



Huawei Pioneers a New Computing Architecture for the AI Era: Making One Million Processors Work as One Computer – Summary of Roundtable Meeting Between Eric Xu and International Media Outlets on September 17

Descrizione

COMUNICATO STAMPA – CONTENUTO PROMOZIONALE

SHANGHAI, CHINA – Media OutReach Newswire – 30 September 2026 – During HUAWEI CONNECT on September 17, 2026, Eric Xu (Huawei’s Rotating Chairman) and Dr. Liao Heng (Chief Scientist of HiSilicon) had a Q&A session with journalists of media outlets to discuss Huawei’s Peerium Computing Architecture for the AI era as well as UnifiedBus (UB).

At last year’s HUAWEI CONNECT, Eric Xu gave a keynote speech and released Ascend chips and their roadmap (including the 950, 960, 970, and SuperPoDs and SuperClusters). He also announced the open release of the UnifiedBus protocol.

At this year’s HUAWEI CONNECT, Huawei announced the Ascend 950 and the Atlas 950 SuperPoD. While the industry is talking about SuperPoDs or supernodes, the ways of making these products differ hugely. The key of SuperPoDs is about allowing several thousand or even tens of thousands of processors to work as one computer.

Today, Huawei announced a SuperPoD based on the Ascend 950 chip. The SuperPoD is powered by a new computing architecture that Huawei has pioneered. This architecture now can make one million processors work as one computer. This architecture is called Peerium, and it can achieve strong scaling to the million-processor level through nested parallelism, unified memory addressing, and peer-to-peer interconnect. It extends the parallelism paradigm of Turing with the introduction of Nested BSP (Nested Bulk Synchronous Parallel).

This architecture also extends the von Neumann single-machine architecture and overruns the master-slave design paradigm that has prevailed for decades. Given that the available technologies in the industry could not support this architecture, Huawei invented a key interconnect technology – UnifiedBus. Ultimately, one million processors can truly work as one larger computer.

UnifiedBus is a high-speed bus that uses an open protocol to scale without a limit to connect CPUs, NPUs, memory, SSDs, interface cards, and switches. With UnifiedBus, peer-to-peer interconnect is achieved across compute, storage, and networking. This fundamentally overruns the traditional master-slave architecture.

The Atlas 950 SuperPoD is a product built upon the Peerium Computing Architecture. A SuperCluster with 256,000 computing cards is currently being deployed and tested.

Q: Dr. Liao, your paper mentioned a SuperPoD that's made up of 256,000 cards. Is that the upper limit of your technology? Can you tell us more about market demand for your products?

Liao Heng: 256,000 or roughly 200k, this number is not an arbitrary number.

First of all, the purpose of these large clusters is for training. It has a hardware infrastructure that can support the training of foundation models, which today are already in the 5T parameter range.

Over the next two years, there will be roughly 6 or 7 frontier AI labs in China, and all will aim at training foundation language models of sizes ranging from 10 to 40 trillion parameters. These numbers kind of match up to the memory capacity and the size of the SuperPoD. So you need such a scale of infrastructure to train these models.

The second thing is that in China, most data centers have a national plan on this power grid. So 200,000 is a relatively conservative number, given the design of the power delivery system in a single autonomous zone and single data center campus.

As I said, everybody must be quite familiar with the AI model race that's happening across the globe. And in China, there are at least 3 or 4 tier-1 players that are already well recognized and have achieved tier-1 status.

Q: What are your plans to encourage Chinese model providers to use Ascend chips for training?

Eric Xu: For the Ascend 950, we first launched the 950PR to the market. It's primarily used for AI inference, and currently the total volume of supply available in the market is not very big. The SuperPoDs based upon the Ascend 950DT are mainly used for AI training. Right now, they are under testing, and their large-scale supply will start at the end of this year or early next year. We are having extensive dialogues and discussions with AI model providers. The testing result of using the Ascend 950DT for training is pretty good. I believe that starting from next year, a lot of AI model training will be based upon SuperPoDs that use the Ascend 950DT. I think at the end of the day, it's not about how hard we are going to push model providers, but about how much supply capacity we will have. We can only supply based on how much that's produced. And we hope the value chain will expand its capacity faster to increase supply.

Q: What is Huawei's take on recent calls by leading AI model providers in the US to slow down AI development?

Eric Xu: We need to look at the level and cadence of AI development in China and the US. AI models in China are largely open-source models and their progress is relatively transparent. If we look at the mainstream model providers in the US, they have more computing power, so maybe only they themselves know where they are in terms of the level of sophistication of their models. The market feeling about AI risks might be weaker in China than in the US. Model providers in China need to speed up their pace to reach a level where they could also feel the risks from AI development. And then, maybe they will feel the same as the leading US model providers. But I always believe that we need to strike a balance between driving AI development and managing AI risks. At the very least, the bottom line is we need to ensure AI for good, not AI for evil.

Q: Does Huawei have plans to bring the Atlas 950 SuperPoD to other markets outside China? If yes, which markets are you going to target? Can you share with us data about your market share in China? What is your take on China's push for chip self-sufficiency?

Eric Xu: Since we don't have enough capacity to even satisfy the demand in China, we don't have plans to expand to the international market in a fully-fledged way. But indeed, there are several countries that have very strong demand. We are doing some testing, and we are providing some supply to those countries, but the volume is quite limited.

Right now, it's pretty hard to collect data about the market share of NVIDIA in China. But based on the data we have collected, Ascend has surpassed NVIDIA.

China's push for chip self-sufficiency is an inevitable path forward. China is a country of 1.4 billion people and it's an industrial powerhouse, so the way people work and live is closely tied to chips. Chinese people have a keen sense of urgency, and will not accept a future in which others decide whether or not we can have access to certain products. Even though our chips may be less advanced, at least their supply is assured, so that you don't have to worry about chip supply day in and day out. I think that's certainly the path forward. For the Chinese government, for the domestic industry, and for Huawei, the path forward is undoubtedly to push for full self-sufficiency in terms of chips and the entire semiconductor value chain. We also believe this will become a reality sooner or later.

To read the full Q&A session:

<https://www.apmultimedianeewsroom.com/multimedia-newsroom/huawei-pioneers-a-new-computing-architecture-for-the-ai-era-making-one-million-processors-work-as-one-computer-2>

Contatti:

Immediapress

comunicati@immediapress.it

Media contact: HuaweiCorporate.comms@huawei.com

COMUNICATO STAMPA - CONTENUTO PROMOZIONALE

Responsabilità editoriale di Immediapress

â??

immediapress

Categoria

1. Comunicati

Tag

1. ImmediaPress

Data di creazione

Ottobre 1, 2026

Autore

redazione

default watermark