Solution Brief

Infrastructure Processing Units (IPU)
Cloud Data Centers

Powering Infrastructure to Help Shape the Data Center of the Future

Creating a New Intel® IPU-based Computing Platform for Optimal Cloud Platform Management and Cost-effectiveness

Contents

Overview

In order to make cloud services more efficient and flexible to provide support for business innovation and intelligent transformation in various industries, especially to move enterprises and edge computing to the cloud, many cloud service providers (CSP) are working on the cloud data center architecture innovation and software tuning. Those efforts can eliminate the bottleneck of the infrastructure and make it more efficient and intelligent, thereby bringing users better service and performance, higher stability and security, and better management and control efficiency.

In response to this need, Intel has introduced a new infrastructure processing product – Intel® Infrastructure Processing Unit (Intel® IPU). Through a hardware acceleration engine composed of the programmable FPGA (Field Programmable Gate Array) logic chip plus Intel® Xeon® D Processor, or a standalone ASIC (Application Specific Integrated Circuit) chip, Intel® IPU can offload and accelerate networking, storage, and security workloads from the host server processor in a programmable manner, so that it not only unleashes the compute power of the cloud data center, but also offers more flexible resource allocation and faster transition to fully virtualized architectures of storage and networking for cloud service providers. In addition, as a new infrastructure acceleration engine, Intel® IPU can meet users' cloud service needs in more scenarios by working with Intel® Xeon® Scalable Processors.

As partners, Baidu AI Cloud and Intel share the same mission and vision in the future direction of cloud data center infrastructure. Through discussions and exchanges at the in-depth technical level, Baidu AI Cloud leveraged the Intel® IPU platform reference design and independently developed the Baidu Taihang DPU1.0 to actively address a series of challenges encountered in their real-world cloud service deployments. Through deployment and appllication to services in bare-metal server and cloud host scenarios, the new product has proved to be a reliable technical foundation for its future AI cloud services.

Background and Challenges

Innovative Enterprise Applications Demand Higher Capabilities from Cloud Data Centers

Benefiting from significant technology improvement and market expansion, cloud service is now providing service to many industries for their business operation, business expansion and technology development. It offers high-efficiency, elastic, and enormous compute power support so that more users can deploy their enterprise grade applications in the cloud environment flexibly. While applications with cutting-edge technologies such as 5G, Artificial Intelligence (AI) and edge computing are deployed in more industries, many new changes and trends have been seen in enterprise level cloud services, including:

- **More distributed cloud service deployment:** The development of edge computing and other technologies promotes the collaborative evolution of cloud services from the traditional central cloud to the "cloud edge", and more cloud services are deployed at the edge close to applications;
- **Broader application scenarios in cloud services:** Cloud services are becoming one of the IT infrastructures in more industries, playing the key role in new formats such as smart manufacturing and smart finance, and becoming a solid technical foundation for various enterprise-level applications;
- **More complex workload requirements:** Different workloads may have different requirements for resources. Some workloads require more compute power, while other workloads require larger storage or acceleration units. How to meet the requirements of different workloads has become one of the considerations for enterprises to move to the cloud;

Instance Diversification and Resource Pooling Drive Transformation to Data Center of the Future

The changes in the above deployment schemes and application scenarios have also introduced more challenges in managing and applying servers in cloud data centers. For example, the deployment of more virtual machines (VMs) complicates management tasks such as virtual machine management. At the same time, its capability output has gradually developed from providing traditional monolithic VMs to microservices, which further increases the complexity of cloud data center management tasks, which in turn takes up a lot of processor resources.

Today, cloud service providers not only need to efficiently manage instances such as virtual machines, microservices, and even bare-metal servers, but also need to accelerate networking, storage, and other infrastructure workloads in cloud data centers. The level of complexity and resource overhead continues to increase. How will the traditional infrastructure architecture still meet the requirements of the future data center? In other words, what technologies do we need to transform to the data center of the future?

As microservice models are increasingly adopted in cloud data center applications, it can be predicted that the future data center may have the following technical characteristics:

- The emergence and rapid growth of cloud-native applications has driven the need for dedicated infrastructure, and dramatically increased the agility of cloud services and the efficiency of cloud data centers;
- As more cloud services embrace microservices, it drives the development of distributed heterogeneous computing environments, and each microservice should run on a more suitable acceleration node;
- The widespread adoption of the microservice model has also led to a data center orchestration system that automates and manages the distribution of microservices among servers in the heterogeneous computing environment;
- The growing use of microservices, virtual machines, containers and container orchestration is driving the development of service meshes. Service meshes can simplify microservice-to-microservice communication and make it more efficient. Service meshes are now a standard part of the cloud-native stack;

At the same time, as resource requirements vary with different workloads and different time periods of users, resource pooling allows for better resource allocation to meet the requirements of elastic expansion.

Therefore, the future data center architecture will evolve to be capable of supporting different instances and increasing resource pooling, so as to make better use of the accelerated processing capabilities brought by heterogeneous computing. The core of this heterogeneous computing architecture will be composed of traditional processor platforms and infrastructure processing equipment such as IPUs introduced by Intel. In a data center based on this architecture, servers are interconnected with IPUs that perform networking and storage acceleration. Not only can the processing capabilities of traditional server nodes be enhanced with the support of dedicated computing nodes, but also storage, network services, and cloud service management including bare-metal servers and microservices can be accelerated.

■ Intel® IPUs Help Data Centers Meet Transformational Challenges

As various types of instance services and resource pooling are more widely used in cloud data center applications, storage, and networks, traditional data centers are transforming to data centers of the future to meet higher demands and challenges. In the architecture of traditional data centers, conventional network interface card (NIC) products are usually used to perform data traffic processing at the physical layer and data link layer, while functions at higher layers require the participation of computing resources such as processors. For example, in the data storage process, each I/O operation requires multiple context switches and memory copies by processors to exchange data between user mode and kernel mode in the "interrupt" mode.

It is even more prominent when cloud services are delivered. Not only virtual switching technologies such as Open vSwitch (OVS), storage transmission protocols such as RDMA over Converged Ethernet (RoCE), and corresponding data security technologies have been introduced into the system architecture, but also different cloud products such as bare-metal server and container cloud services have been derived. The introduction of these new products and technologies has made data processing in cloud data centers more complex.

This trend, along with the increasing data scale (port bandwidth is gradually evolving from 25G to 100G or even higher), is bringing huge challenges to cloud data centers, including:

- The growing speed of data processing is continuously higher than that of compute power, so its use of processor resources is increasing. Statistics show that 30% of compute power in data centers is used on traffic processing, which is evocatively referred to as Datacenter Tax.[1]
- To ensure efficient execution of core services, data centers have to acquire more processor resources, which increases the total cost of ownership (TCO) of cloud services, which in turn increases the cost of deploying and using cloud services for end users;
- Greater data scale and processing complexity also affect the performance of various infrastructure workloads in data centers, including network I/O operation, forwarding, storage, security, and management. It is preventing cloud data centers from achieving their performance optimization goals in increasing network throughput, reducing network latency, etc.

In this context, the ever-innovating cloud data centers urgently seek a new intelligent architecture to solve the above challenges. "Expansion and empowerment" of conventional NIC products is the industry's initial solution to solve the challenge in data processing overloading. This product, called SmartNIC, adds functions to the NIC, such as processing and forwarding various types of data for respective ports on the NIC card. By doing so, it releases more compute power, improves network performance, and reduces processing latency. However, due to the lack of compute complex, SmartNIC can only offload the data plane, and leaves the control plane and infrastructure management workloads such as the hypervisor to the host processing unit. It cannot achieve full offloading of infrastructure processing.

By further consolidation and analysis of the infrastructure capabilities of data center, Intel launched Intel® Infrastructure Processing Unit (Intel® IPU), which offloads data plane, control plane, infrastructure management workloads like Hypervisor and other cloud service infrastructure capabilities from the host server processors. Offloading those workloads from the servers, on the one hand, allows more valuable computing resources to be used for the CSP's key business. On the other hand, it also accelerates those capabilities by offloading them to dedicated devices. It effectively improves the execution efficiency of various data processing workloads, and comprehensively helps cloud data centers to build new acceleration engines for their infrastructure capabilities.

Solution: Leveraging Intel® IPU, Build a Strong Engine for Cloud Data Center

IPU is a brand-new product format. It can bring a higher level of security and control to the cloud data center of the future while covering the functions of SmartNICs. Its advantages include:

- It offloads intensive infrastructure workloads to the IPU and accelerates them, such as encryption/decryption and packet processing;
- In extreme cases, the IPU can offload the entire hypervisor, freeing up all cores of the host server processors, which is critical for bare-metal server service and supports applications and microservices;
- It reduces the overhead of hypervisors and infrastructure stack, and offloads storage stack in the host server processors, making more processor resources available for application and tenant workloads;
- It decouples infrastructure management from tenant applications for higher security and control;

- It creates possibility for a unified cloud management and control platform for both bare-metal servers and virtual machines.

As shown in Figure 1, Intel® IPU not only includes FPGA chip or dedicated hardware accelerators in ASIC, but it also includes general purpose processors to offload various infrastructure control plane functions. This combination of dedicated programmable hardware and general-purpose processors not only enables the acceleration of infrastructure capabilities in the cloud data center to meet the growing demand for network data processing, but also enables flexible management of control plane functions. It enables better system-level security, control, and management.

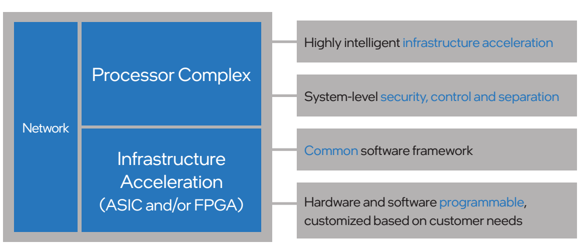

Figure 1. Intel® IPU Architecture

Let's take as an example, Intel® FPGA IPU C5000X-PL (code-named "Big Spring Canyon"), the commercially available hero IPU product. As shown in Figure 2, it provides two 25G ports (using SFP28 optical ports), which can provide up to 50G network throughput. Its core processing power is provided by Intel® Stratix® 10 DX FPGA chip and Intel® Xeon® D processor.

Figure 2. Intel® FPGA IPU C5000X-PL

Among them, Intel® Stratix® 10 DX FPGA programmable logic chip can give full play to its hardware programmability. This device features advanced architecture design, advanced packaging technology, more transceivers than the previous generation of FPGA, and support for hardcore PCIe Gen4 interfaces. Thereby, it helps achieve higher bandwidth, high throughput and low latency with customized design. By completing I/O virtualization, OVS forwarding and other tasks, it is adequate for offloading infrastructure management, networking, and storage workloads.

Intel® Xeon® D processor, a highly integrated design with excellent single-core performance, cannot only perform various control plane functions, but also supports Hypervisor. With the x86 compatibilities and the excellent ecosystem formed together with other Intel® architecture hardware, it enables users to quickly migrate system codes or application compatibilities, thereby improving their offloading efficiency.

Offloading Infrastructure Workload and Releasing compute power

With the above two chips, as shown in Figure 3, the cloud data center can effectively offload various workloads such as networking, storage, security, and infrastructure management from the host server processors to the IPU, thereby releasing compute power and accelerating various infrastructures capabilities. These capabilities include:

- **Network acceleration:** Offloading virtual switch software such as OVS, which hosts network I/O operation, data forwarding and other functions, from the host server processors to IPU, improving network throughput and reducing network processing latency;

- **Storage acceleration:** Moving storage interfaces and protocol stacks such as virtio-blk and NVMe-oF from the host server processors to IPU, improving storage elasticity and flexibility, and reducing system complexity and overhead;
- **Security acceleration:** Offloading functions such as encryption/decryption, and compression from the host server processors;
- **Infrastructure processing:** Offloading cloud service management functions from the host server processors to IPU, thus making the allocation and management of virtual machines, containers, or bare-metal servers more efficient.

In addition to introducing hardware with higher performance and flexible programmable features to optimize for and accelerate specific functions, and releasing more valuable compute power, Intel also endows IPU with a rich software ecosystem. Currently, Intel is leveraging Acceleration Development Platform (ADP) to optimize the ecosystem. It helps partners to rapidly develop and deploy their performance acceleration solutions in cloud data centers. Currently during the planning stage, the ADP platform will be supported by Intel® Open FPGA Stack (Intel® OFS), which will enable board hardware design, software, drivers, and associated technical design. Partners will be able to leverage these software tools and technical support to accelerate the development of IPU-related products and bring products to market quickly.

Coordinating with Intel® Xeon® Scalable Processors to Boost Compute Power for Cloud Data Centers

In the efficient cloud data center ecosystem solution provided by Intel, Intel® IPU can be used for offloading functions and accelerating performance of the infrastructure. In addition, Intel® Xeon® Scalable Processors can be introduced to further enhance the compute power. The effort to have more compute power while reducing its use will lead to faster data processing, greater bandwidth access capabilities and lower network latency.

The performance enhancements that the 3rd Gen Intel® Xeon® Scalable Processor brings to servers include:

- More cores and better architecture bring about a significant improvement in computing performance, which can effectively meet the needs for high-density computing tasks;

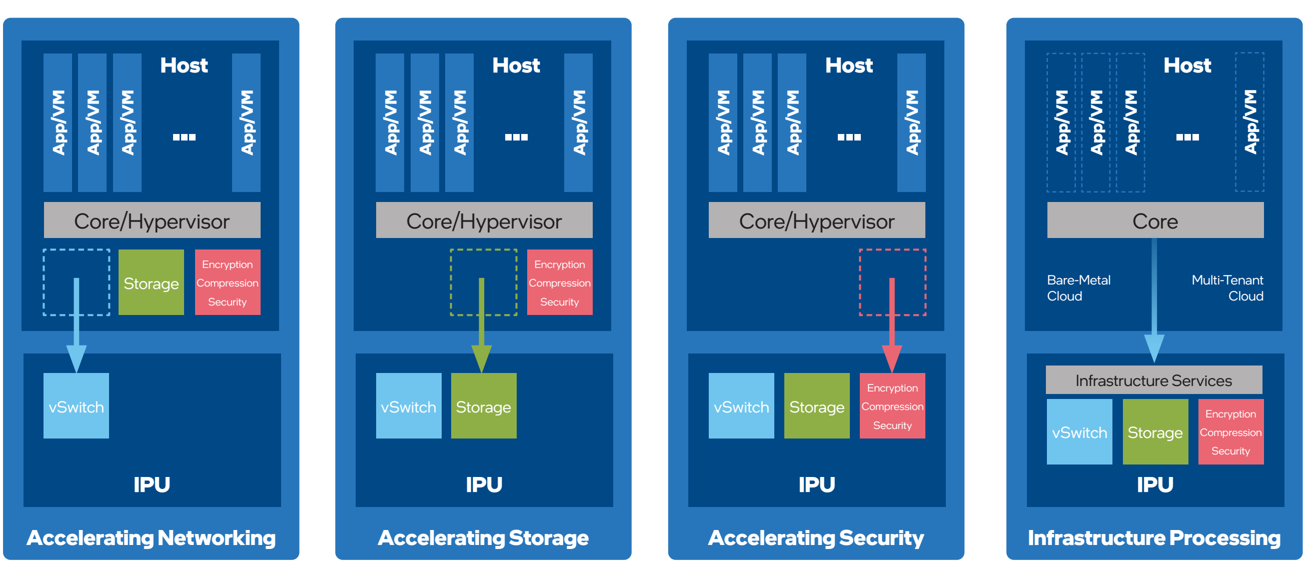

Figure 3. Intel® IPU use cases offloading different basic functions

- Support for more memory and PCIe Gen4, achieving higher I/O bandwidth per core;
- Multiple built-in enhancement technologies, such as Intel® Deep Learning Boost Technology (Intel® DL Boost), which can provide powerful acceleration in scenarios such as artificial intelligence.

Another advantage of introducing Intel® Xeon® Scalable Processors in this computing architecture is that it can effectively improve the efficiency and smoothness of workload offloading. Since both the host server and IPU use Intel®-based processors (Intel® Xeon® Scalable Processors and Intel® Xeon® D Processors), it is very handy to move applications running on the host server to IPU, even without compiling. For example, software such as DPDK and SPDK can be directly migrated from the host server to run on IPU, so as to greatly improve the efficiency in development, testing, and deployment.

This computing architecture is helping cloud data centers to provide more efficient throughput. Taking bare-metal servers as an example, they have the benefit of both high performance and security of physical server machines. That is why bare-metal servers are more and more adopted in critical business scenarios. However, traditionally, the cloud management of the bare-metal server is hosted by the processors, while the bare-metal server user monopolizes resources of the server. In that case, the traditional bare-metal server cannot provide compute power to the user as conveniently and flexibly as virtual cloud hosts.

By offloading of cloud management workloads to Intel® IPU, the bare-metal server can be physically isolated between the infrastructure management plane and tenants. Intel® IPU can also enable elastic service by using the hot-swappable feature of VirtIO devices. Together with the 3rd Gen Intel® Xeon® Scalable Processors, the bare-metal server can provide more compute power for users, which can help cloud data centers to provide users with more performance and security for their critical business. With all above, cloud data centers can improve the utilization efficiency of computing resources, providing high-performance and cost-effective cloud services.

Partnering with Baidu AI Cloud to Build More Efficient and Reliable Cloud Service

> **"Together with other technology waves, cloud computing is driving the digital and intelligent transformation in various industries. In this process, we launched our first proprietary DPU device – Baidu Taihang DPU 1.0. It can support Taihang's elastic bare-metal server services, network offloading, hot-swappable PCIe device feature and other features, allowing the virtualization overhead of cloud computing to be offloaded on the network card, which greatly improves the performance of our cloud products."**
>
> **Hou Zhenyu**
> **Vice President, Baidu Group**

As a leading cloud service provider, Baidu AI Cloud shares the same vision with Intel in the future direction of data center. That is, all overhead related to infrastructure workloads should

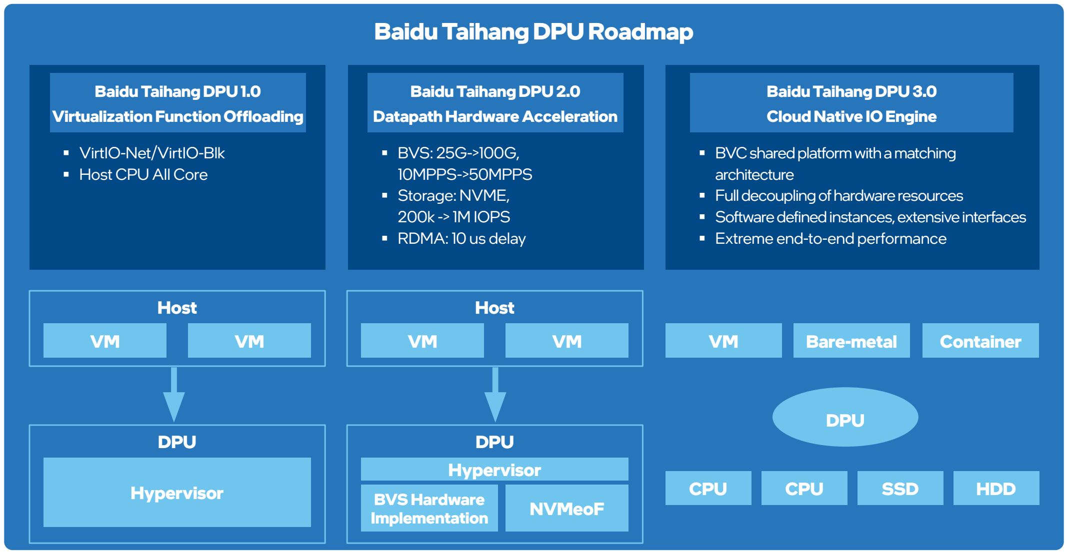

Figure 4. Baidu Taihang DPU Roadmap

be offloaded from the host server processors to dedicated infrastructure processing units. Such implementation will bring performance acceleration to free up compute power and optimize overall performance to create an efficient, secure and programmable data center of the future.

The practice from Baidu AI Cloud shows that cloud services, which continue to extend to various fields and host more and more critical business, are facing the challenge of requiring more powerful computing resources, higher bandwidth access, and lower network latency. Taking the leading AI cloud service from Baidu AI Cloud as an example, currently the demand for AI compute power has increased by several orders of magnitude, and the compute power demand of the Metaverse will increase by another 1,000 times.[2] Therefore, to support the above demand, a more powerful technical foundation is needed.

To achieve this, Baidu AI Cloud is committed to building a combined software and hardware infrastructure for cloud computing. By leveraging Baidu's rich experience in the cloud services field and meeting a large number of real-world needs, Baidu's proprietary Taihang DPU series can address cloud services computing power, networking, storage, security and other challenges.

As shown in Figure 5 and Table 1, Baidu Taihang DPU 1.0 is equipped with two 25G optical ports. It has Intel® FPGA and Intel® Xeon® D Processors at the core of its computing power to offload various virtualization functions. In addition to providing users with elastic networking and storage capabilities, it supports 1,024 hot-swappable devices. At the same time, by offloading the cloud management and control plane to Baidu Taihang DPU 1.0, virtual cloud hosts and bare-metal servers can be put in the same resource pool. In so doing, both bare-metal servers and virtual cloud hosts can be centrally managed and flexibly scheduled. Test result from Baidu AI Cloud shows that Baidu Taihang DPU 1.0 can achieve 10 million PPS forwarding rate and 200,000 IOPS in storage performance.[3]

Figure 5. Baidu Taihang DPU 1.0 Product Image

Baidu Taihang DPU Feature Highlights and Configurations
Feature Highlights - Enabling Elastic Bare-metal Server and Virtual Machine Deployment in All Scenarios
▪ Support for network offloading and acceleration, and uniform physical machine models for both bare-metal servers and virtual machines. ▪ Support for storage offloading and acceleration. Just like virtual machines, bare-metal servers can rapidly launch via boot from cloud disks, use cloud disks as data storage, as well as perform disk snapshots, data evacuation and other capabilities. ▪ Support for elastic hot-swappable feature. Key capabilities such as elastic NIC and storage expansion for bare-metal servers and virtual machines are provided. ▪ Support for hot upgrade, hot recovery, and hot migration (virtual machines).
Baidu Taihang DPU 1.0 - Virtualization Function Offloading
▪ Elastic NIC and storage, with 1,024 hot swappable devices ▪ Offloading the cloud management and control plane to DPU, so that virtual machines and bare-metal servers share the same resource pool ▪ 2*25G, 10MPPS, 200K IOPS

Table 1. Baidu Taihang DPU Feature Highlights and Configurations[4]

At present, Baidu Taihang DPU 1.0 has been deployed in various cloud service scenarios such as Taihang elastic bare-metal servers, and has received positive feedback from users. Practice has proved that Baidu Taihang DPU 1.0, which is independently developed by Baidu, can effectively remove the barriers in management and virtualization in existing cloud data centers and help Baidu AI Cloud maintain its advantages. In Baidu AI Cloud's plan, the DPU product will serve as the core component of its AI cloud service, helping it build a unified and elastic base for Infrastructure as a Service (IaaS). Relying on it, Baidu AI Cloud will provide efficient, reliable and high-performance services to their users in their various business innovations and intelligent industrial transformation.

Summary and Outlook

As Intel® IPU and products developed based on its reference design, has shown significant performance advantages and market value in the infrastructure innovation at cloud data centers, it gradually formed a sustainable ecosystem. Intel, together with its strategic partners like Baidu AI Cloud, is also developing more software and hardware optimized solutions for future cloud services. For example, the new Intel® IPU platform provides

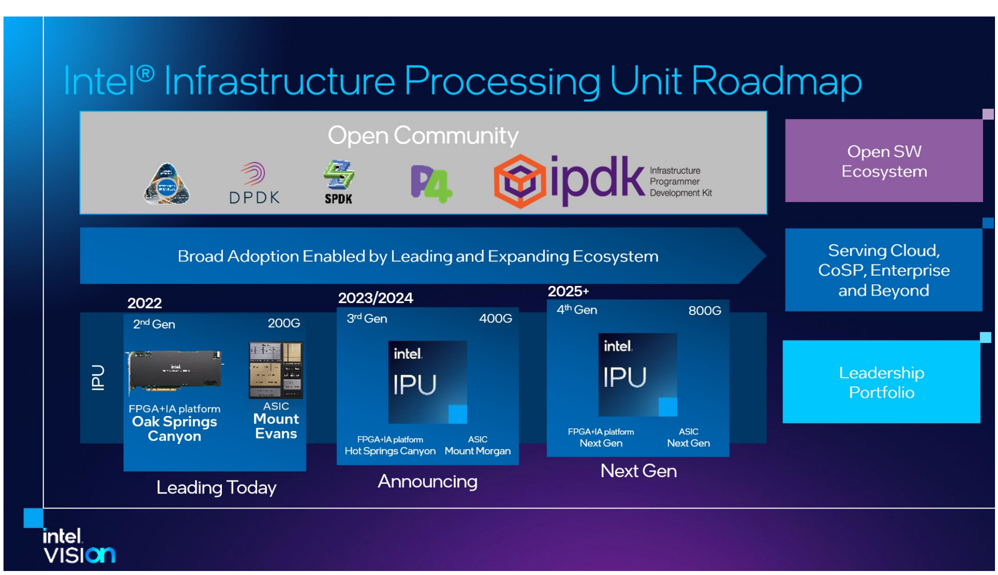

Figure 6. Intel® IPU Product Roadmap

greater network throughput. This fully decouples infrastructure management, and leads to end-to-end acceleration with extremely low latency.

Looking into the future, Intel also plans to continue to make efforts in the IPU product roadmap. As shown in Figure 6[5], Intel will launch two 2nd Gen 200G programmable IPUs in 2022 codenamed "Oak Springs Canyon" and "Mount Evans". Oak Springs Canyon is powered by Intel® Xeon® D Processors and Intel® Agilex™ FPGA chips, while Mount Evans is Intel's first IPU product based on the ASIC chip. Both products will have the infrastructure acceleration capabilities to offload 200G workloads, and both can work with Intel® Xeon® Scalable Processors closely. In the next two to three years, Intel will launch IPU products supporting 400G. In 2025 or later, Intel will launch a new generation of 800G IPU product based on FPGA and ASIC.

At the same time, with support from its rich open-source software such as the International Package and Development Kit (IPDK), SPDK, and DPDK, Intel is committed to the development of an open ecosystem, multilateral collaboration, and active community interaction. This will allow CSPs like Baidu AI Cloud, cloud service end users, and other partners to accelerate their development of IPU-related solutions, thus making greater value from the continuously innovating infrastructure of cloud data centers.

intel. 百度智能云

[1]This viewpoint is quoted from the article *Talking about the Computing System Change from the Rise of DPU* published in media: https://aijishu.com/a/1060000000228825.
[2]The data is quoted from public media reports: https://finance.sina.com.cn/chanjing/cyxw/2022-03-10/doc-imcwiwss5271286.shtml
[3,4]Baidu Taihang DPU product data from https://live.baidu.com/m/media/pclive/pchome/live.html?room_id=5073343376&source=h5pre
[5]Data quoted from the article "Intel Unveils Multi-generation Infrastructure Processing Unit (IPU) Roadmap at Intel Vision 2022" https://www.intel.cn/content/www/cn/zh/now/data-centric /rolls-out-multi-generation-ipu-roadmap-at-vision.html

Performance varies by use,configurationand other factors. For more information, see www.Intel.com/PerformanceIndex.
No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document. Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.
Any forecasts of goods and services needed for Intel's operations are provided for discussion purposes only. Intel shall bear no liability for any purchase made in connection with the forecasts published in this document. All information provided here is subject to change without prior notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.
Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No product or component can be absolutely secure. For more information, please check with your OEM or retailer, or visit intel.com.
Tests document performance of components on a particular test, in specific systems. Any differences in your system hardware, software, or configuration may affect your actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more information on performance and results of benchmark test procedure, please visit www.intel.com/benchmarks
Intel does not control or audit third party data. You should review this content, consult other sources, and confirm whether referenced data is accurate.
© Intel Corporation Intel, the Intel logo, and other Intel trademarks are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.