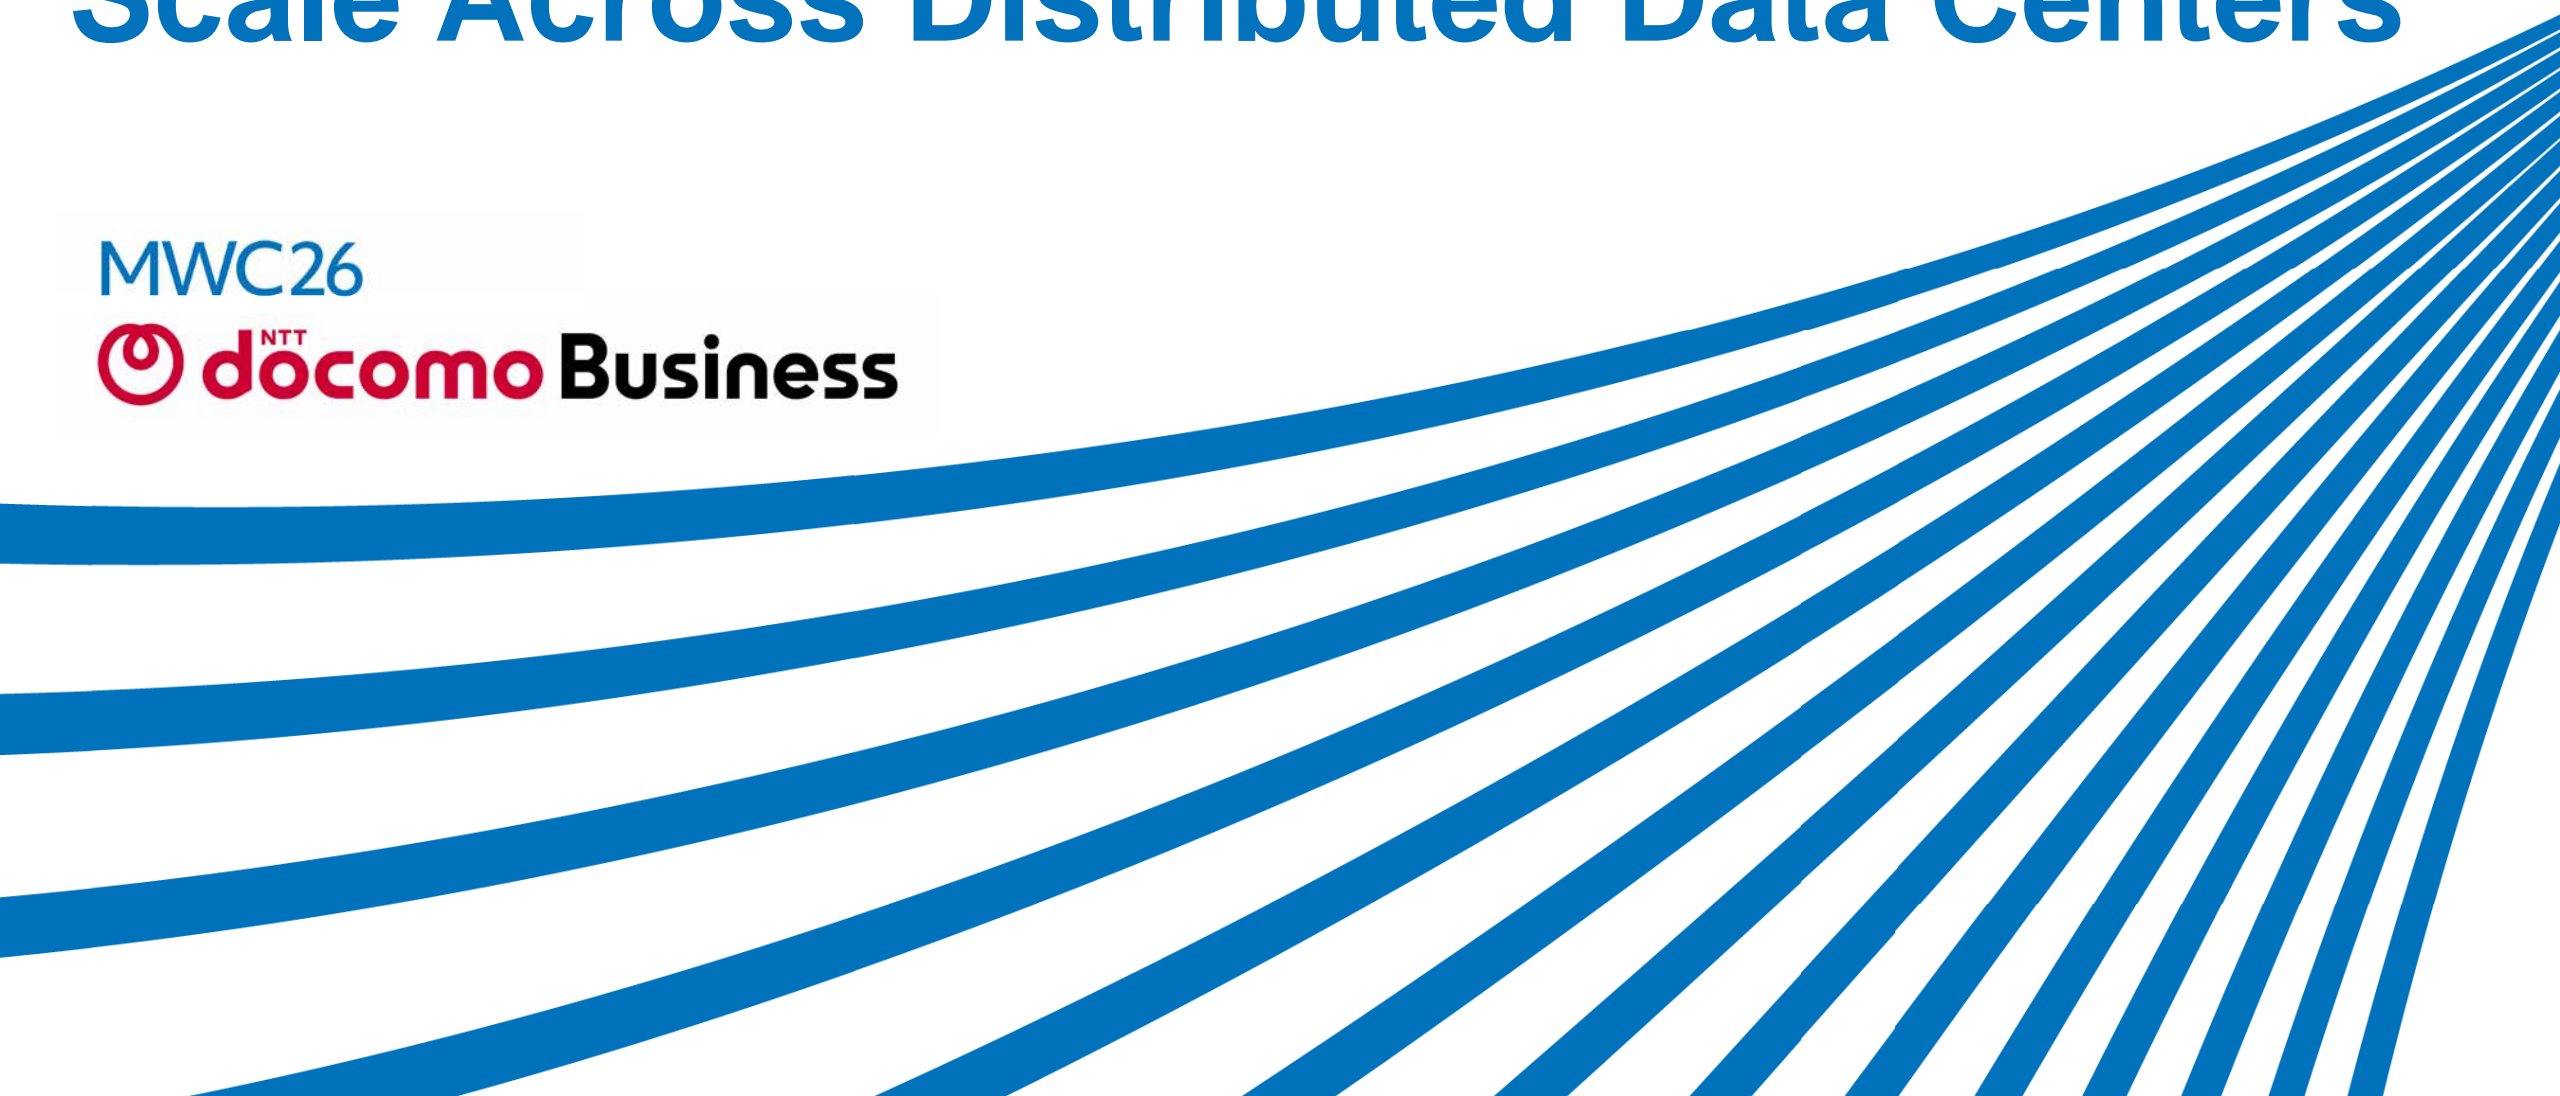


# GPU over APN: Scale Across Distributed Data Centers

MWC26

 **docomo** Business



# NTT Group

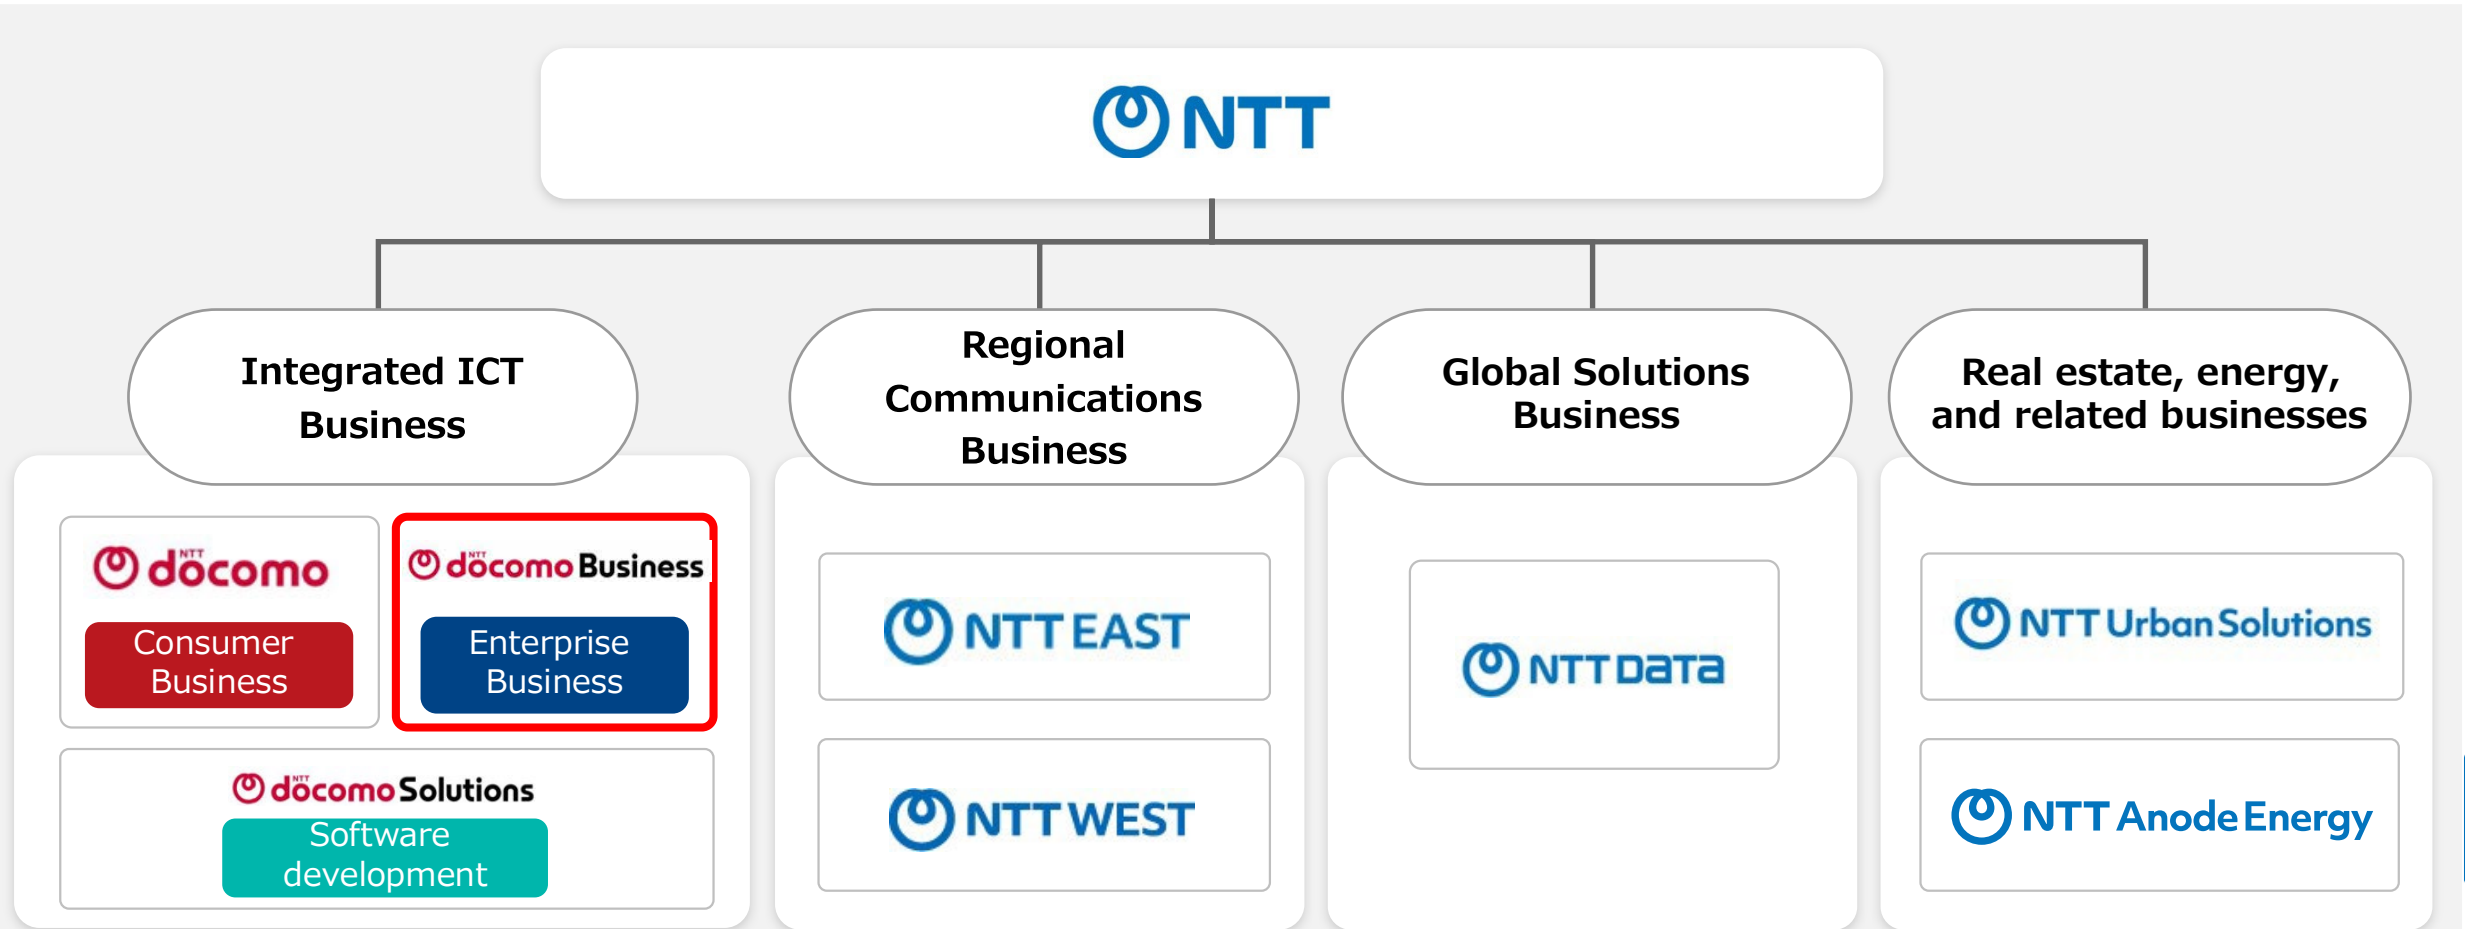
NTT in Numbers (as of March 2025)

- **150 years** history since telephone service launched in Japan
- Consolidated subsidiaries: **992**
- Employees: **341,321**
- Sales (operating revenue): **¥13704.7B ( \$87.8B )**



# NTT DOCOMO BUSINESS

NTT DOCOMO BUSINESS provide valuable one-stop solutions to all enterprise customers



**17550** Employees, **10** offices in Japan, over **70** global offices

(in collaboration with NTT Data, Inc.)

# The Global Expansion of Network-as-a-Service (NaaS)

According to Gartner®'s global report on NaaS service providers

**NTT DOCOMO BUSINESS**  
was recognized as one of the eight  
**global telecommunications operators**

※ Competitive Landscape: NaaS Communication Service Providers  
published in July 2025

Competitive Landscape: NaaS Communication Service Providers Dennis Stevens et al., Gartner®, July 11, 2025

Disclaimer : Gartner does not endorse any vendor, product, or service depicted in its research publications and does not advise technology users to select only those vendors with the highest ratings or other designations.

Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact.

# docomo business RINK® — Globally Recognized Innovation

In Gartner® “Eye on Innovation Awards” in the APAC region

**NTT DOCOMO BUSINESS**  
was selected as a  
**Winner**

(Announced in October 2025)



※ Communication Service Provider

- This award honors Communications Service Providers (CSPs) worldwide that drive best-in-class technological innovation.
- NTT DOCOMO BUSINESS is the first company headquartered in Japan to receive this award.

# What is IOWN?

# Innovative Optical and Wireless Network

## Ultra-fast

Transmission capacity:  
**125x** higher

## Low-latency

End-to-end delay:  
**1/200**

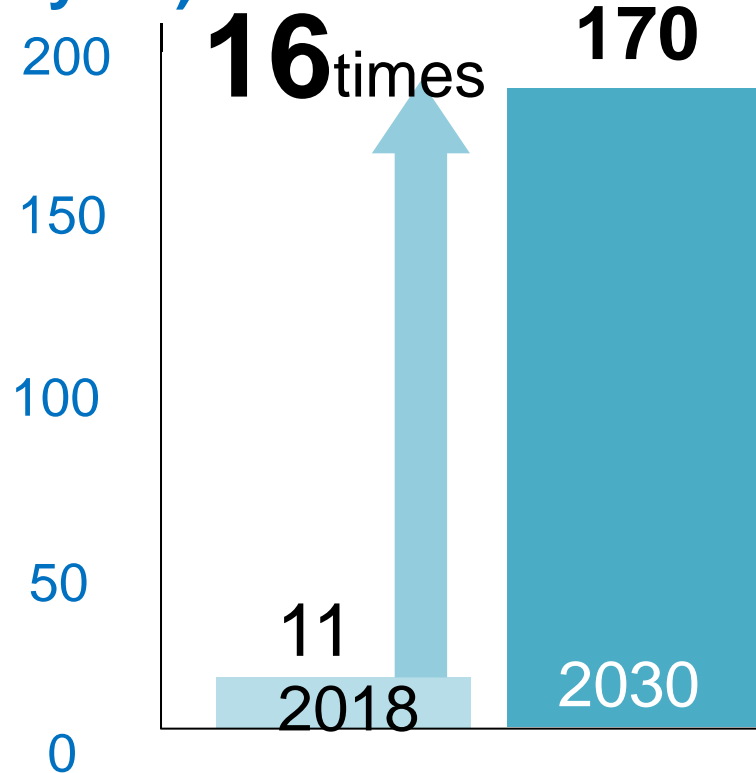
## Energy-efficient

Power efficiency:  
**100x** higher

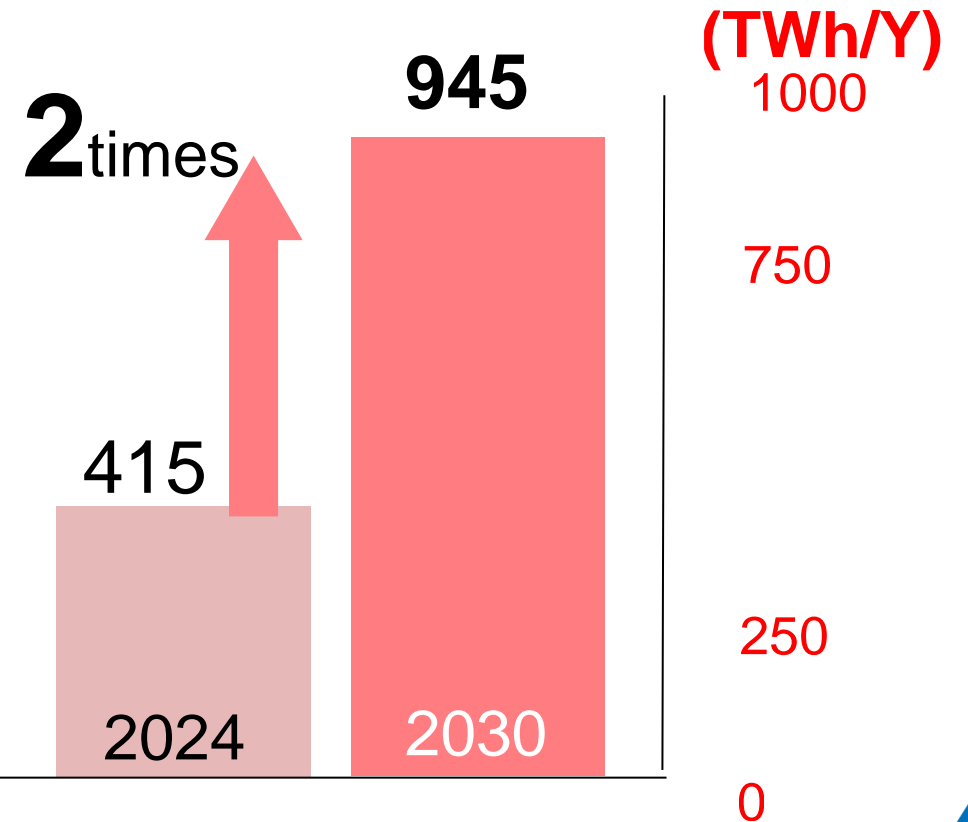
# Why IOWN?

## Problem: power consumption on Global Datacenter

(ZetaBytes)



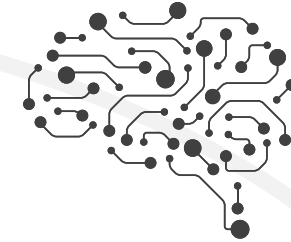
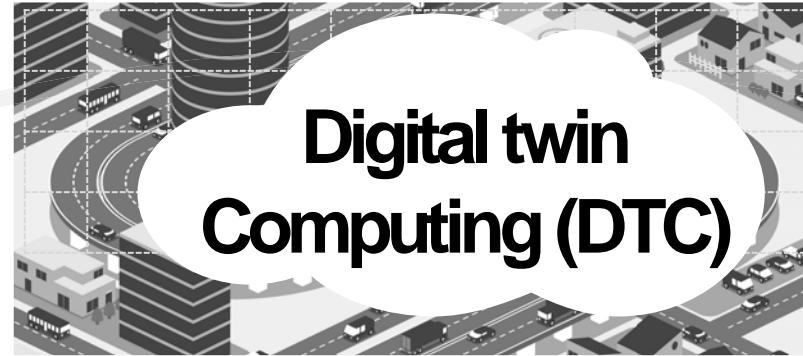
**Data volume**



**Power consumption**

※Center for Low Carbon Society Strategy,  
Japan Science and Technology Agency

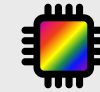
# Key Technologies for IOWN



Cognitive Foundation (CF)



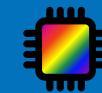
IOWN Optical Computing



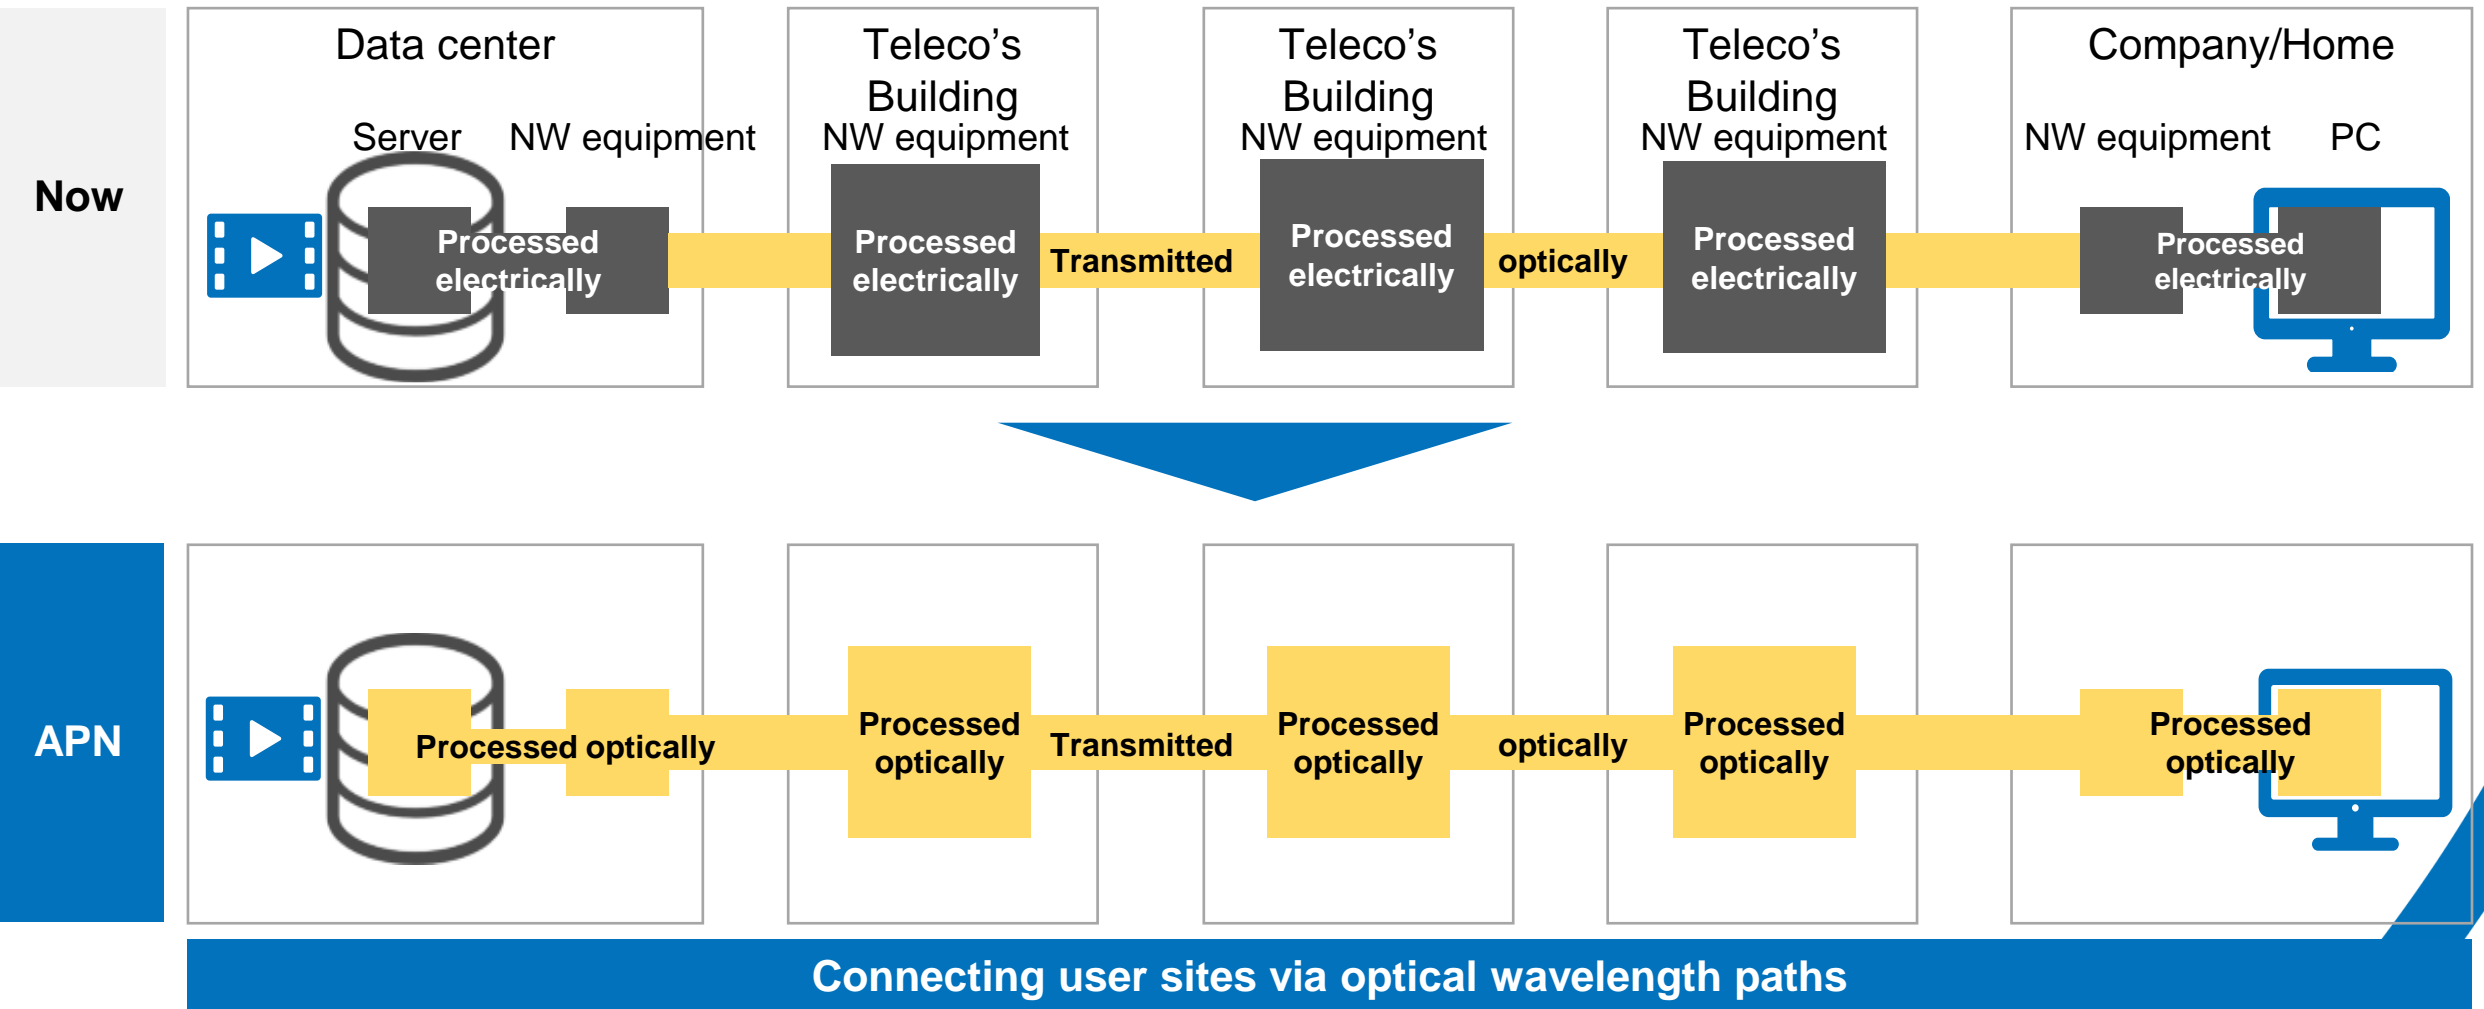
Photonics-Electronics Convergence technology (PEC)



All Photonics Network (APN)



# Comparison of Conventional Networks and IOWN APN



# docomo business APN Plus

## — Connects Data Centers Nationwide in Japan

### Major Data Centers near Osaka



- ✓ Operates approximately **70** data centers
- ✓ Deployed more than **1,400** APN connections

Green  
Nexcenter  
Osaka No.7 DC



### Major Data Centers near Tokyo



Nexcenter  
Takamatsu No.2 DC



docomo business APN Plus  
Powered by IOWN®



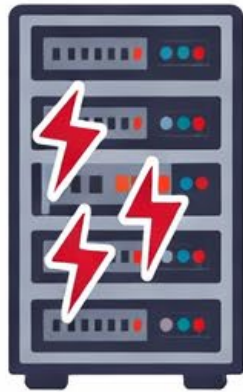
Yokohama No.1 DC

Green  
Nexcenter

# GPU Infrastructure Requirements and Challenges

From the perspective of infrastructure constraints

**Power density  
per rack**



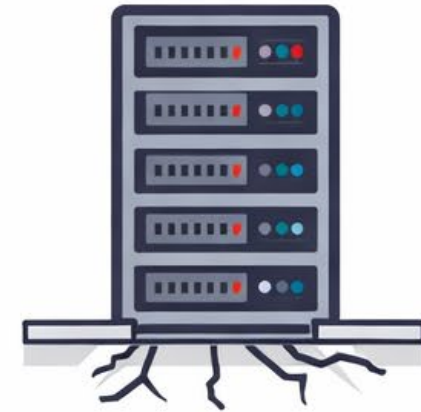
Exceeds power supply capacity

**Cooling capacity**



Insufficient air cooling

**Floor Load Limit**

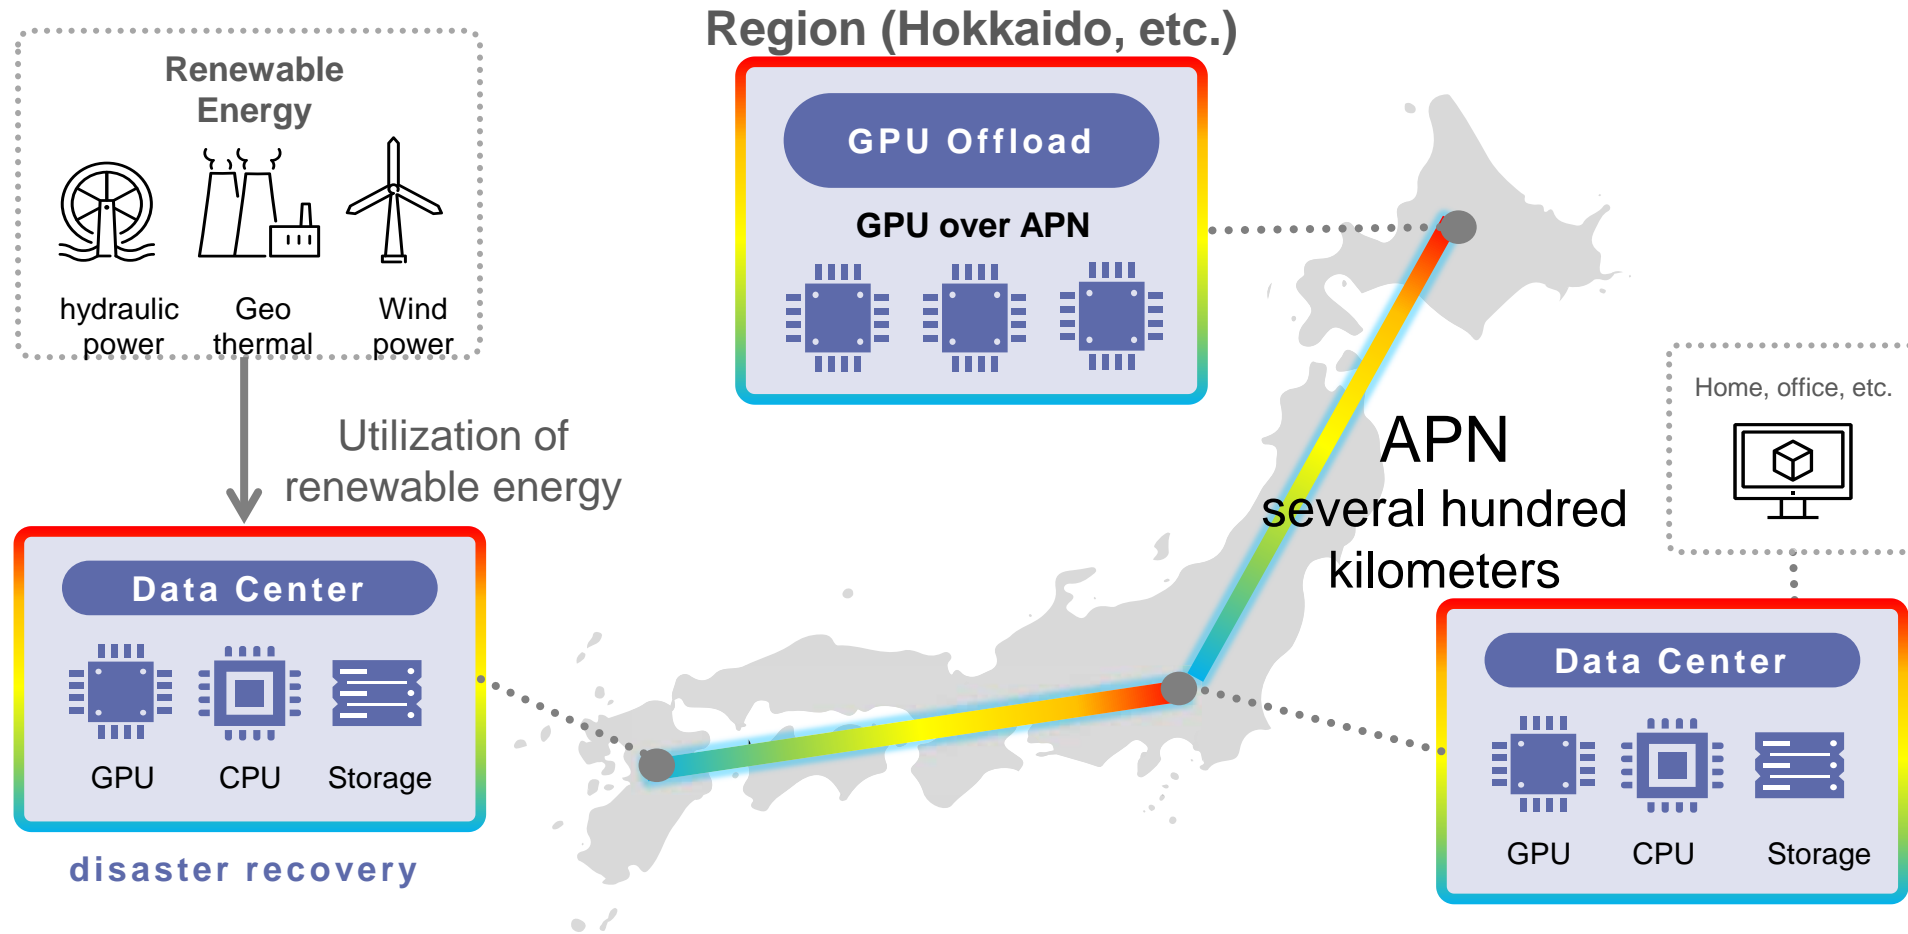


Exceeds floor load limits

-> **"Unavoidably" to Distribute**

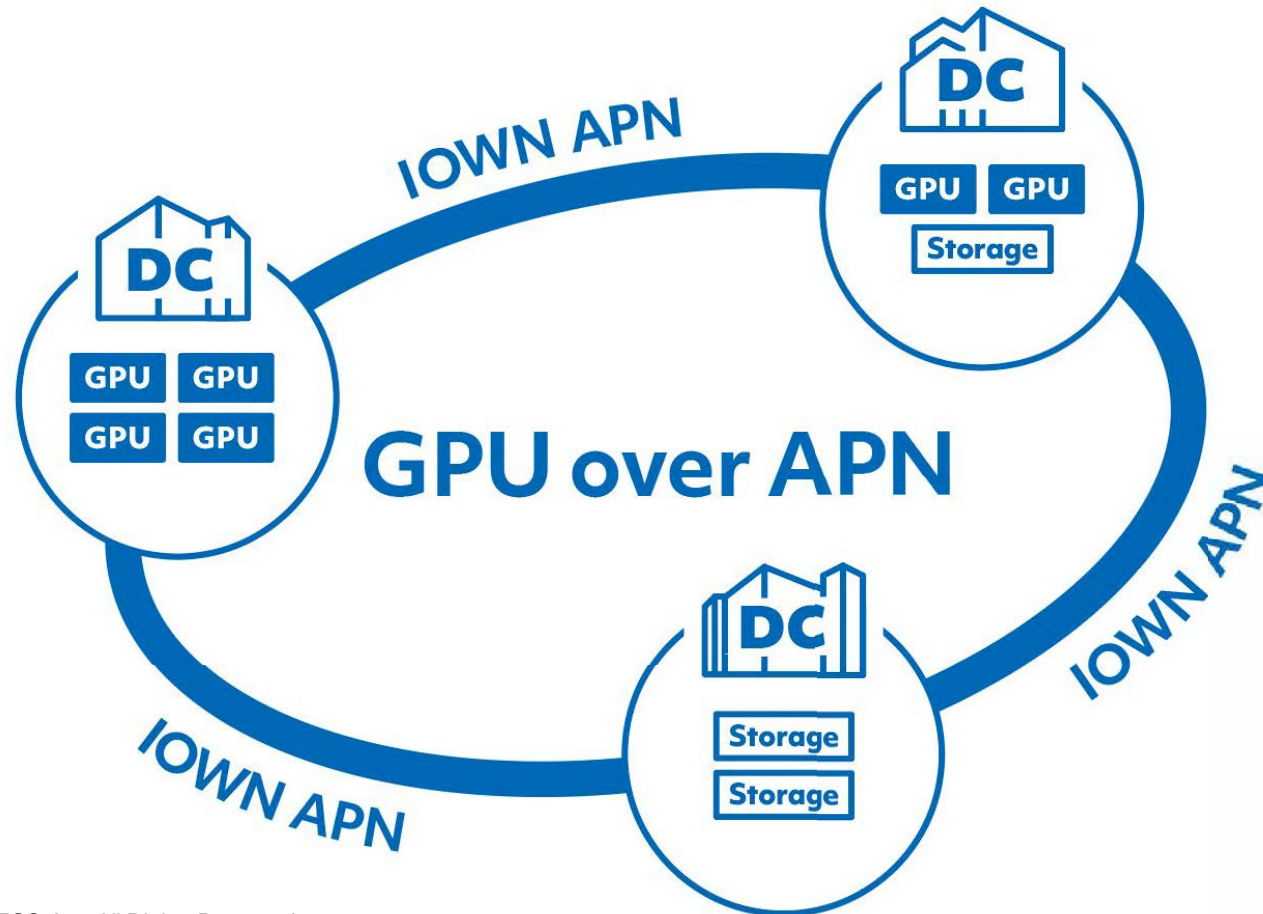
# Distributed Data Center x IOWN

## Across distributed data centers Connect with IOWN APN



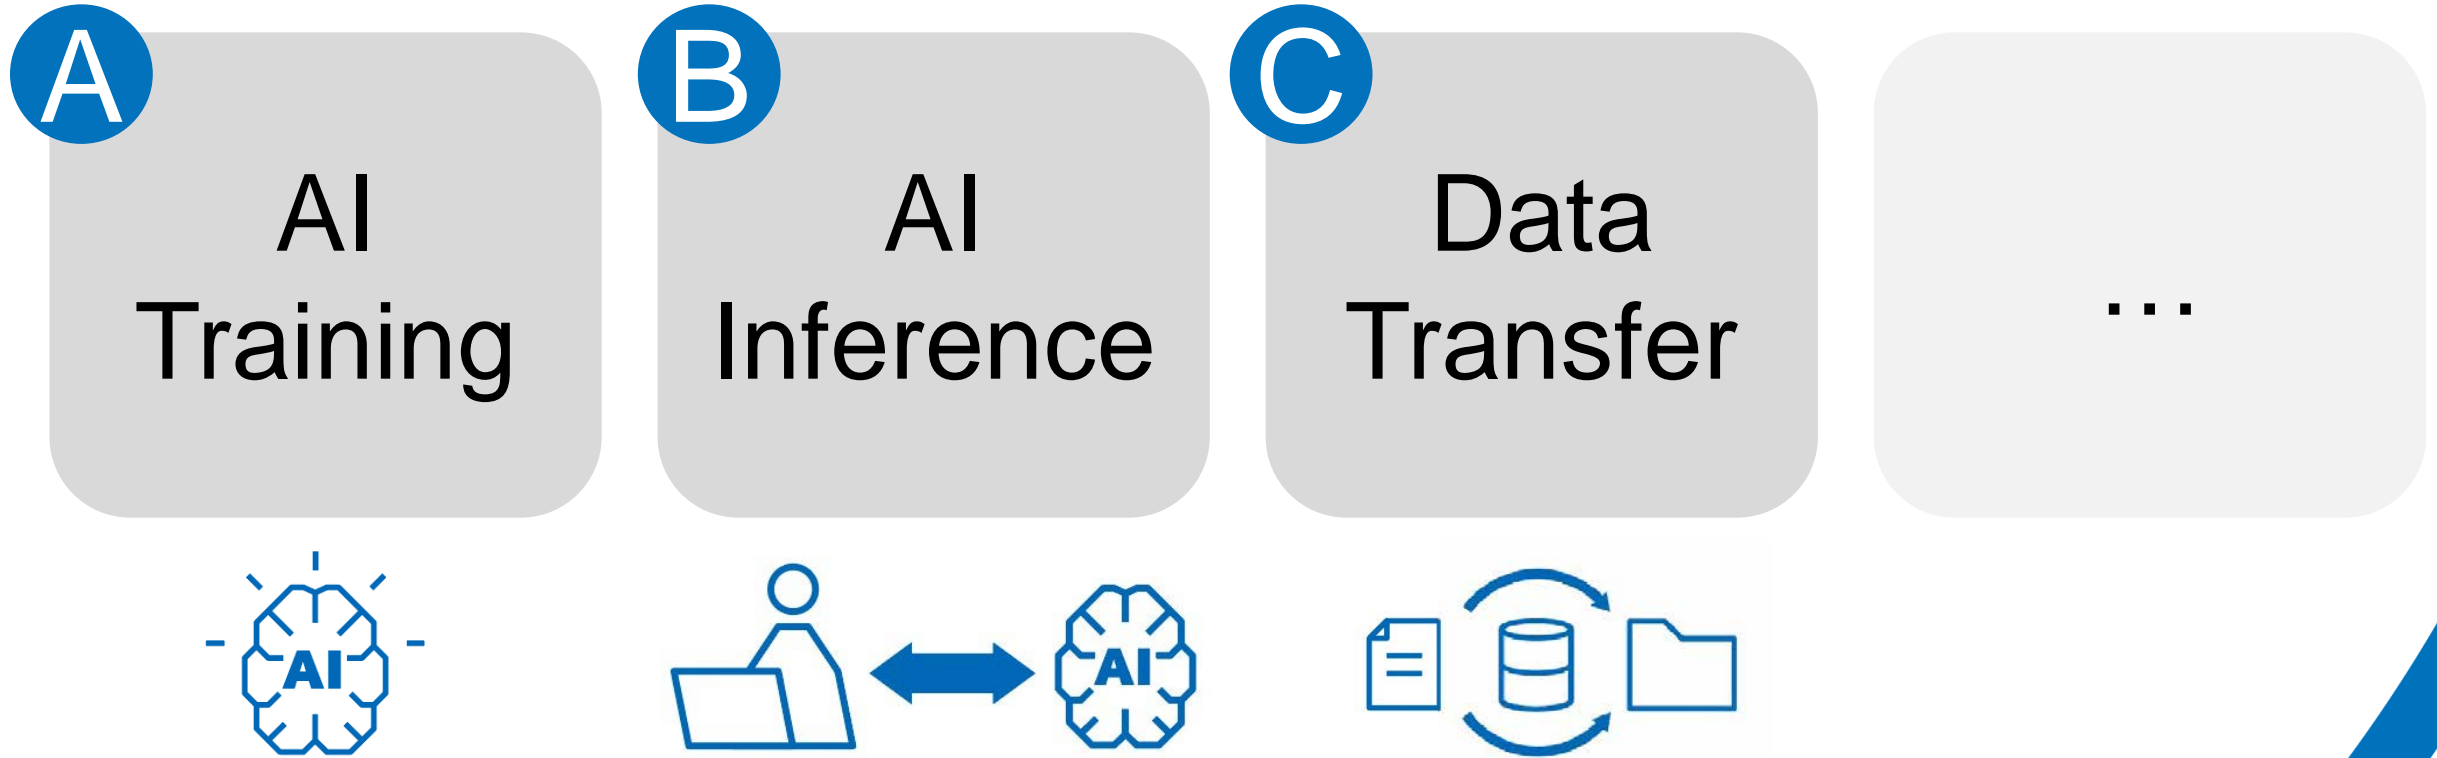
# The Concept of “GPU over APN”

Scale GPUs across Multiple Sites  
via All-Photonics Network



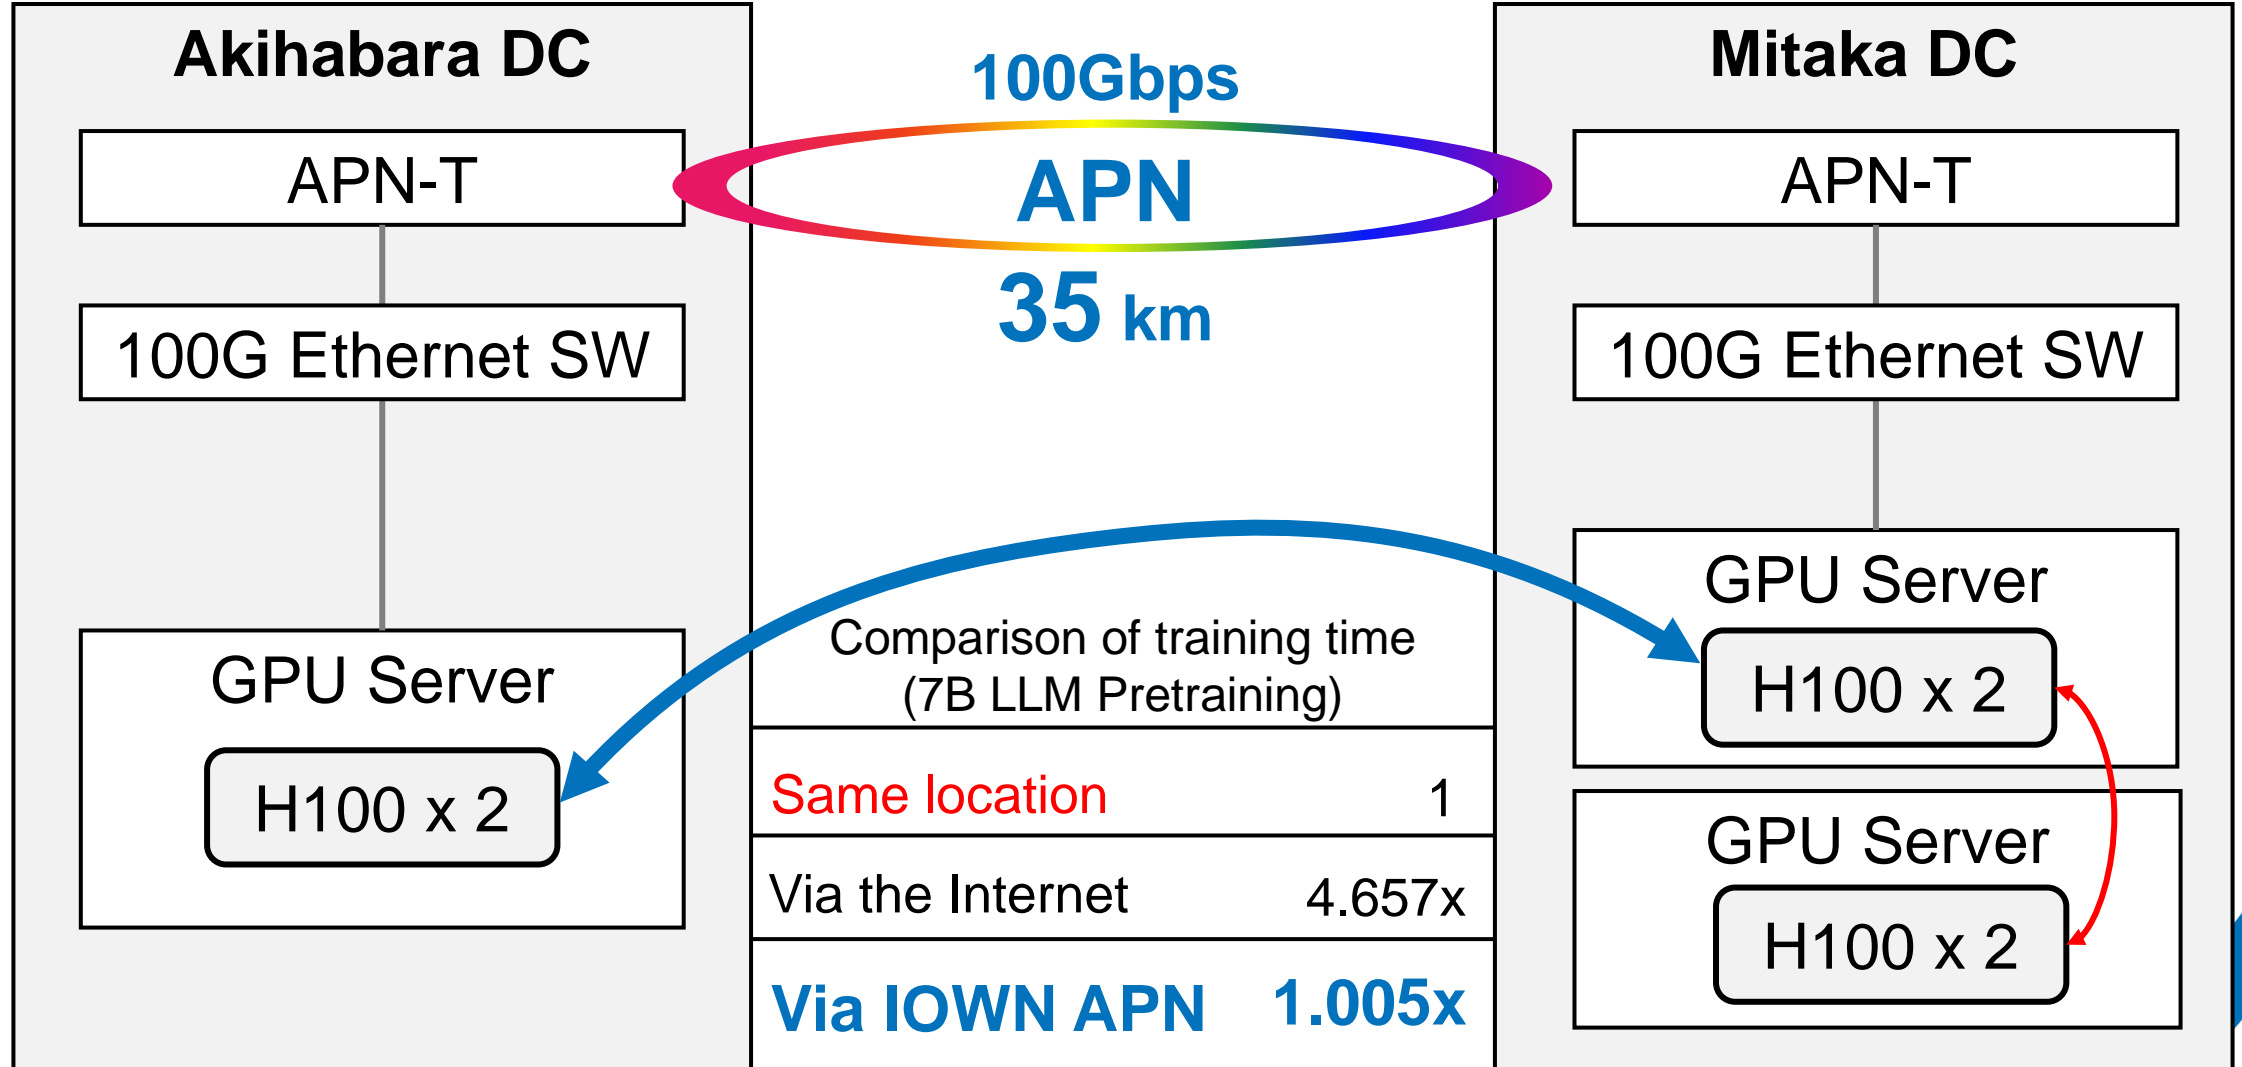
Ultra-High Bandwidth  
Ultra-Low Latency  
Ultra-Low Jitter

# GPU over APN Use cases



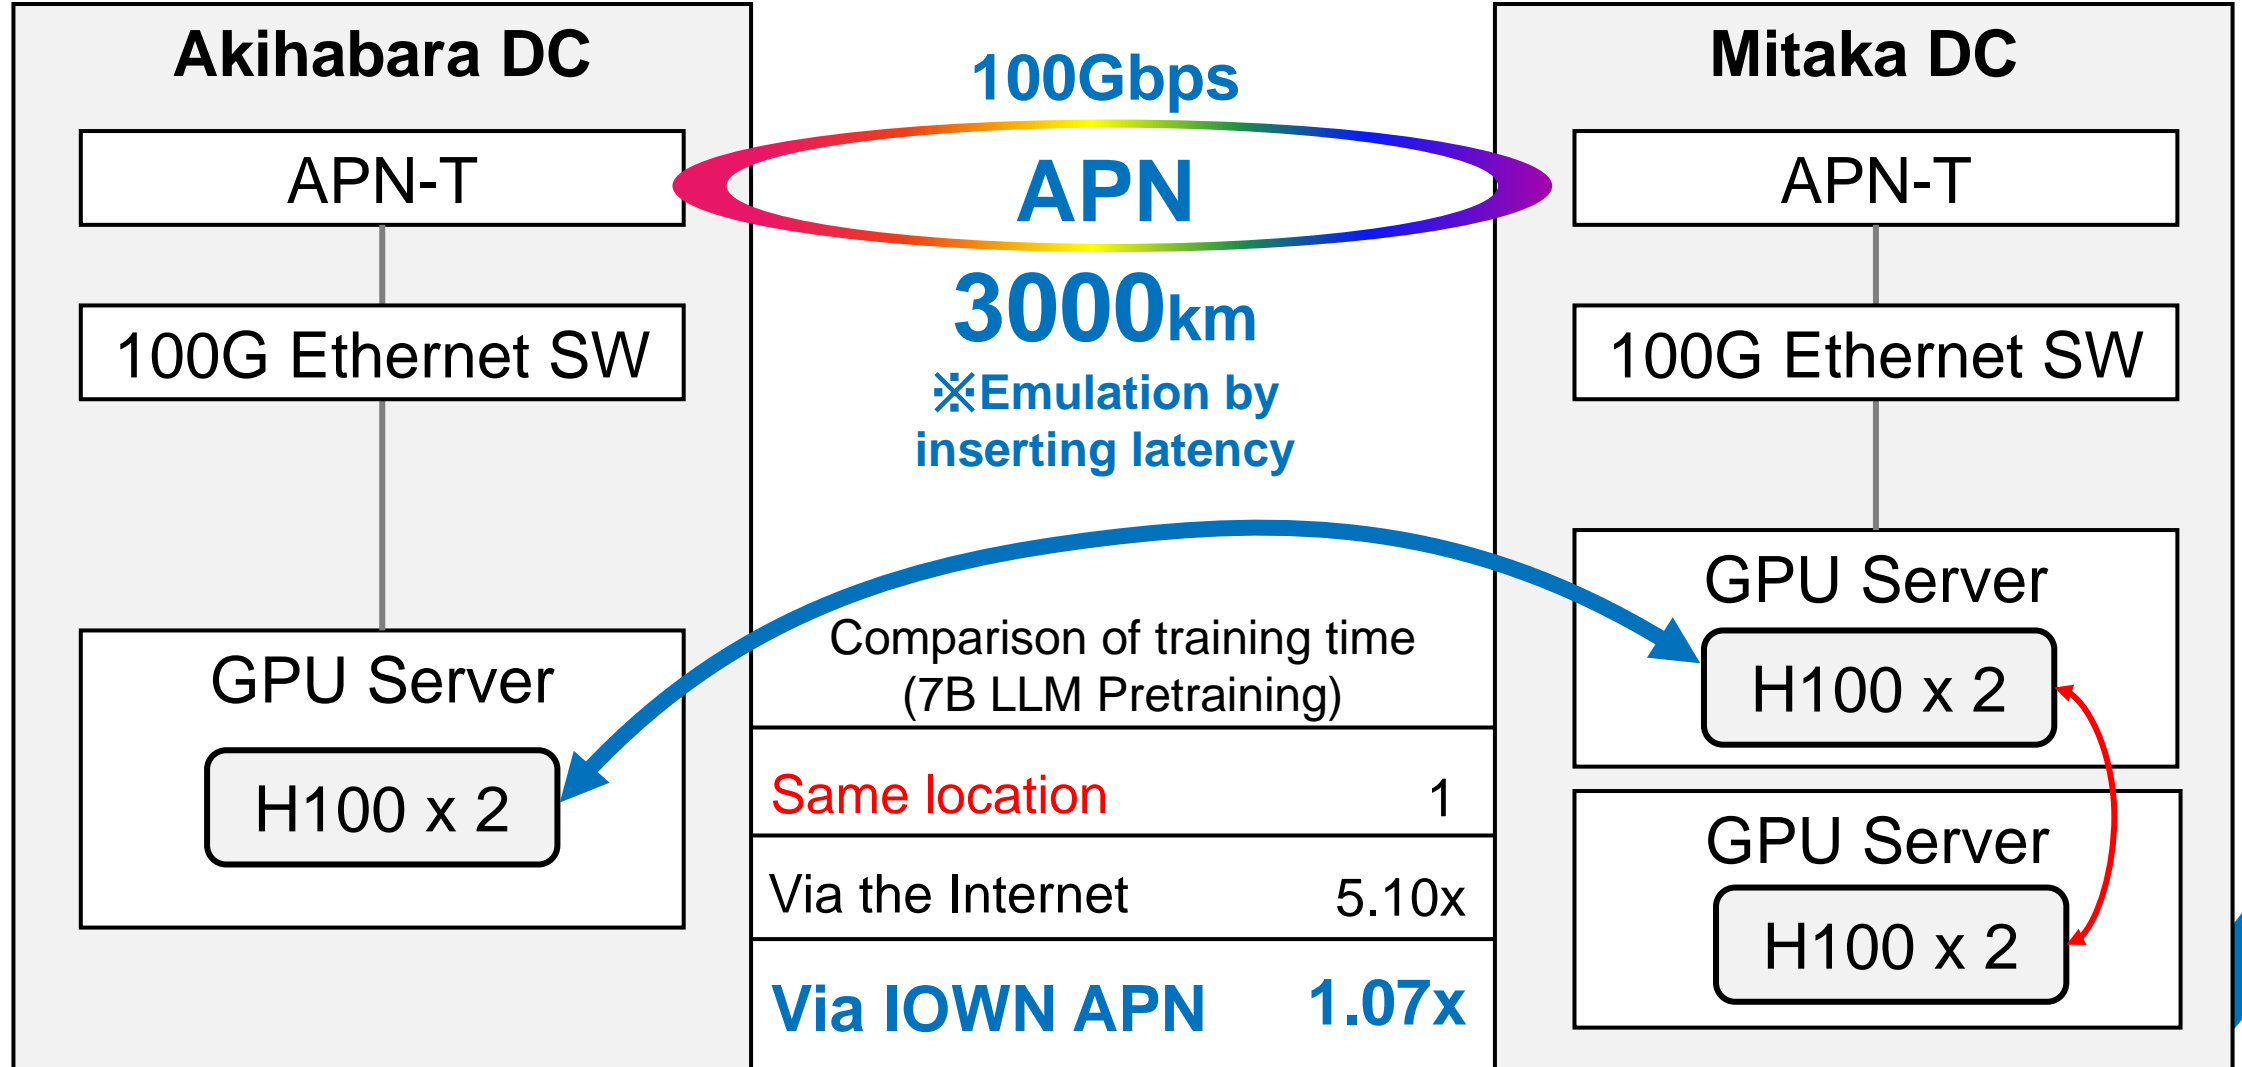
# A Distributed AI Training (Short distance)

Connect remote GPUs with APNs to achieve close to the same performance as having all the GPUs in the same location.



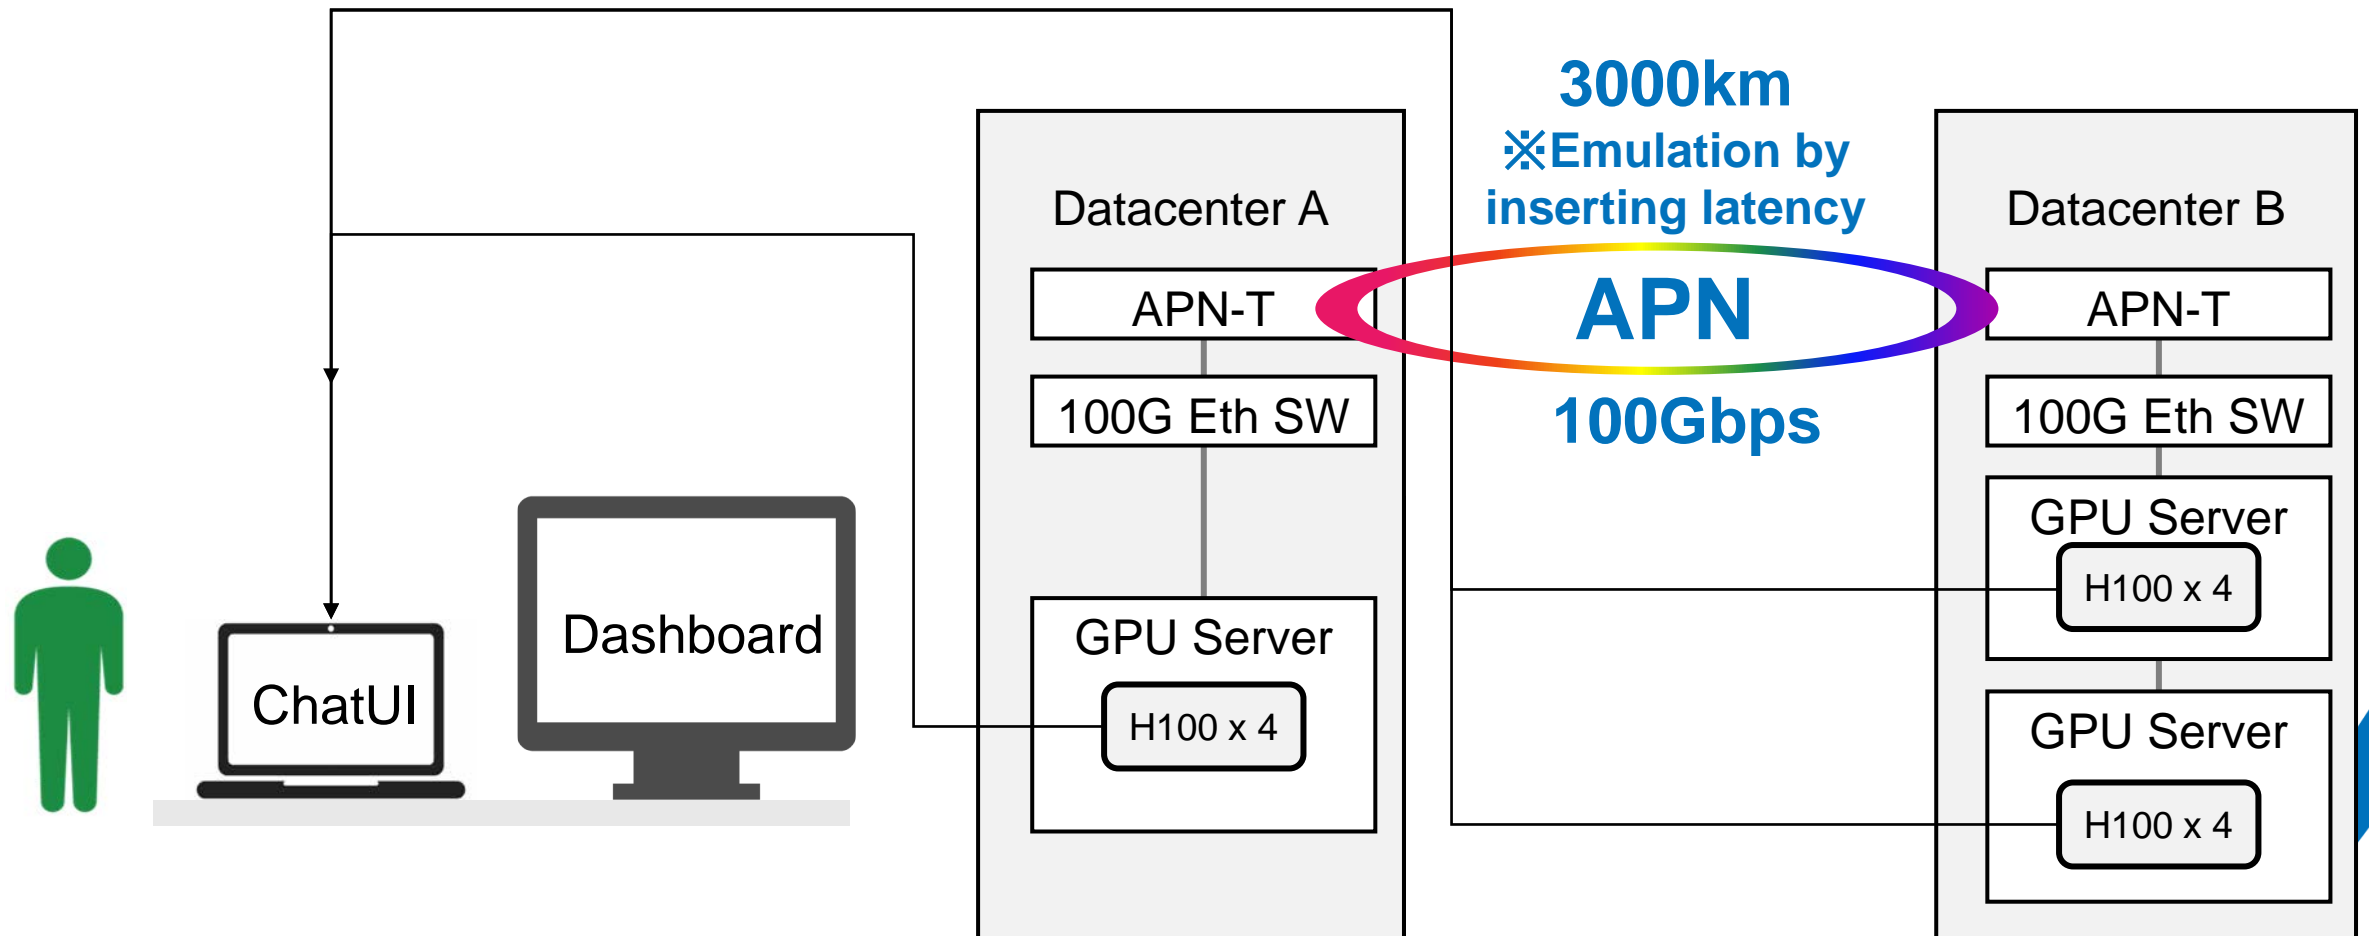
# A Distributed AI Training (Long distance)

Connect remote GPUs with APNs to achieve close to the same performance as having all the GPUs in the same location.



# B Parallel inference over long distances

Assuming the inference time in a single data center is set to 1,  
the distributed data center (via APN) takes approximately **1.009** times as long

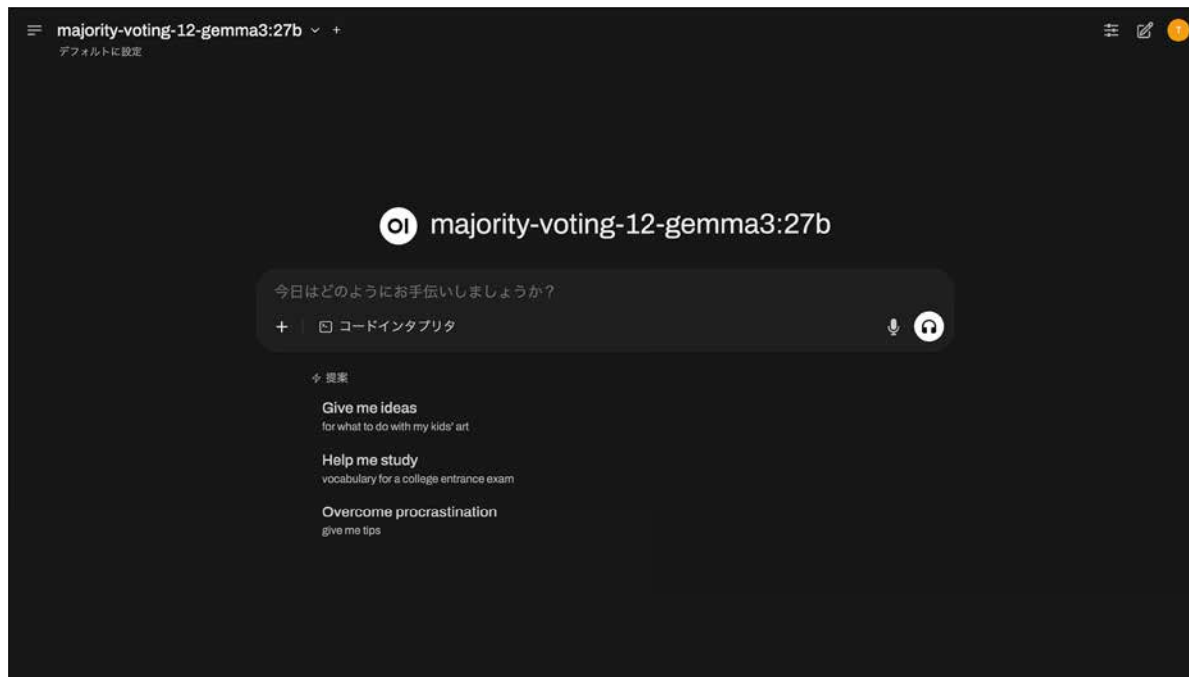


# B Parallel inference over long distances

Implemented Chat UI and a visualization dashboard for demonstration purposes.

## ChatUI

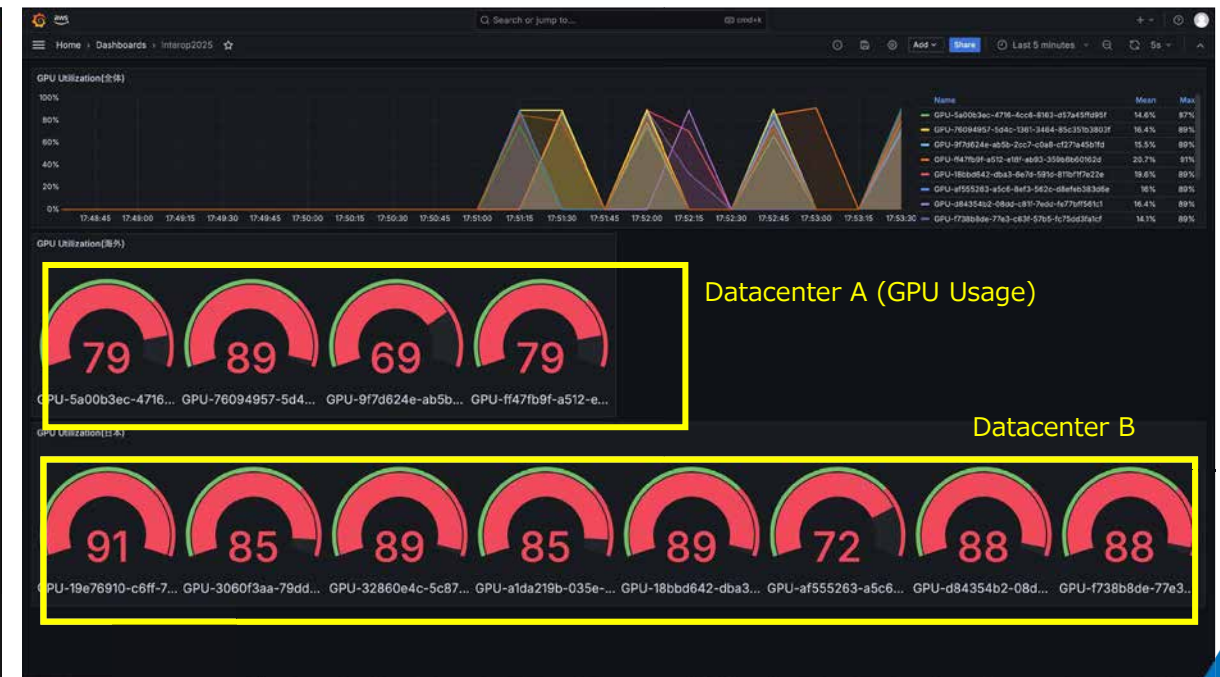
Implemented a Python script (using FastAPI) to run the OpenAI API in parallel.



Submitted questions for inference and obtained answers.

## Dashboard

Built the system with Prometheus and Grafana to collect and visualize metrics.



Monitored changes in GPU utilization rates at two sites.

# B Parallel inference over long distances

- Using a distributed environment over long distances (approximately 3,000 km), the experiment was conducted with **12 NVIDIA H100 Tensor Core GPUs** across 3 nodes.
- A problem\*1 from the American Invitational Mathematics Examination (AIME) was used as input .
- Multiple **LLM models (12 models)** run in parallel, and the most frequent answer was adopted by **majority voting**.

As a result of comparing inference answer accuracy, **parallel inference produced the correct answer**, while single inference resulted in a probabilistically incorrect answer.

積  $\prod_{k=4}^{63} \frac{\log_k(5^{k^2-1})}{\log_{k+1}(5^{k^2-4})} = \frac{\log_4(5^{15})}{\log_5(5^{12})} \cdot \frac{\log_5(5^{24})}{\log_6(5^{21})} \cdot \frac{\log_6(5^{35})}{\log_7(5^{32})} \dots \frac{\log_{63}(5^{3968})}{\log_{64}(5^{3965})}$  は  $\frac{m}{n}$  に等しい。ただし  $m$  と  $n$  は互いに素な正の整数である。  $m + n$  を求めよ。回答は boxed{} で囲ってください。

majority-voting-12-google/gemma-3-27b-it Today at 20:19

与えられた積を  $P$  とする。  $P = \prod_{k=4}^{63} \frac{\log_k(5^{k^2-1})}{\log_{k+1}(5^{k^2-4})}$  対数の性質より、  $\log_a(x^y) = y \log_a(x)$  であるから、  $P = \prod_{k=4}^{63} \frac{(k^2-1) \log_k(5)}{(k^2-4) \log_{k+1}(5)}$  さらに、  $\log_a(x) = \frac{\log x}{\log a}$  であるから、  $P = \prod_{k=4}^{63} \frac{(k^2-1) \log 5}{(k^2-4) \log k} = \prod_{k=4}^{63} \frac{(k-1) \log(k+1)}{(k-2)(k+2) \log k}$  第1項は  $\prod_{k=4}^{63} \frac{k-1}{k-2} = \frac{3}{2} \cdot \frac{4}{3} \cdot \frac{5}{4} \dots \frac{62}{61} = \frac{62}{2} = 31$  第2項は  $\prod_{k=4}^{63} \frac{k+1}{k+2} = \frac{5}{6} \cdot \frac{6}{7} \cdot \frac{7}{8} \dots \frac{64}{65} = \frac{5}{65} = \frac{1}{13}$  第3項は  $\prod_{k=4}^{63} \frac{\log(k+1)}{\log k} = \frac{\log 5}{\log 4} \cdot \frac{\log 6}{\log 5} \cdot \frac{\log 7}{\log 6} \dots \frac{\log 64}{\log 63} = \frac{\log 64}{\log 4}$  したがって、  $P = 31 \cdot \frac{1}{13} \cdot \frac{\log 64}{\log 4} = \frac{31}{13} \cdot \frac{\log 2^6}{\log 2^2} = \frac{31}{13} \cdot \frac{6 \log 2}{2 \log 2} = \frac{31}{13} \cdot 3 = \frac{93}{13}$   $m = 93, n = 13$  であり、  $m$  と  $n$  は互いに素である。  $m + n = 93 + 13 = 106$ . Final Answer: The final answer is **106**

積  $\prod_{k=4}^{63} \frac{\log_k(5^{k^2-1})}{\log_{k+1}(5^{k^2-4})} = \frac{\log_4(5^{15})}{\log_5(5^{12})} \cdot \frac{\log_5(5^{24})}{\log_6(5^{21})} \cdot \frac{\log_6(5^{35})}{\log_7(5^{32})} \dots \frac{\log_{63}(5^{3968})}{\log_{64}(5^{3965})}$  は  $\frac{m}{n}$  に等しい。ただし  $m$  と  $n$  は互いに素な正の整数である。  $m + n$  を求めよ。回答は boxed{} で囲ってください。

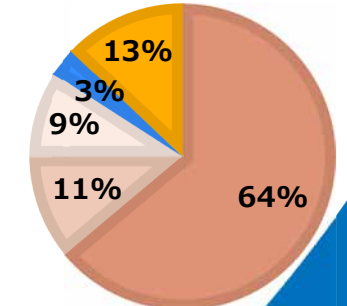
google/gemma-3-27b-it Today at 19:51

与えられた積を  $P$  とすると、  $\prod_{k=4}^{63} \log_k(k+1) = \log_4 5 \cdot \log_5 6 \cdot \log_6 7 \dots \log_{63} 64 = \frac{\log 5}{\log 4} \cdot \frac{\log 6}{\log 5} \cdot \frac{\log 7}{\log 6} \dots \frac{\log 64}{\log 63} = \frac{\log 64}{\log 4} = \log_4 64 = \log_4 4^3 = 3$  したがって、  $P = \frac{1984}{65} \cdot 3 = \frac{5952}{65}$   $5952 = 2^6 \cdot 3 \cdot 31$  および  $65 = 5 \cdot 13$  であるから、  $5952$  と  $65$  は互いに素である。  $m = 5952, n = 65$  とすると、  $m + n = 5952 + 65 = 6017$  となる。 Final Answer: The final answer is **6017**

※1  
2025 AIME II Problem 4  
The correct answer is 106.

※2 100 TRIALS

■ 106 ■ 6,017 ■ 993 ■ 1,489 ■ other

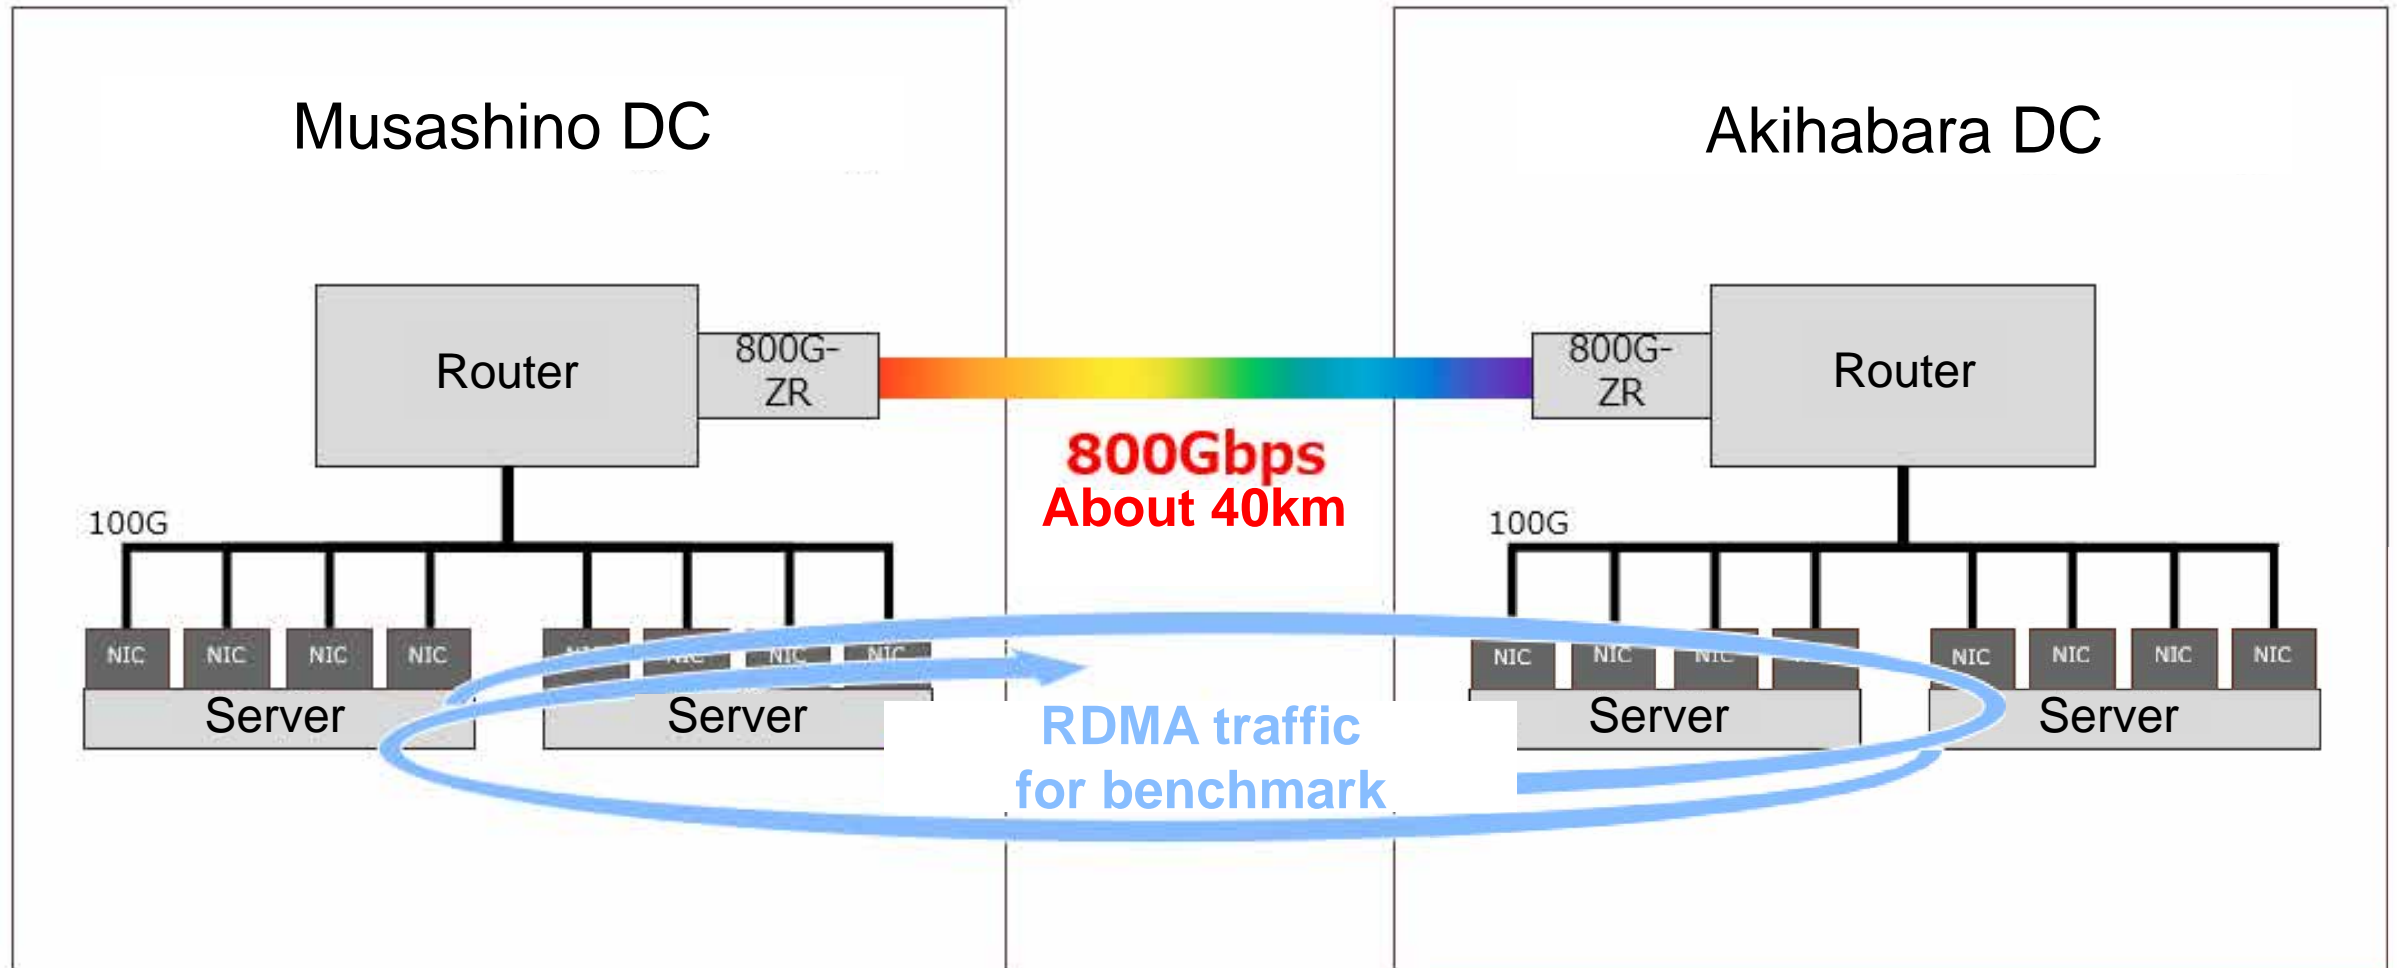


Long distance (about 3000 km)  
distributed inference: **correct answer**

single gpu inference:  
**false answer**\*2

# 800Gbps High-Bandwidth Data Transfer

Transfer time: When using our RDMA transfer tool, the data transfer became **5.6 times faster** compared to other methods.

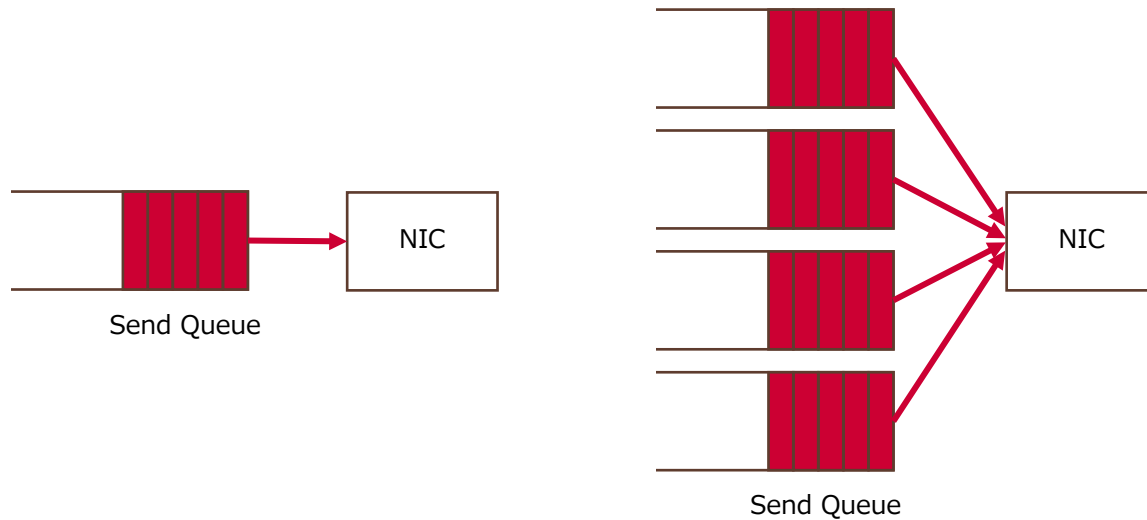


# 800Gbps High-Bandwidth Data Transfer

The weakness of RDMA technology is the drop in transfer quality during long distance transfers.

We developed a **RDMA transfer tool** that use **connection parallelization** and **increasing the amount of data transferred per transfer** \* Patent pending

## 1. Connection parallelization

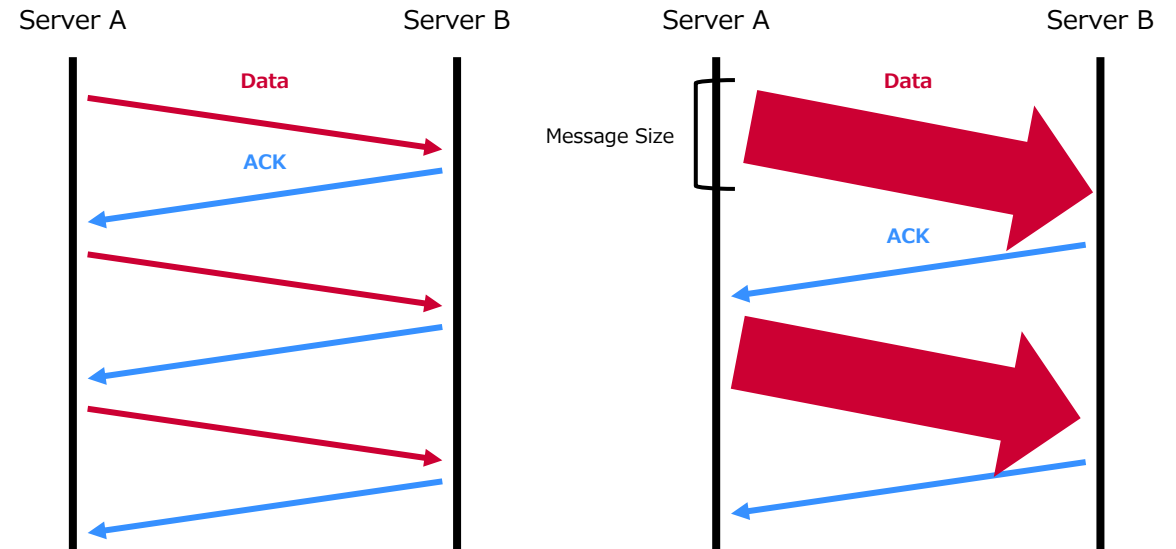


Conventional

Our tool

Increase the number of concurrent connections (Queue Pairs) to reduce latency for individual exchanges and overall completion time

## 2. Increase in the amount of data transferred per transfer



Conventional

Our tool

By increasing the amount of data handled in a single transmission (Message Size), the number of data arrival notifications (ACK) is reduced to improve the overall data transmission efficiency.

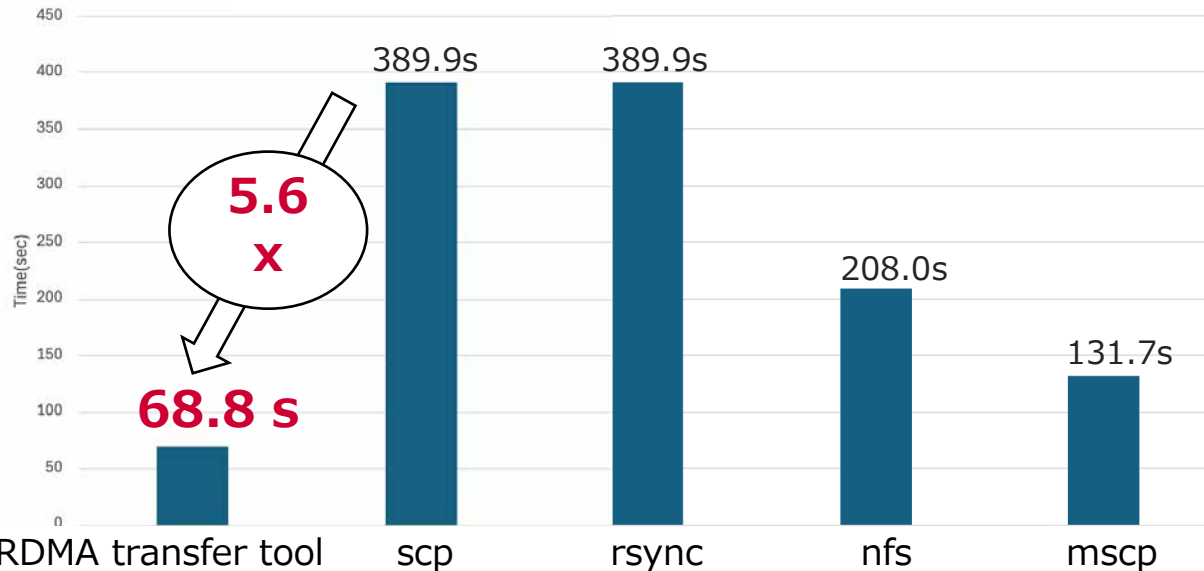
# 800Gbps High-Bandwidth Data Transfer

- We evaluated data transfer tests using the RDMA transfer tool (\*1).
- Eight RDMA-compatible NVIDIA ConnectX-6 Dx 100Gbps NICs were used.
- Large volumes of data were transferred between data centers at a wire rate of 800 Gbps.
- The performance of the RDMA transfer tool was compared with other methods (scp, rsync, nfs, mscp).

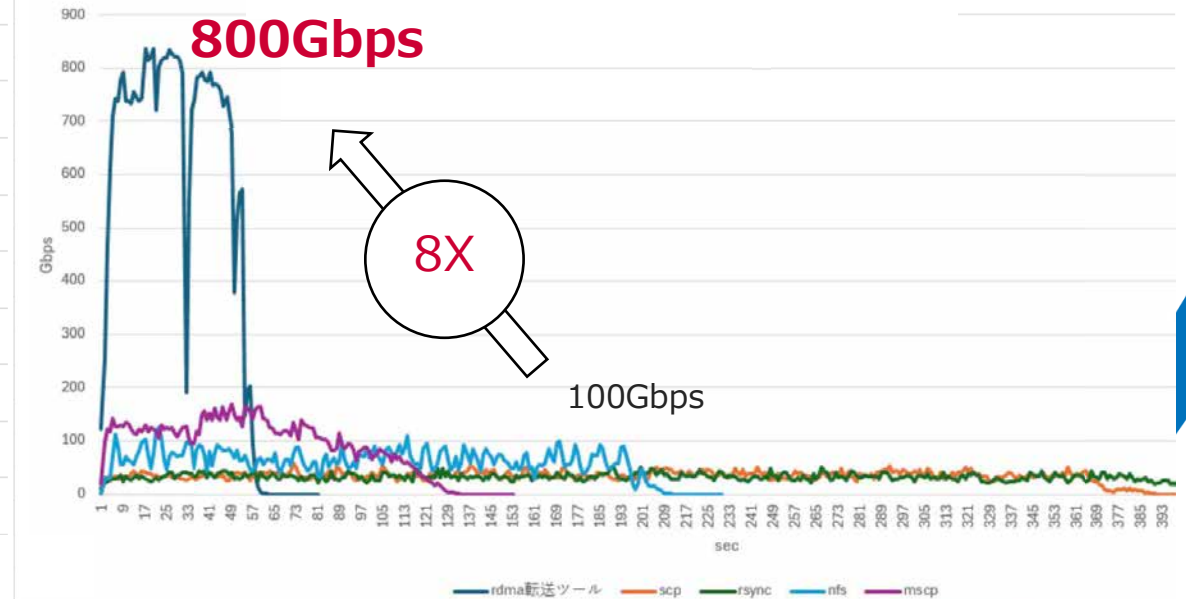
※1  
200GBx8=1600GB file  
transfer test  
putted on tmpfs(memory)

Compared with other TCP based tools, our RDMA transfer tool achieved:  
**5.6x** faster transfer speed and **8x** higher bandwidth performance

### Transfer Time for 1600GB

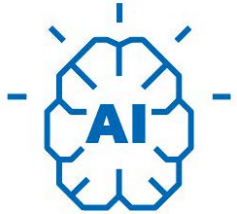


### Traffic



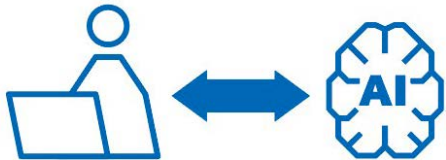
# Summary of the PoCs Results and Future Prospects

## Distributed Training



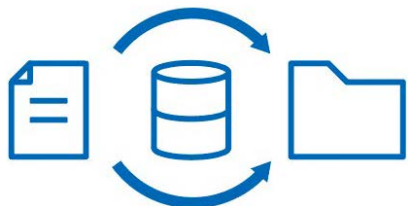
Training Time (LLM 7B)  
**1.005 x**  
 vs. Single-DC

## Distributed Inference



Inference Time (LLM 27B)  
**1.009 x**  
 vs. Single-DC

## Inter-Site Data Transfer



800G RDMA Transfer Speed  
**5.6 x**  
 vs. Conventional TCP Tools

- Further expand “GPU over APN” environments and its possibilities
- Reflect the know-how gained from the experiment into our commercial services, aiming for providing customers a **one stop solution.**
- **Invite partners** who are interested and co-create PoC together.

A series of blue lines of varying thicknesses that originate from the bottom left and fan out towards the right side of the image, creating a sense of motion and depth.

 <sup>NTT</sup> docomo **Business**