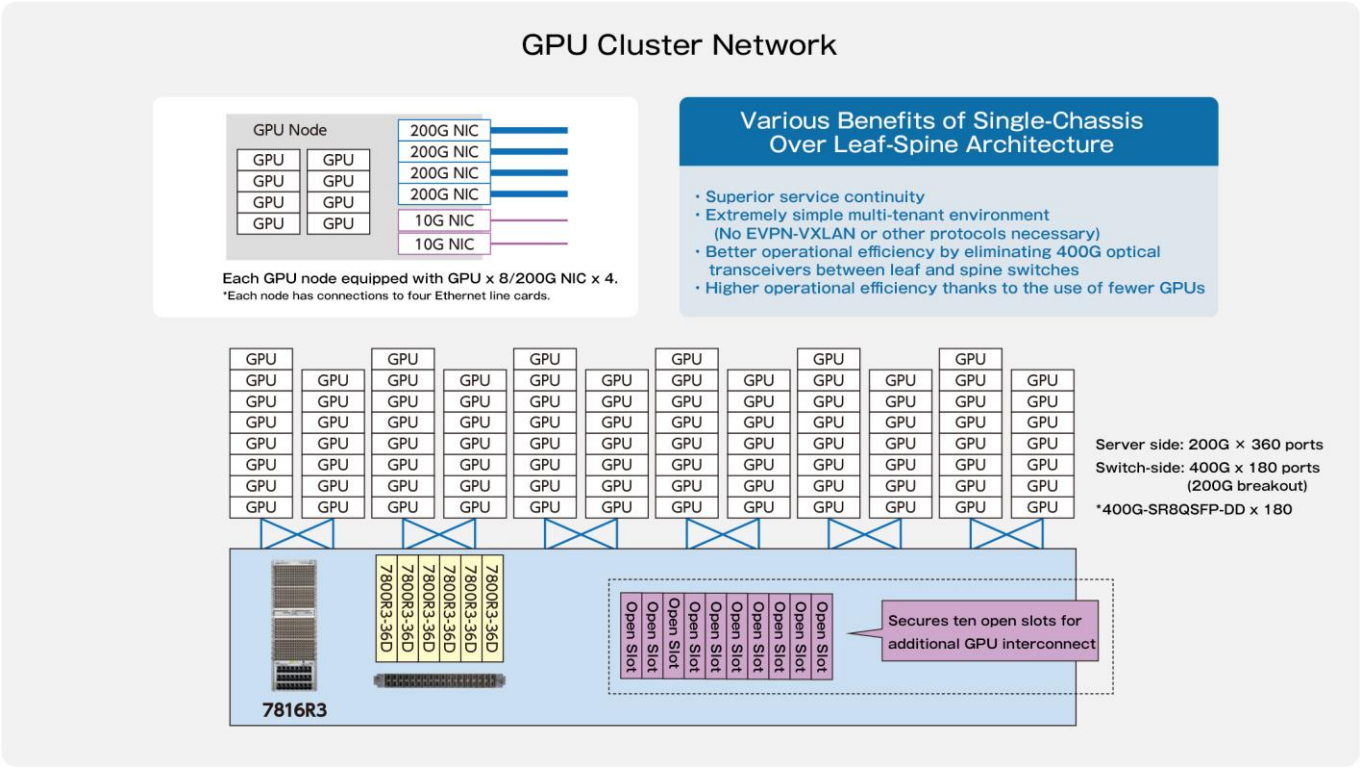

Koukaryoku PHY System Configuration

SAKURA internet rolled out the current system configuration around June 2023, and Koukaryoku PHY services went live in January 2024. The team used the time to ensure the launch of services would be successful, considering the scale of the system. Takashi Inoue, Leader of Infrastructure Development in the Cloud Service Department of the Cloud Business Division at SAKURA internet Inc. comments on the Koukaryoku PHY system configuration: “SAKURA internet has been providing these types of services that shift physical servers to cloud computing through its dedicated PHY servers since its founding. Fundamentally, this system configuration is an extension of those services that incorporates functionality unique to GPUs.” The two main goals were realizing a multi-tenant environment in a system connecting GPUs in a service providing multi-tenant infrastructure at its core, and keeping everything as simple as possible, which was necessary in able to configure everything in the short 6 month lead time.

The Koukaryoku PHY provides just under 70 GPU servers each equipped with eight NVIDIA H100 Tensor core GPUs. Each GPU server also has four 200Gb ENIC ports, which provide almost 300 ports to connect servers. Takashi Inoue wanted these connections to realize a simple way to distribute the traffic necessary in a multi-tenant environment. That is why the network configuration employs a single chassis system rather than the often-used leaf-spine architecture to directly connect switches to the Ethernet line card port and server-side NIC.

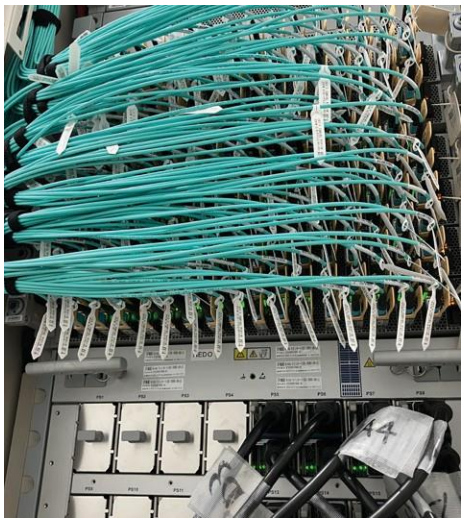

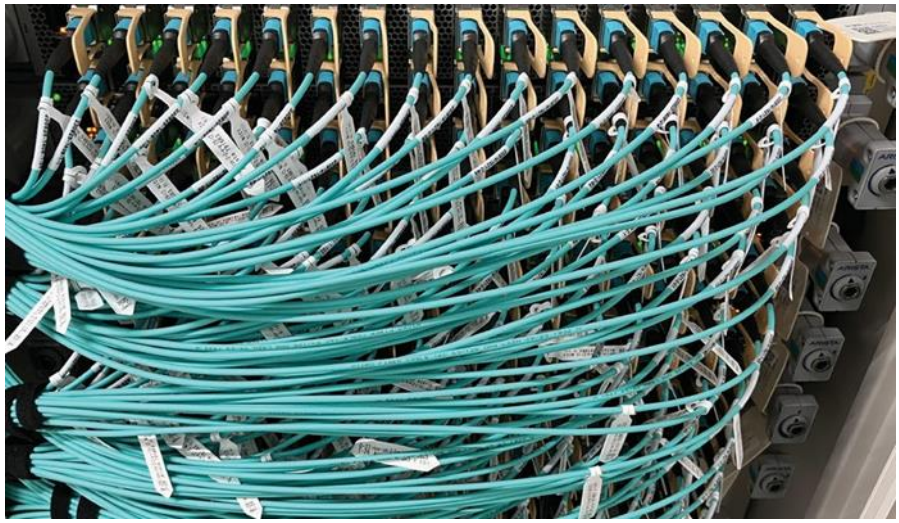

Arista Networks switch architecture does not require upper-layer protocols to distribute the network, and can independently control communication between ports because it can switch connections for all ports inside of the chassis, even in large-scale configurations that straddle Ethernet line cards. “It takes a lot of time to configure and validate EVPN, VXLAN and other distributed networks using protocols, but our system architecture did not require any of this extra work” explains Takashi Inoue. Another benefit he mentioned is Arista Networks switches provide superior schedule consistency when allocating resources between ports, which is a key point to offer uniform services in multi-tenant environments.

Takashi Inoue goes on to say, “Users who run large-scale computations take advantage of group communication using numerous servers, but the slowest server can hold up the rest during these calculations. From a consistency standpoint, that is why we want the overall latency to be mostly equal.

Leaf-spine infrastructure could handle far-away locations by incorporating other switches and devices, but a single chassis system has the benefit of sustaining the same latency regardless of the combination of servers that run these computations."

Ongoing Koukaryoku PHY Innovations

User demand has been high since the release of the Koukaryoku PHY service launched in January 2024 - only six months after its announcement in June 2023 - that SAKURA internet has already begun adding additional architecture. Current plans have enabled SAKURA internet to offer a total of 2,000 GPUs as of June 2024. These innovations also aim to increase the connection interface to 400GbE while still utilizing the same basic configuration. Moreover, SAKURA internet has already announced its plans to realize total computational capabilities of 18.9 EFLOPS by aggregating around 10,000 next-generation GPUs, including the NVIDIA HGX B200 system.

A wide range of industries are seeing a trend toward the use of regenerative and other AI technologies in the market which is increasing the importance of securing sufficient computational resources in order to not stifle innovation. As GPUs and the network infrastructure to efficiently connect a large number of GPUs become more important, products and technology from Arista Networks that can realize simple configurations for high-performance services will support the Koukaryoku PHY service.

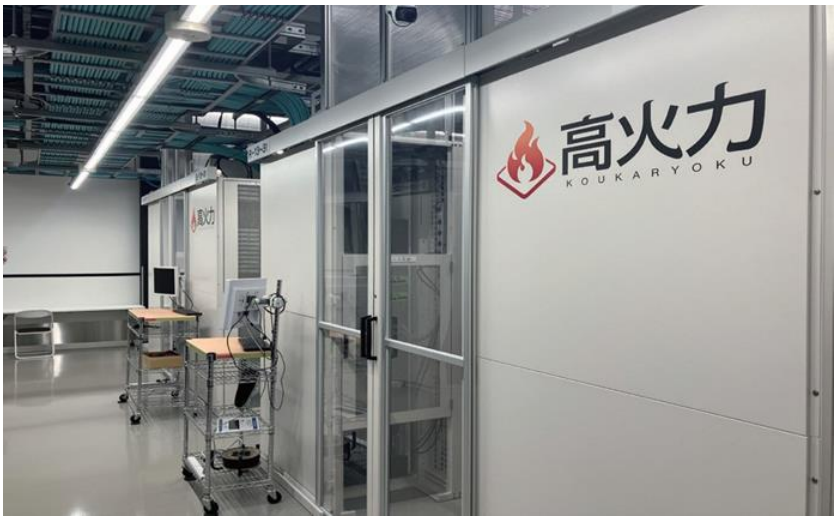

Takashi Inoue commented on the latest technological innovations in modern networks for AI saying, "In the past, the introduction of new technologies started in the telecommunications market before expanding out into other markets, but I think modern GPU and AI networks are now driving more and more technological innovations. I look forward to exchanging information and working closely with Arista Networks in the future in the pursuit of these revolutions in technology." Takefumi Sudo added, "It is extremely hard to guess what cloud services will be used on the other side of the internet. The benefits to using cloud services are obvious from a rational economic perspective, but it is impossible to gain the ability to build these systems regardless of how much more knowledge is gained about using these cloud services. Complete dependency on technologies from overseas comes with risks. That is why we agree with the Japanese government initiatives, that building these systems domestically are so important."

The Koukaryoku PHY that has been set up as vital infrastructure for government-level IT strategies also employs Arista Networks technologies. The partnership between Arista Networks and SAKURA internet for configuring these cutting-edge GPU/AI networks hopes to bring about new innovations through a collaborative relationship.

Left to Right:
SAKURA internet Inc.
Cloud Service Department
Cloud Business Division
Assistant Director
Takefumi Sudo

Infrastructure Development Group
Cloud Service Department
Cloud Business Division
Leader
Takashi Inoue

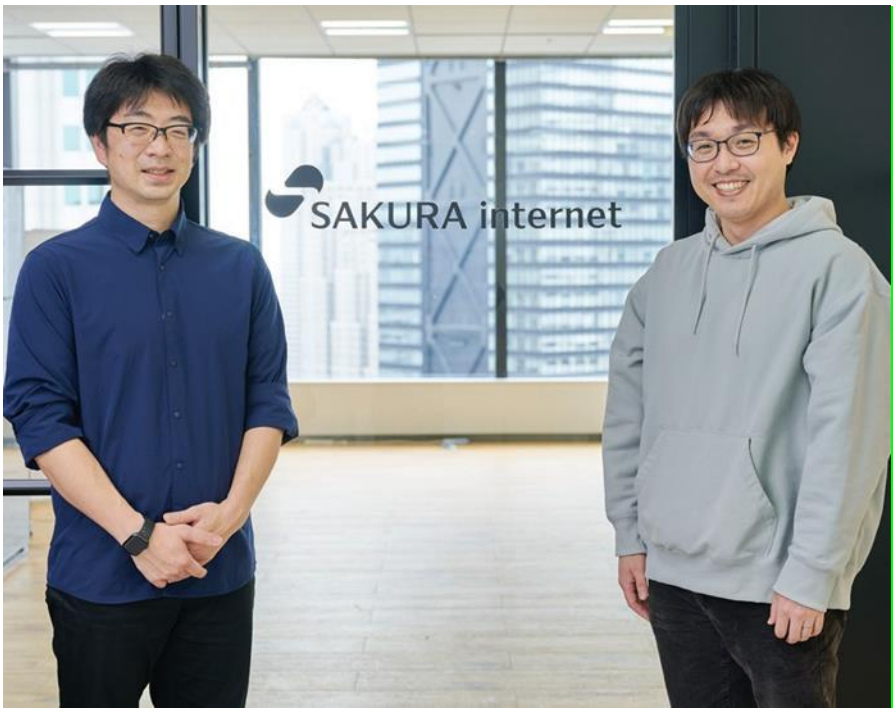

contact-japan@arista.com

www.arista.com/jp

Tokyo
Level27, Tokyo Sankei Building,
1-7-2 Otemachi, Chiyoda-ku,
Tokyo 100-0004, Japan

Osaka
Level19 Hilton Plaza West Office Tower,
2-2-2 Umeda Kita-ku
Osaka 530-0001, Japan