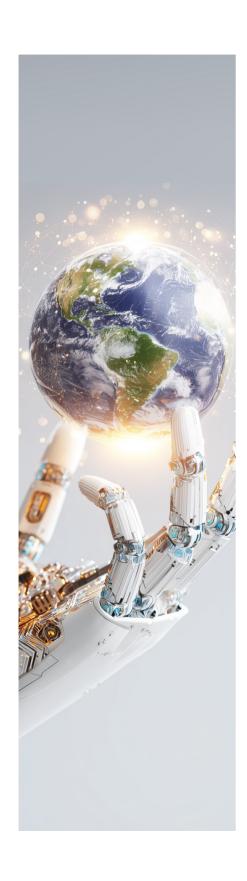


## **DIGITAL NAVIGATOR**



- P01 The Era of Soaring Computing: How H3C Leads the Infrastructure Revolution
- P26 Data Express: High-Efficiency Transmission of Massive Data Across Wide Area Networks
- P35 Reflections on the Evolution of Computing Power Network Architectures



## **CONTENTES**

The Era of Soaring Computing:How H3C Leads the Infrastructure Revolution

### 7 Flagship Innovations

- DDC: Redefining Intelligent ComputingNetworks with a Breakthrough Architecture
- Innovation in Liquid Cooling:Powering Sustainable AI in the Net-Zero Era
- 15 Full-Stack AI, Application-First:The H3C LinSeer Cube Delivers on the Promise of AI
- 18 A New Paradigm for IT Operations:One Map, One Brain, Two-Way Synergy
- The Wireless Iron Triangle:Empowering Enterprise Digital Transformation

### **1** Tech Insider

- 26 Data Express: High-Efficiency Transmission of Massive Data Across Wide Area Networks
- 35 Reflections on the Evolution of Computing Power Network Architectures

### **About H3C**



H3C is the industry leader in the provision of Digital and Al Solutions, committed to becoming the most trusted partner for customers in their business innovation and digital transformation. Through the deep layout of the whole industrial chain of "cloud-networking-security-computing-storage-term inal', we are constantly improving the capability of digitalization and intelligence. We offer comprehensive digital infrastructure capabilities, spanning across computing, storage, networking, 5G, security, terminal and related domains, and provide a one-stop digital platform that includes cloud computing, big data, artificial intelligence (AI), industrial internet, information security, intelligent connectivity, edge computing and unified operations, as well as end-to-end technical services.

We view technological innovation as the driving force behind our growth. Over 50% of our workforce is dedicated to R&D, and we have filed more than 16,000 patents, with over 90% of them being invention patents. Facing the Al-driven intelligence

era, H3C, with its profound technical and industry expertise, continues to deepen "Al in ALL" and accelerate "Al for ALL". By fully harnessing the multiplier effect of computing power and connectivity, we foster an open innovation to synergize computing power, data, large models, and ecosystems. This unlocks the full potential and efficiency of Al, making our products and solutions more intelligent, empowering smarter customer applications, and supporting the digital transformation and intelligent upgrading of diverse industries – including carriers, government, finance, healthcare, education, transportation, manufacturing, power, energy, internet, and construction.

"Shaping the Digital Future for a Better Life" is the corporate vision of H3C. Looking ahead, we will use digital solutions as the brush and the power of synergy as the palette, working closely with our customers and partners to cultivate wisdom for the era through practical and dedicated efforts!

# The Era of Soaring Computing: How H3C Leads the Infrastructure Revolution



■ By Liu Xinmin, Vice President of H3C, President of Technology Strategy Department

When the "Jevons Paradox" projects itself from the Industrial Revolution into the modern world of Al compute, the symbiotic relationship between increased efficiency and expanded demand is once again proven. As innovative models like ChatGPT and DeepSeek optimize to lower the computational cost of training, they empower more businesses and professionals to experiment and explore. This, in turn, fuels a broad industry transformation and drives a massive leap in demand for computing power, undoubtedly presenting a significant challenge to infrastructure.

## The Shift to "Centralized Training + Edge Inference"

The explosion of generative AI, particularly the remarkable emergence of DeepSeek, is fundamentally reshaping the compute landscape.

While more efficient models reduce individual training burdens, they introduce system-wide complexities. Global demand for intelligent compute

is growing at a compound annual rate of 78%, with over one-third of that growth attributable to applications built on DeepSeek-style models. Though DeepSeek's mixture-of-experts (MoE) architecture slashes per-inference computation to just 20% of traditional models, the explosion of open-source adoption, 5 billion daily API calls, has tripled total compute consumption.

We're witnessing a structural shift toward "centralized training and edge inference."

As resource requirements for training stabilize, hyperscale data centers are pivoting from pure capacity expansion toward efficiency optimization. Meanwhile, inference demand is exploding, driven by lower barriers to entry, with surging concurrent requests necessitating high-throughput parameter processing. Cloud providers are rearchitecting toward distributed inference, and data centers are pushing toward the edge.

Different applications now require tailored compute solutions. Autonomous driving and robotics rely on low-latency, high-concurrency inference servers, while sectors like fintech and medical imaging prefer turnkey, energy-efficient appliances. These workloads also place new demands on high-performance storage and Al-optimized networks. Optimized interconnects, improved storage protocols, and smarter networks have become essential to raising the efficiency of intelligent computing centers.

Going forward, we expect a hybrid ecosystem of both large and small compute centers. Major data centers will continue to anchor large-scale Al training, leveraging high-density racks and liquid cooling, while the surge in inference will fuel the growth of smaller, edge-oriented facilities, enabling SMEs to deploy Al locally.

## Four Key Vectors of Compute Infrastructure Innovation

How should enterprises respond to Al's insatiable appetite for compute? H3C believes future-proofed infrastructure will be shaped by advances in four critical domains: scale-up, scale-out, high-performance storage, and wide-area data transmission.

With models scaling to hundreds of billions of parameters, single GPUs can no longer carry the load. Scale-up architectures enable vertical expansion, connecting dozens to hundreds of GPUs to accelerate AI workloads. Scale-out solutions, meanwhile, support horizontal scaling across tens of thousands of cards, delivering the reliability required for massive clusters.



Unifying connection standards will be pivotal to entering the age of AI compute



Training performance also depends on high-speed data access. Storage must deliver high bandwidth, throughput, and low latency which achieved through tactics like hot/cold data tiering, NVMe-oF protocols, and all-flash arrays.

As intelligent computing evolves, we're seeing more use cases that involve long-distance data movement, like cross-center collaborative training, sample ingestion across regions, and geo-distributed inference. Wide-area networks introduce physical latency and reliability challenges, making packet-loss recovery and congestion control essential.

In response, H3C has introduced the "Computing × Connectivity" vision, a commitment to deliver efficient intelligent computing through diversified computational resources and standardized connectivity:

• **Scale-up:** H3C's G7 modular servers, developed in collaboration with more than 10 chip vendors, support cluster training, hybrid training-inference, and edge deployment.

Our UniPoD SuperPod systems enhance scale-up interconnect density by 300%, dramatically boosting communication efficiency and compute utilization for trillion-parameter models.

- **Scale-out:** We unveiled the world's first 51.2T 800G LPO/liquid-cooling switch, purpose-built for high-throughput, low-latency, energy-efficient Al networks. Our end-network converged scale-out solution incorporates a self-developed congestion control algorithm and global network orchestration to maximize cluster throughput.
- **Storage:** We rebuilt our storage software stack from the ground up to create the UniStor Polaris X20000, a next-generation AI storage platform delivering 80 GB/s bandwidth and 2M IOPS per node. It represents a leap in performance, protocol convergence, and architectural flexibility.
- Wide-Area Data Transmission: Using deterministic networking and our "Data Express" solution, with redesigned hardware and software stacks in which we achieve >90% bandwidth utilization over 5,000 km, outperforming traditional protocols by 5x. The result is a high-speed data railway.

Energy efficiency remains paramount. "The end of Al is energy,"as the saying goes. H3C uses Al-driven O&M and compute scheduling to consolidate and power down underutilized servers, boosting utilization and energy efficiency without compromising performance.

## Open Innovation: Depth in Technology, Breadth in Ecosystem

The next generation of computing and networking will look nothing like today's. As infrastructure complexity grows exponentially, H3C is not only building a full-stack portfolio from hardware to

industry solutions, but also advancing through open innovation under the dual principles of technical depth and ecosystem breadth.

The training of large models will require seamless interconnectivity, including within chips, between devices, across data centers, and over thousands of kilometers. Common standards will be critical to ushering in the age of pervasive compute.

That's why we're helping lead standards efforts with organizations like the Ultra Ethernet Consortium (UEC) and Unified Acceleration Link (UALink). We're committed to open, interoperable connectivity, where compatibility coexists with differentiation, enabling truly plug-and-play solutions.

Under our "Al in ALL" strategy, we're leveraging the multiplier effect of "Computing × Connectivity" through technology integration, standards collaboration, and real-world implementation to partner with industries and enable intelligent transformation.

In the new AI era, H3C is not just an infrastructure provider, we are architects of industry-wide advancement. We believe the future of connectivity is not merely about linking compute resources, but about interconnecting everything. True innovation lies not in chasing trends, but in defining the principles of what comes next. With technology as our engine, we are accelerating the compute revolution and helping build what's beyond.

• This article is based on an interview with ZDTech media.

## DDC: Redefining Intelligent Computing Networks with a Breakthrough Architecture

By Wei Yun, H3C Switch Product Line

Amid the ongoing digital and intelligent transformation, intelligent computing is rapidly expanding across sectors from scientific research to commercial applications, driving fundamental changes throughout industries. As training clusters scale to thousands or even tens of thousands of GPUs, networks must efficiently handle massive training data and optimize performance in complex data environments. This poses a critical challenge: how to build networks that keep pace with the demands of large-scale Al training.

#### **Key Challenges in Intelligent Computing Network**

The primary goal of an intelligent computing network is to meet the demands of Al training data transmission. The key requirements for networks in Al applications are as follows:

- ◆ **High Bandwidth:** All training relies on parallel processing across vast GPU arrays, requiring high-speed communication between nodes through network interfaces. With NIC speeds advancing from 200G/400G to 800G, network infrastructure must support 400G/800G connectivity to avoid bottlenecks.
- ◆ **Scalability:** Network architectures must flexibly scale from dozens to hundreds of thousands of GPUs, supporting continuous expansion of training resources without compromising performance.
- ◆ Lossless Transmission: Packet loss can severely degrade training efficiency or halt jobs entirely. Al workloads generate traffic patterns prone to congestion, rendering traditional ECMP inadequate. Advanced congestion control and dynamic load balancing are essential.
- ◆ **Multi-Vendor Compatibility:** Today's AI ecosystems incorporate GPUs and NICs from diverse suppliers. Networks must integrate these heterogeneous components seamlessly, giving users freedom to choose best-of-breed solutions.

Currently, several mainstream network solutions are used in Al scenarios, but each has its own pros and cons:

- Single Layer Chassis Switch + RoCE: CLOS-based chassis switches deliver non-blocking performance within the frame, suitable for Al in principle, but are limited in port density, power-hungry, and operationally complex.
- Box Switches + RoCE: While scalable in spine-leaf topologies, these suffer from hash polarization under ECMP, causing unbalanced traffic and congestion. Large-scale deployments also introduce operational overhead.
- **Proprietary Solutions:** Some vendors offer tightly integrated GPU-network bundles, limiting flexibility and often underperforming in multi-vendor environments.

Vendors that offer both GPU and networking solutions often utilize a Collective Communication Library (CCL) to tightly integrate them, effectively bundling them together and compelling customers to adopt the entire package. While such networking solutions may perform well within homogeneous computing environments of the same brand, they tend to exhibit a closed attitude or reduced technical adaptability when faced with demands for heterogeneous computing resources.

## The DDC Architecture: A New Paradigm for Intelligent Computing Networks

H3C's DDC (Diversified Dynamic-Connectivity) architecture tackles these challenges through groundbreaking design and open standards. Built on the Comware V9 OS, it combines cell-based forwarding and scalable fabric technology to overcome congestion and scaling barriers in Al workloads.

#### **◆** Architectural Innovation

DDC adopts a non-blocking transmission mechanism similar to chassis switches, but uniquely designs the switch fabric and line cards as independent box-shaped units. These units are then interconnected via an 800G Spine-Leaf architecture, which solves the port expansion and single-device power consumption issues of traditional chassis-based solutions.

Furthermore, H3C's integrated control and forwarding design simplifies the main control board mechanism of the chassis, streamlining the network into a two-tier NCP + NCF architecture, where the NCP serves as the line card and the NCF as the switch fabric. This decentralized architecture retains all DDC functionality while substantially improving system reliability and scalability. It is better suited for the evolution of cloud-native, distributed networks and eliminates the risk of management node failures.



Figure 1. DDC Architecture



#### **◆ Ultra-High Performance**

DDC uses VoQ+Credit (Virtual Output Queuing) and cell-based switching mechanisms to ensure traffic load balancing and congestion-free operation.

VOQ is a high-performance queue management technology for congestion control. Its core mechanism uses multiple virtual output queues corresponding to multiple output ports. Before sending data, the sending end requests bandwidth resources from the receiver. Once resource confirmation is received, it flexibly schedules queues to complete traffic forwarding. This allows for orderly traffic scheduling within the DDC system, leading to:

- **Non-Blocking Transmission:** Ensures packet forwarding within the DDC without any blocking through precise traffic scheduling.
- **Priority Scheduling:** Supports priority-based traffic scheduling, ensuring critical data flows are transmitted first.
- Improved Network Utilization: Avoids congestion and optimizes data flow paths to enhance overall network bandwidth utilization.

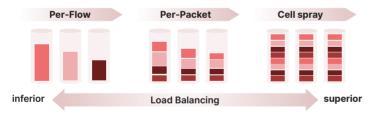


Figure 2. Cell Switching Mechanism

Cell-based switching is a key technology for achieving multi-path load balancing. Data packets are fragmented into equal-length cells by the NCP and sprayed uniformly across multiple paths, ensuring balanced traffic distribution. Compared to traditional flow-by-flow or packet-by-packet forwarding, cell switching achieves higher multi-link utilization, with a 100% load balancing effect.

Through the cell-based switching mechanism, DDC significantly reduces the risk of network congestion and improves network stability and throughput. The specific forwarding steps are as follows:

- 1.Server 1 sends a packet to Server 4. NCP 1 looks up its forwarding table to get the SystemPort and remote NCP encapsulation information index.
- 2.NCP 1 queues the packet in a VoQ based on the SystemPort. If it has a credit it segments the packet, adds a cell header containing information like ModID and PortID, and sprays it to the NCF through the cell port.
- 3.Upon receiving the cell, the NCF looks up its local cell forwarding table to find the egress interface and sends it to Server 4.
- 4.NCP 4 receives the cells and reassembles them based on their sequence numbers. It then uses the encapsulation index from the cell header to get the encapsulation information, adds the Layer 2 encapsulation to the reassembled packet, and sends it out the specified port to Server 4.



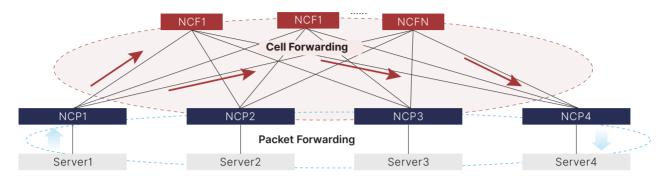


Figure 3. DDC Cell Forwarding

Through VoQ+Credit and cell-based switching, DDC achieves optimal lossless transmission and load balancing in large-scale networks, as evidenced by the test results below.

#### ♦ Open and Decoupled

DDC's open and decoupled capabilities are reflected in its support for decoupling compute facilities and its open heterogeneous capabilities.

## The combination of cell-based switching and the standard RoCE protocol serves as a key alternative to closed solutions.

In intelligent computing networks, the capabilities of the end devices (CCL/GPU/NIC) can impact overall network throughput and load balancing. For example, a packet-by-packet load balancing solution relies on the NIC's out-of-order reassembly capability.

DDC decouples the network from the end-device capabilities. It not only supports interconnecting NICs with various port encapsulations but also integrates the out-of-order reassembly function into the network side by using cell slicing and reassembly capabilities, eliminating the dependency on the NIC's capabilities.

Furthermore, in traditional solutions, the number of flow sessions (QP specifications) for different NICs can lead to varying hash effects. DDC's cell-based forwarding mechanism is independent of the number of QP samples. Even a single QP can be sliced and sprayed, perfectly overcoming the limitations of insufficient NIC QP specifications.

Therefore, the DDC solution perfectly addresses network performance issues caused by a lack of consistent end-device capabilities, breaking the binding of compute and network facilities and freeing customers from their reliance on closed solutions.

Regarding open capabilities within the architecture, H3C fully considered an open ecosystem from the very

beginning of the DDC design. The DDC control plane uses the standard BGP EVPN protocol, which not only enables auto-negotiation and auto-networking between network elements, reducing configuration complexity, but also supports heterogeneous networking with NCP/NCF devices from different vendors.

Based on DDC's open design at the protocol level, in 2024, H3C and its partners submitted multiple standard proposals related to DDC's open architecture and standardization in three areas: applicable scenarios, framework definition, and technical implementation. This is part of a commitment to driving open interconnectivity within the industry.

The draft clearly defines the DDC architecture's implementation framework as OSF (Open Schedule Fabric), which combines a scheduled network architecture with a traditional Ethernet network. This aims to achieve optimal resource utilization and rapid failover, solving the new challenges that Al applications bring to Ethernet networks through network-side optimization. The draft defines the following layers for the OSF network:

- OSF Management Layer: Responsible for monitoring, configuring, and maintaining OSF devices.
- **OSF Control Layer:** Responsible for maintaining OSF network topology information, congestion detection, and fast switching.
- **OSF Data Layer:** Responsible for encapsulating, forwarding, and decapsulating packets based on routing information from the control layer, and sending them to an authorized system for congestion control.

By standardizing the implementation and information exchange mechanisms of each layer, different network vendors can develop NCP/NCF devices based on this standard and ensure interoperability within the DDC architecture, achieving the same heterogeneous interoperability found in traditional Ethernet Spine-Leaf architectures.

Leveraging key capabilities in architectural innovation, extreme performance, and open decoupled, H3C launched the industry's first DDC architecture-based H3C S12500Al series switches, fully validating the design philosophy and delivering customer value.



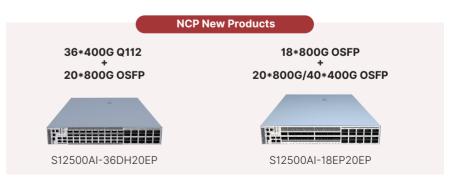


Figure 4. H3C S12500Al Series DDC Switches



#### Practical Achievement: H3C S12500Al Series DDC Switches

H3C's next-generation DDC switch, the S12500Al, delivers a universally applicable "box-to-box" architecture for hyperscale networking. The NCF provides 128-port 800G OSFP800, while the NCP downlink ports support both 400G QSFP112 and 800G OSFP800, fully meeting the connectivity requirements of mainstream NIC form factors.

The DDC network architecture offers expansion capabilities far beyond single-chassis systems, overcoming the inherent port capacity limitations of traditional modular devices and enabling flexible horizontal scalability. Under the two-tier NCF+NCP architecture, it supports up to 9,216 ports of 400G or 4,608 ports of 800G. In multi-cluster DDC interconnection deployments, the network can scale to a maximum of 70,000+ ports.

The DDC architecture has successfully passed performance validation by Tolly Group, based on multi-GPU NVIDIA Collective Communication Library (NCCL) benchmarks as well as large-model (Llama3) testing. NCCL All-to-All results demonstrate that the H3C DDC network architecture delivers an average 2.5% higher bus bandwidth (busbw) compared to InfiniBand solutions, with particularly notable advantages in large-message transmission scenarios.

Currently, H3C's DDC products have been deployed at scale in multiple industry customers. This practice fully validates that DDC technology, built on open standards, offers stronger adaptability and competitiveness in the market.

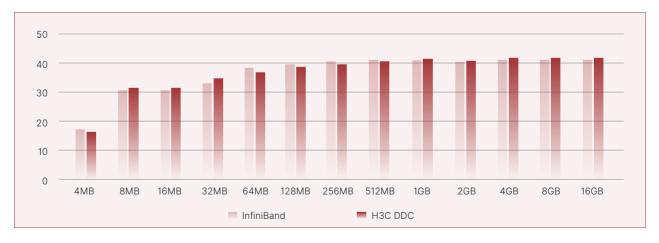


Figure 5. DDC Network Performance Test (All-to-All)

## Innovation in Liquid Cooling: Powering Sustainable Al in the Net-Zero Era

■ By Wu Hao, H3C Cloud and Computing Storage Product Line

The rapid development of artificial intelligence is reshaping the global economy, with computing power serving as the core resource of this transformation. As a result, it has become a key indicator of national competitiveness. As chip power consumption grows exponentially, traditional air-cooling technologies are reaching their efficiency limits. This, combined with strict net-zero goals, is pushing data centers toward a critical efficiency revolution.

In this transformation, liquid cooling technology, with its disruptive energy efficiency, has emerged as the core engine for building green computing centers and a sustainable path for Al-driven digital economies.

## Hidden Concerns Behind the Al Computing Boom

The widespread adoption of AI technology and mature open-source tools have significantly lowered the barrier to entry for computing power.

Cost-effective solutions based on models like

DeepSeek have driven explosive growth in AI

inference, but this technological boom has created a major problem: the conflict between soaring chip performance and unrestrained power consumption.

The combination of multi-core processors, high-speed I/O interfaces, and complex computing tasks pushes chip temperatures close to their critical threshold of 70-80°C. Experiments show that for every 2°C increase in temperature, a chip's performance can drop by approximately 10%. Traditional air-cooling struggles with frequent local hotspots, which not only shortens device lifespan but also makes it impossible for cooling systems to support single-cabinet power densities exceeding 100kW

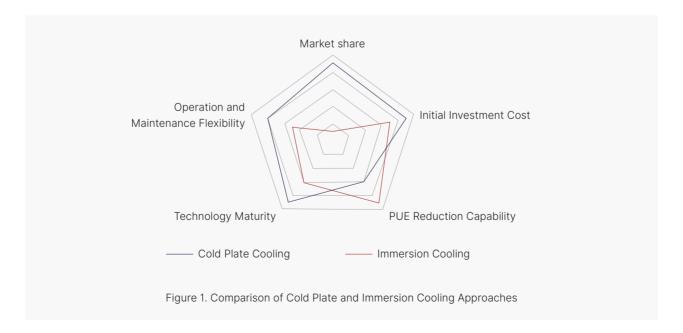
At the same time, national climate strategies impose rigid constraints on data center carbon emissions. This goal directly conflicts with the increasing energy consumption of data centers-some hyperscale data centers already consume more electricity than entire major cities. Balancing efficiency improvements with green transformation has become an urgent industry challenge.

#### **Liquid Cooling: The Green Game-Changer**

The rise of liquid cooling technology is a precise solution to this problem. Liquid mediums offer thermal conductivity hundreds of times higher than air, and their direct contact with heat sources precisely eliminates local hotspots, breaking the physical limits of traditional cooling. Data center liquid cooling has evolved over the years, forming an industry landscape with two main approaches: cold plate and immersion cooling.

Cold plate liquid cooling is the current mainstream for large-scale applications, accounting for over 90% of the market in 2024. This technology uses customized components to target high-heat-flux parts like CPUs and GPUs. Its core advantages include compatibility with existing air-cooled infrastructure, lower retrofitting costs (40%-60% less than immersion cooling), simplified maintenance, and significant energy savings. Current advancements focus on optimizing cabinet structures, moving from traditional decoupled designs to high-density integrated forms.

Immersion cooling fully submerges devices like servers in a dielectric coolant with high insulation and low boiling points. This achieves near-zero thermal resistance, making it ideal for ultra-efficient data centers with a PUE below 1.1. H3C's phase-change immersion cooling solution, for example, leverages the latent heat of vaporization to achieve a single-U cooling capacity exceeding 3.5kW, stably supporting high-power chips up to 1376W. This technology also ensures zero coolant evaporation through dual-seal designs, reducing coolant loss by 90% while enhancing device reliability. However, immersion cooling has high initial costs and challenges with coolant compatibility and operational complexity, making it better suited for new data centers or noise-sensitive environments.



## H3C Liquid Cooling Solutions: The Green Engine for Computing

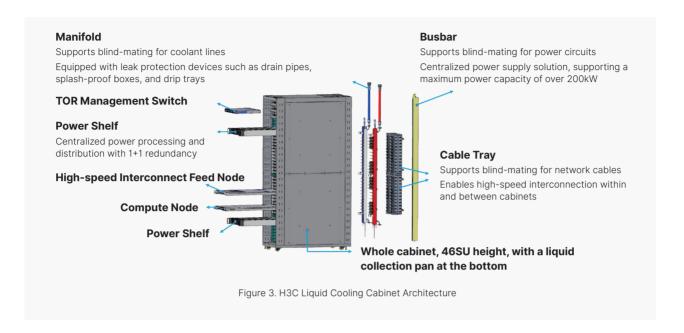
In a landscape of surging computing demand and limited space, high-density deployment is inevitable. H3C introduces a portfolio of high-density liquid cooling and Al liquid cooling cabinets that feature ultimate cooling, ultra-high density, simplified maintenance, and high-speed interconnectivity. Through the deep integration of architectural innovation and liquid cooling, these solutions provide a high-performance, reliable, and easy-to-deploy intelligent computing foundation for next-generation data centers.



Figure 2. H3C Liquid Cooling Cabinet Solution

The H3C high-density liquid cooling cabinet uses a 1U2N high-density server architecture, deploying dual-node quad-socket servers within a single U space. Each cabinet can support up to 32 server units, delivering 256-core cluster computing power and up to 120kW of power. The cabinet uses a 46SU (Scorpio U) standard height, which significantly improves I/O expansion through an optimized layout. Thanks to factory pre-assembly and full-cabinet delivery, users can achieve plug-and-play deployment, drastically shortening service launch cycles. The liquid cooling system covers core components like CPUs and memory, with cabinet-wide liquid cooling accounting for 75%-80%. The system includes a three-tier leak detection mechanism and an anti-spray design, along with optional rear door heat exchanger (RDHx) modules, further enhancing cooling efficiency and safety.

For AI, the liquid cooling cabinet solution enables full interconnectivity of 80 GPUs per cabinet via a 112G high-speed backplane. High-speed cable interconnects between dual-cabinet switch nodes break traditional bandwidth limits and fully unleash computing performance. Each cabinet can configure two 32-GPU clusters, combining centralized power and a high-density layout with blind-mate water, power, and network buses that adhere to ORV3 standards for simplified maintenance. This significantly improves operational ease, space utilization, and energy efficiency, delivering unparalleled computing density and stability for AI training-a future-proof choice for data centers.



## **Building Green Computing Centers: From Innovation to Collaboration**

Constructing green computing centers is a systemic project that requires alignment across technology, policy, and the ecosystem.

Technology Innovation: Technical approaches are the foundation for diverse use cases. Cold plate solutions cover 75%-80% of heat sources, balancing efficiency and retrofitting costs, while immersion solutions unlock ultimate computing power with 100% liquid cooling. H3C's "ALL in GREEN" full-stack liquid cooling portfolio, spanning cold plate and immersion technologies, forms a closed-loop system from chips to cabinets. H3C has also launched innovative products like full cold plate liquid-cooled servers, cold plate and immersion switches, and various liquid cooling cabinets to ensure end-to-end product support from design to implementation.

Lifecycle Management (LCM): Leveraging over a decade of industry expertise, New H3C Group has built a comprehensive LCM service system covering data center planning, design, implementation, verification, maintenance, and energy efficiency optimization. H3C has served over 500 clients, including internet giants and state-owned banks, with over 800 global case studies, patented LCM technologies, and end-to-end service capabilities. H3C drives intelligent iteration, flexible resource expansion, and low-carbon operations across the data center lifecycle, becoming a key enabler of sustainable growth for its customers.

#### **Ecosystem Collaboration and Standards: The**

large-scale implementation of liquid cooling requires deep cooperation among upstream infrastructure providers, midstream ICT and chip manufacturers, and downstream service providers. H3C actively participates in the formulation of over 40 technical

standards-including national, industry, and group standards-providing critical technical guidance for the green development of data centers. In 2024, H3C and Intel jointly launched the G-Flow synthetic oil solution, which addresses the heat dissipation challenges of high-power chips through innovative architectural design. For liquid-cooled cabinets, H3C's adherence to the ORV3 open power standard significantly reduces comprehensive development costs and shortens the product R&D cycle, demonstrating the enabling value of unified technical standards for industrial upgrading.

#### **Conclusion: From a Cooling Revolution to Green Infrastructure**

The evolution of liquid cooling technology is a microcosm of computing centers transitioning from high energy consumption to near-zero carbons. It is no longer just a cooling method; it has evolved into essential infrastructure that supports the green development of the digital economy. Through the continuous iteration of its full-stack liquid cooling solutions, H3C has not only verified the technical feasibility of green computing centers but has also unlocked diverse technical pathways from decoupled to high-density modular deployment via architectural innovation.

In the future, as policy incentives take effect, technology costs decline, and ecosystems mature, liquid cooling will push the boundaries of computing economics beyond physical and energy constraints. It will provide a solid foundation for humanity's march toward carbon neutrality. In this Al-driven transformation, green computing centers will stand as the best testament to technology empowering sustainable development.

## Full-Stack AI, Application-First: The H3C LinSeer Cube Delivers on the Promise of AI

As the application of Al becomes more widespread, large model technology faces a dual challenge: pushing innovation while ensuring practical implementation. Although the underlying architecture continues to advance computing power, three critical bottlenecks persist on the application side.

- ♦ **High Deployment Barriers and Complexity:** Traditional large model deployment requires intricate hardware and software adaptation, along with lengthy testing and verification, which delays business launches and fails to meet rapid business demands.
- ◆ Insufficient Scenario Adaptation: General-purpose models often fail to meet the needs of specific industries. Extensive customization is required, which is costly and time-consuming. We lack end-to-end solutions that integrate industry data for fine-tuning and create low-barrier, cost-effective Al applications.
- ◆ Fragmented Workflows and High O&M Costs: Performance tuning and continuous iteration after deployment require specialized technical teams, creating an "implementation equals liability" dilemma.

## Full-Stack Intelligence, Ready Out-of-the-Box: LinSeer Cube All-in-One Al Machine Accelerates Al Adoption

To address these three challenges, H3C is launching the "LinSeer Cube". Leveraging its full-stack AIGC ecosystem and deep industry expertise, this All-in-One AI machine is the first in the industry to deeply integrate intelligent computing, algorithms, and governance. Powered by robust computing, it fully supports DeepSeek models and is compatible with various GPU brands globally. LinSeer Cube redefines large model implementation with three core values: out-of-the-box full-stack delivery, vertical industry customization, and end-to-end lifecycle services.



The H3C LinSeer Cube

#### ◆ Full-Stack Delivery: A Ready-to-Use Intelligent Engine

H3C's "hardware + algorithm + platform" full-stack Al solution is built on its ICT ecosystem and "Intelligent Computing × Connectivity" vision.

Infrastructure: The machine comes with a pre-integrated intelligent computing cluster, featuring unified orchestration for networking, storage, and security, ensuring a seamless setup.

Algorithm Services: It is preloaded with DeepSeek base models and industry-optimized versions, ready for further distillation and fine-tuning.

Application Enablement: An embedded H3C Al Store platform provides over 1,000 free open-source datasets and models. Users can also access commercial resources with a single click, cutting asset acquisition time from days to hours.

This turnkey delivery model enables rapid deployment with a web-based front end and standardized APIs, reducing application launch time from one week to under a day.

#### ◆ Scenario Intelligence: Vertical Industry Customization

H3C focuses on government and enterprise markets, using platforms like pilot labs to tailor models for finance, healthcare, and manufacturing, bridging the final mile of implementation. The LinSeer Cube supports DeepSeek and other models for distillation, fine-tuning, inference, and workflow orchestration. It enables end-to-end industry Al delivery with modular model tuning and flexible workflows.

#### ◆ Continuous Evolution: End-to-End Lifecycle Services

H3C's "hardware + model + full-stack + service" offering includes pre-consulting (customized implementation plans), mid-deployment (elastic scaling and API encapsulation), post-O&M (AI-driven automation), and value-added services (computing optimization, model tuning), ensuring localized deployment and support.

#### Application-First, Accelerated Implementation: H3C Empowers Industries to Harness Al Benefits

The LinSeer Cube's out-of-the-box approach not only meets diverse Al deployment needs but also ensures data and business security through on-premises models, providing unmatched stability even in low-connectivity environments. As DeepSeek expands into key industries, our tailored solutions and ecosystem synergy help clients bypass complex setup and fine-tuning, accelerating Al development.

Partnering with industry leaders, H3C delivers end-to-end Al capabilities across government, education, healthcare, finance, and enterprises, achieving successful localized DeepSeek deployments with exemplary results.

In education, H3C has enabled many top-tier Chinese universities to apply AI in research and teaching, enhancing efficiency while safeguarding academic data privacy.

In healthcare, hospitals leverage it for medical model training and diagnostics, improving clinical and operational workflows.

In finance, the AI empowers banks in smart customer service, risk control, and marketing, achieving vertical integration for core financial use cases. In government, some big cities in China deploy AI for document processing, public services, and decision-making, creating replicable smart governance models.

For SOEs and manufacturing, DeepSeek's NLP and task-decomposition capabilities optimize R&D, production, and decision-making, driving intelligent transformation.

H3C has also actively collaborated with carriers to advance multiple use case tests based on the LinSeer Cube, aiming to bridge the final mile of scenario-based large model applications in the carrier enterprise sector.

The H3C LinSeer Cube meets the stringent demands of enterprise scenarios, and its powerful intelligent computing infrastructure fully unleashes the potential of large models. Looking ahead, H3C will continue to embody its philosophy of "Anchored in Pragmatism, Empowering the Era with Wisdom" by accelerating product evolution and refining its partnership framework. This will enable high-quality large models like DeepSeek to drive comprehensive industry digitization.

## A New Paradigm for IT Operations: One Map, One Brain, Two-Way Synergy

■ By Ao Xiangqiao, General Manager of Intelligent Management and Operations
Product Line of H3C

Nowadays, Artificial intelligence is advancing at an unprecedented pace. In the past, high training costs and demanding performance requirements for training chips limited widespread adoption across industries. However, with the maturation of technologies like optimization algorithms, mixed-precision training, gradient compaction, and model distillation, these barriers have been significantly lowered. The reduction in costs and optimization of computational thresholds have not only accelerated the widespread adoption of AI but also created unprecedented opportunities for development across various sectors.

In IT operations, for instance, advancements in computing power and algorithms have paved the way for operations large models to become a reality. These models are now maturing and are being deeply integrated into daily operations management. However, truly improving reasoning and decision-making efficiency still relies on a high-quality digital foundation. To address this need, H3C proposed "Global Ops Topology" is gaining industry attention.

This solution integrates all network resources, providing critical data and tools for operations large models while enabling mutual empowerment through a two-way synergy. The Global Ops Topology continuously optimizes data to enhance the inference capabilities of operations large models, while the evolution of these models further drives innovation and deeper application of the map. Within this collaborative system, the Global Ops Topologymust possess a series of essential capabilities to ensure the efficiency and sustainability of intelligent operations:

**Panoramic Visibility:** Imagine an ultra-low-level operations map that clearly displays the operational status of every component in the system. This allows operations personnel to grasp the situation at a glance, providing an accurate basis for decision-making.

**Predictive and Troubleshooting Capabilities:** Like any complex machine, IT systems are prone to faults. With these capabilities, it can proactively identify potential issues, quickly locate problems when they occur, and resolve them efficiently to minimize losses.

**Efficient Synergy with Large Models:** This integration optimizes user interaction, enhancing the convenience and intelligence of operations tasks. It also enables deep empowerment of intelligent operations through large model-driven intelligent decision-making and automated workflows.

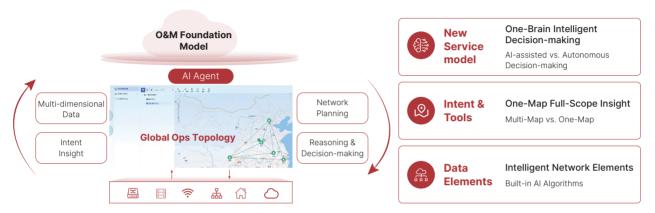


Figure 1. "One Map" collaborates with "One Brain" to illustrate the blueprint of intelligent transformation across industries.

#### Massive ICT Information, Clear at a Glance

TheGlobal Ops Topology, as a new benchmark in operations visualization, stems from a profound understanding of modern IT complexity.

Technologies like cloud computing, big data, and microservices have fully penetrated modern IT, and cross-domain collaboration is now frequent.

Traditional operations software is simply no longer adequate. Operations teams urgently need a tool that can intuitively visualize the system status from a global perspective to aid in problem location and resource optimization. TheGlobal Ops

Topologyemerged to meet this demand.

Based on a distributed system architecture, it integrates multi-source data across technology domains, such as device management, monitoring, configuration management, and log analysis to build a unified, real-time system view. Adhering to the "what you see is what you get" principle, it enables operations personnel to clearly grasp the IT environment and make precise decisions.

The architecture of the Global Ops Topologyis divided into four layers:

- 1.The data acquisition layer collects real-time monitoring data from IT components.
- 2.The data processing layer cleanses, integrates, and analyzes data using big data and machine learning technologies to extract key information.
- 3. The presentation layer displays analysis results in a graphical map interface, supporting interactive operations like zoom, drag, and filter, allowing users to flexibly view information across different hierarchies and dimensions.
- 4. The application layer provides rich function interfaces, such as alarm notifications and automated script execution, to meet the needs of various operations scenarios.

In practice, the Global Ops Topologyis widely used for fault diagnosis, performance optimization, and resource management. It quickly locates problem nodes on the map and identifies fault causes through correlative analysis. It identifies performance bottlenecks, guiding optimal resource allocation and

tuning, and enhances IT resource management efficiency to reduce waste.

## **Efficient Fault Diagnosis Across All Scenarios**

For network operations personnel, identifying the root cause of issues amidst massive data, complex network structures, and intricate fault information remains a highly challenging task. Efficient, intuitive visual fault diagnosis and troubleshooting capabilities are an urgent need to resolve these challenges.

In campus scenarios, the Global Ops
Topologyemploys in-depth end-to-end path analysis
combined with root cause tracing technology for
connectivity faults. Supplemented by Tracert
detection, it accurately pinpoints the root causes of
connectivity issues. It integrates network topology
for visual presentation, enabling operations
personnel to take timely and effective repair

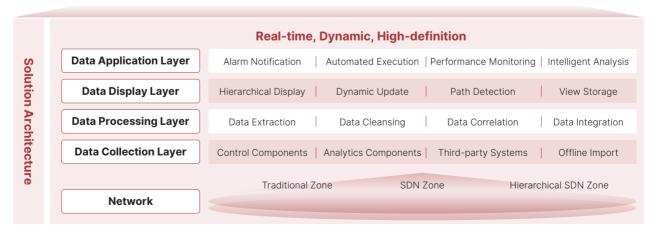


Figure 2. The Architecture of Panoramic O&M Solution



measures, significantly improving the efficiency and accuracy of campus network operations.

In data center scenarios, the path detection capability of the Global Ops Topologyplays a critical role. It uses data plane verification (DPV) technology to verify whether forwarding behavior aligns with business intent through formal modeling. Once anomalies occur, the Global Ops Topologycan precisely outline the fault path and identify the cause, ensuring stable data center operation.

## Two-Way Synergy Elevates User Experience

The tight synergy between the Global Ops
Topologyand operations large models is the core
element for optimizing user experience. Their
seamless integration brings significant advantages to
the field of intelligent operations.

Intelligent Question Answering (IQA): The IQA function provides convenient, efficient information retrieval. Operations personnel can query the system in natural language, such as, "Which network devices experienced loss of connectivity in the past 24 hours?" The system automatically analyzes and highlights the problematic devices on the map, providing detailed reports on outage times and frequency, making problem location intuitive and rapid.

**Automated Troubleshooting:** This function greatly simplifies fault handling workflows. When the system detects a fault, the operations large model

automatically generates troubleshooting scripts or recommendations based on fault type identification and historical data analysis. Operations personnel need only acknowledge and execute them, reducing human intervention and error probabilities.

**Intelligent Optimization Suggestions:** This function offers data-driven recommendations for resource optimization and configuration adjustments through in-depth analysis of historical and real-time data, enabling more optimal and effective use of system resources.

## Practical Application: Solving Financial IT Operations Challenges

The financial industry is a prime example of a highly complex IT environment with demanding maintenance requirements. With the rapid development of digital finance, business scale continues to expand, and multi-data center collaborative operations have become the norm. Traditional maintenance methods struggle to adapt to this complexity, resulting in low efficiency, limited stability, and numerous challenges. For instance, under a multi-data center architecture, business access paths are long, involving multiple device nodes and physical regions. The lack of unified management and visualization tools makes network topology maintenance difficult and fault resolution cycles excessively lengthy. Meanwhile, independent maintenance analyses across different networks and third-party systems hinder integrated management, further increasing IT maintenance complexity.

The introduction of the Global Ops Topologyprovides a groundbreaking solution for IT maintenance in the financial industry. It fully integrates network-wide resources and offers a highly visualized maintenance perspective, helping maintenance personnel monitor network, device, and business flow directions in real time, significantly improving efficiency. Its path detection and simulation functions enable enterprises to quickly identify and predict faults, reducing the risk of service outages. Moreover, leveraging the intelligent analysis capabilities of the operations large model, the Global Ops Topologyachieves precise closed-loop management, fundamentally enhancing the stability and reliability of financial application systems and helping enterprises maintain robust operations in the digital era.

#### **Closing Remarks**

The rapid advancement of AI technology is driving IT maintenance toward intelligent and highly efficient evolution. The synergistic effect of technological progress and tool innovation has formed a clear trend. The technical architecture, with the Global Ops Topologyas its foundation and the operations large model as its engine, is reshaping the underlying logic of IT maintenance systems, ushering in a new paradigm of intelligent decision-making and precise execution for the industry.

## The Wireless Iron Triangle: Empowering Enterprise Digital Transformation

As enterprise digital transformation accelerates, network infrastructure and IT operations teams are under unprecedented pressure. IT managers face several persistent challenges in deploying and maintaining wireless campus networks:

## 1. Ensuring seamless access for a massive number of mobile terminals (TEs)

With the proliferation of mobile work and IoT adoption in enterprise campuses, a wide variety of mobile terminals-diverse in type, vendor, hardware configuration, operating system, access behavior, and roaming characteristics which require stable wireless connectivity. Guaranteeing optimal access, usage, and handover for these devices has become a primary concern for IT managers.

## 2. Maintaining peak wireless performance across diverse usage scenarios

Wireless networks surpass wired networks in offering unrestricted mobility, enabling ubiquitous access. However, this advantage also imposes higher operational demands. Networks must

self-adjust and adapt rapidly to sustain peak performance across multiple contexts, such as high-density meeting rooms, training facilities, logistics warehouses, factories, mobile healthcare wards, remote branches, and rail transit systems.

## 3. Self-optimization of wireless networks to support emerging services

Traditional wireless networks primarily supported internal data sharing and Internet access. In the era of digital transformation, however, new services often sensitive to latency, jitter, and roaming delays and depend on wireless performance. For instance, cloud desktops and video conferencing are disrupted by packet loss, while delayed roaming can critically affect mobile healthcare and automated guided vehicle (AGV) operations.

#### The Wireless Iron Triangle Framework

To address these challenges, H3C proposes the Wireless Iron Triangle framework, which comprises three fundamental pillars:

#### 1. High-quality, scenario-specific WLAN products

Digital transformation diversifies campus environments, requiring WLAN devices tailored to specific scenarios. For example:

- High-density APs increase user capacity through multi-radio designs and optimized algorithms.
- Hotel faceplate APs emphasize aesthetics, thermal design, and energy efficiency.
- Rail transit APs strengthen vibration resistance, reliability, antenna performance, and high-speed switchover algorithms to ensure secure communication.

#### 2. Optimal deployment for each scenario

Wireless deployment is significantly more complex than wired infrastructure. Coverage, user density, bandwidth, building layout, physical obstructions, power supply, cabling, and AP operating conditions all affect the final outcome. Even identical AP models from the same vendor yield vastly different results depending on deployment strategy, much like cooking, where identical ingredients produce varied flavors depending on heat, cookware, and technique.

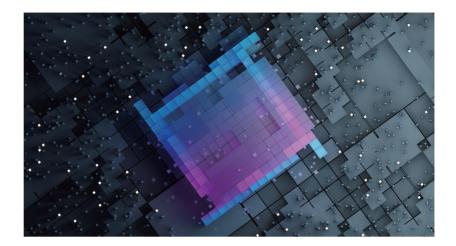
A classic example is hotel and dormitory coverage. Historically, hallway-installed APs minimized unit count but relied on high transmission power to penetrate walls, often leading to interference and poor in-room coverage. H3C's innovative "Terminator" solution places optimized APs inside rooms, eliminating wall penetration. This approach ensures dedicated, high-quality coverage per room while walls act as natural interference barriers. Currently, this deployment model has become the industry standard, replacing outdated hallway designs.

#### 3. The 4i Cloud-Network-Terminal Intelligent Operations Solution

Just as preventive maintenance enhances vehicle performance, intelligent wireless operations enable proactive optimization. Leveraging more than 20 years of experience and over 10 million deployed devices, H3C introduces a 4i-based solution, comprising four modules:

**iRadio:** Expands resource management from 2D AP-centric mapping to 4D optimization by integrating terminal and cloud intelligence. Terminals report real coverage data, while cloud AI analyzes long-term usage patterns, predicts demand, and preemptively allocates resources.

**iStation:** Monitors terminal access and roaming in real time, ensuring seamless service continuity. It identifies terminal types, vendors, hardware configurations, and behavioral patterns, enabling fine-grained optimization and differentiated roaming strategies.



**iEdge:** Enhances service assurance for latency- and jitter-sensitive applications such as audio and video conferencing. By intelligently identifying applications at the edge and prioritizing traffic through fast processing, it ensures smooth user experiences. Cloud-based visualization further enables 24/7 monitoring of latency, jitter, and packet loss at campus, AP, and individual user levels.

**iHeal:** Creates a self-healing and predictive wireless environment. By combining terminal feedback with cloud analytics, it not only identifies and resolves issues automatically but also continuously adapts to dynamic air-interface conditions, user scales, and service requirements, progressing toward "zero-touch" operations.

#### Conclusion

Through the Wireless Iron Triangle concept, encompassing high-quality WLAN products, scenario-optimized deployment, and intelligent operations that enterprise IT managers can effectively overcome wireless campus network challenges during digital transformation. This framework enhances user experience, supports emerging services, and drives enterprises toward a more adaptive and resilient digital infrastructure.



# Data Express: High-Efficiency Transmission of Massive Data Across Wide Area Networks

■ By Qin Juan and Gao Feng, H3C Technology Strategy Department

#### **Abstract**

In the digital and intelligent era, the efficient and secure cross-wide area network (WAN) transmission of massive data has become an urgent need across industries. However, traditional TCP protocols suffer significant performance degradation in high-latency, packet-loss WAN environments. This document proposes an efficient Data Express solution, which combines the QUIC protocol and WAN RDMA technology. Through protocol optimization and software-hardware collaboration, it significantly improves WAN data transmission throughput and reliability. The Data Express system supports multi-task management, breakpoint resumption, and file verification, making it widely applicable to scenarios such as scientific data exchange, media content distribution, and intelligent computing across WANs. It provides industries with an efficient, stable, and reliable data transmission solution, thereby supporting digital transformation and business innovation.

#### **Keywords**

Data Express; QUIC; WAN RDMA



#### Introduction

In the digital and intelligent era, data has become the core driver of enterprise development. Whether it's massive scientific data or high-definition media content distribution, efficient and secure WAN transmission of massive data is now an urgent need across industries. Particularly in intelligent computing WAN scenarios, cross-data center collaborative training, edge transmission and training of sensitive samples, and model distribution impose higher demands on online data migration efficiency. Traditional TCP protocols have revealed performance bottlenecks in WAN data transmission. This document addresses these bottlenecks by proposing an efficient Data Express solution and will detail its core features, software architecture, and technical implementation principles. It also analyzes the application value of Data Express in various scenarios and explores its future development trends, opening new paths to improve cross-WAN data transmission efficiency.

## 1. Traditional WAN Data Transmission Solutions

In WAN data transmission scenarios, users need to transfer large volumes of data over long distances while maintaining high transmission rates. Traditional WAN transmission software like FTP and SFTP rely on TCP as the transport protocol. In high-latency or packet-loss scenarios, performance degrades due to the need to wait for RTT acknowledgment. Especially in complex, high-latency, unstable, and high bit error rate (BER) WAN environments, TCP's performance bottlenecks become more pronounced, resulting in effective throughput far below the available bandwidth and inefficient network utilization.

To ensure user service experience, end-to-end optimization at the protocol level is required, including:

◆ **Protocol optimization:** Improve existing TCP protocols or introduce new transport protocols to enhance data transmission efficiency and reliability in WAN environments.

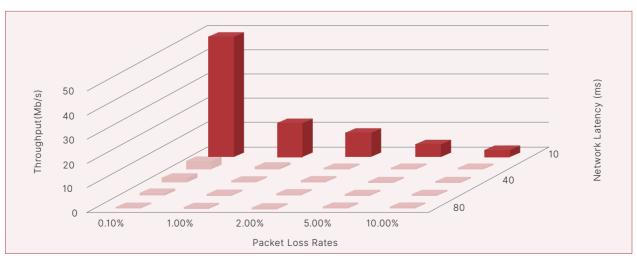


Figure 1. TCP Performance on a 50Mbps Link Under Different Latency and Packet Loss Rates



◆ End-to-end performance: Optimize end-to-end data transmission performance by reducing protocol overhead, improving congestion control algorithms, and enhancing packet transmission reliability and integrity verification mechanisms.

Through innovation and optimization at the protocol level, the issues of throughput degradation and insufficient performance in traditional WAN data transmission can be effectively addressed, meeting the higher demands of the digital and intelligent era.

#### 2. Data Express Transmission System

The Data Express transmission system is designed to address the need for large-scale data transmission in WAN environments, providing efficient, secure, and reliable data transfer services. The system features modularity and high scalability, combining multiple technologies and protocols to achieve end-to-end data transmission optimization. The transmission software running at the application layer controls the overall data transfer process. By leveraging underlying OFED (OpenFabrics Enterprise Distribution) and QUIC interfaces, it enables high-efficiency data transmission.

#### 2.1 Key Features of Data Express

**Transport Protocol:** Based on high-efficiency transport layer protocols, it improves performance in big data flow scenarios by optimizing congestion control algorithms, flow control, and ACK processing.

**File Transfer Control:** Supports multi-task transmission, with each task enabling multi-connection file chunking and transfer control algorithms. It utilizes multithreading to maximize multi-core resources, boosting file transfer speeds.

**Write Optimization:** Enhances file system write performance and accelerates disk storage through optimized file write control.

**Multi-Mode Support:** Features a client-server architecture, providing both a user-friendly client UI and a command-line interface.

#### 2.2 Data Express Software Architecture

The Data Express system software consists of four main layers: application, interface, business, and protocol.

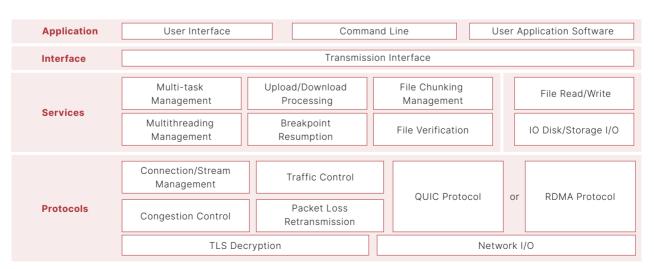


Figure 2. The Software Architecture of Data Express System



#### **Application layer**

It offers a user-friendly UI and command line, supporting both graphical and command-line system operations and control. It also provides SDKs for integration with existing user systems.

#### Interface layer

It serves as the entry point for data transmission, providing interfaces for external systems or applications to receive and initiate data transfer requests.

#### **Business Layer**

- Multi-Task Management: Supports concurrent processing of multiple data transmission tasks, improving system parallel processing capability and efficiency.
- Multithreading Management: Uses multithreading to enable concurrent processing of data transmission tasks, enhancing speed and performance.
- Upload/Download Processing: Handles large file upload and download requests, ensuring data integrity and reliability during transmission.
- Breakpoint Resumption: Allows resuming interrupted file transfers from the last breakpoint, avoiding redundant operations and time waste.
- File Verification: Ensures data integrity and prevents malicious modifications through file checksum validation.
- File Chunking Management: Splits large files into smaller chunks for transmission, reducing the impact of single-file failures and improving efficiency.
- File Read/Write: Manages file read/write operations and interacts with the file system.

• Disk/Storage I/O: Manages input/output operations for drives or other memory devices, supporting high-efficiency data read/write.

#### **Protocol Layer**

Connection/Stream Management: Manages network connection establishment, maintenance, and termination, controlling data flow transmission and reception (Rx).

- Traffic Control: Monitors and controls transmission flow to prevent network congestion and ensure stability.
- Packet Loss Retransmission: Retransmits lost packets during transmission to ensure data integrity and reliability.
- Congestion Control: Uses advanced congestion control algorithms to dynamically adjust transmission rates and avoid network congestion.
- TLS encryption/decipher: Provides secure encryption and decipher at the transport layer, ensuring the confidentiality and security of data during transmission.
- QUIC transmission: Selects the QUIC protocol for data transmission based on scenario requirements, suitable for real-time or large-volume data transfer needs.
- RMDA: Selects the RDMA protocol for data transmission based on scenario requirements, suitable for real-time or large-volume data transfer needs implemented via hardware (HW).
- Network IO: Interacts with network hardware (HW) and drivers, handling underlying network input/output operations.

#### 2.3 Based on the QUIC Protocol

QUIC, short for Quick UDP Internet Connections, is a UDP-based transport layer protocol. QUIC packets consist of two parts: header and data.

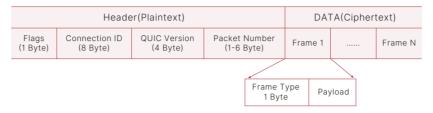


Figure 3. QUIC Packet Format

#### **Details**

- Header: Plaintext, containing 4 fields: Flags, Connection ID, QUIC Version, and Packet Number.
- Data: Ciphertext, may contain one or more frames, each divided into type and payload.

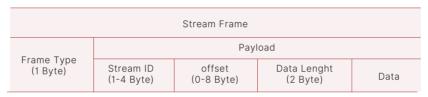


Figure 4. Frame Format (FF)

- FrameType: Frame types include Stream, ACK, Padding, Window Update, Blocked, etc.
- Stream ID: Identifies (ID) the stream to which the packet belongs.
- Offset: Indicates the byte-level offset of the data within the Stream, used for data ordering.
- Data Length: Indicates the length of the actual application data.
- Data: The actual application data.

The QUIC-based data delivery solution does not rely on hardware (HW) NICs. The QUIC protocol incorporates multiple optimizations, effectively addressing high latency and packet loss in wide area network (WAN) environments, significantly improving data transmission efficiency and reliability. This solution leverages QUIC's advantages, such as fast connection establishment, multiplexing, traffic control, and built-in encryption, while employing advanced congestion control to maintain superior performance under complex network conditions.

**Fast connection establishment:** QUIC uses UDP as the underlying transport protocol, avoiding TCP's three-way handshake. The first connection completes key negotiation and data transmission within 1-RTT. QUIC also caches the current session's context. When the client reconnects, it can directly use cached data to compute the communication secret-key encryption and send application data, achieving 0-RTT handshake transmission.

**Multiplexing:** QUIC supports multiplexing, allowing multiple data flows to transmit in parallel over a single connection, avoiding TCP's head-of-line blocking.

In multipath scenarios, relying solely on strictly increasing Packet Numbers cannot guarantee data order and reliability. QUIC introduces the Stream Offset concept, using the Stream's Offset to ensure application data order.

#### **Example:**

Assume Stream X's data is sent in three frames with Offsets of 0, 100, and 200:

Path A: Sends [Packet 1: Stream X, Offset=0, Data=0~99 bytes]
Path B: Sends [Packet 2: Stream X, Offset=100, Data=100~199 bytes]

Path A: Sends [Packet 3: Stream X, Offset=200, Data=200~299 bytes]

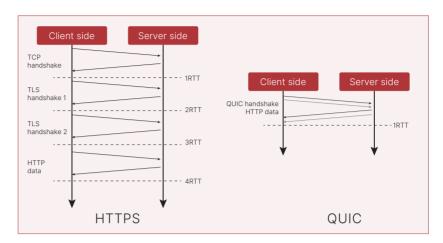


Figure 5. Process of TCP and QUIC Connection

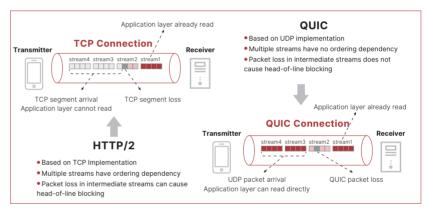


Figure 6. TCP vs. QUIC Multiplexing Transmission



Figure 7. Ensuring Data Order via Offset

#### Receiver behavior:

- 1. Upon receiving Packet 2 (Offset=100), it detects that Offset=0 data has not arrived and buffers Packet 2.
- 2. After receiving Packet 1 (Offset=0), it immediately delivers bytes 0~99 to the application layer and checks the buffer:
- Discovers buffered Offset=100 data and continues delivering bytes 100~199.
- 3. Finally, upon receiving Packet 3 (Offset=200), it directly delivers bytes 200~299.



**Built-in encryption:** QUIC integrates TLS encryption at the protocol layer, ensuring data security while reducing handshake latency.

**Traffic control:** Since QUIC supports multiplexing, it enables two levels of traffic control: connection-level and stream-level. QUIC implements traffic control via WINDOW\_UPDATE and BLOCKED frames.

The receiver sends a WINDOW\_UPDATE frame to notify the sender of an updated receive (Rx) window, allowing more data transmission. The sender may also send a BLOCKED frame to indicate data transmission is constrained by the receiver's window.

Through its UDP-based design, built-in encryption, multiplexing, and traffic control mechanisms, combined with advanced congestion control algorithms, the QUIC transport layer protocol significantly improves the efficiency, reliability, and security of large-scale data transmission over wide area networks (WANs). Compared to traditional TCP-based software, it offers clear advantages in latency, throughput, and reliability.

## 2.4 RDMA-Based Software-Hardware (HW) Co-Optimization

To further optimize wide area network (WAN) large-data transmission performance, the industry is exploring RDMA (Remote Direct Memory Access)-based software-hardware (HW) co-optimization solutions. Compared to QUIC, RDMA requires additional hardware (HW) NIC support, offloading data processing to hardware (HW) to free CPU

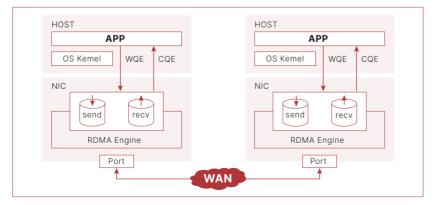


Figure 8. Wide-Area RDMA Architecture

resources while leveraging hardware (HW)'s high-throughput capabilities to significantly enhance data delivery performance.

Using RDMA technology in wide area network (WAN) transmission, especially in high-latency (over 100ms) and packet-loss (0.1%) environments, presents significant challenges. Originally designed for high-performance, low-latency local area network (LAN) or data center environments, RDMA faces issues like high latency, packet loss handling, congestion control, and network jitter in WANs, which differ from its intended design. Specific optimizations are required to enable RDMA to deliver equally efficient data processing performance in WANs as in data centers. These include:

Connection optimization: High latency in WANs (over 100ms) significantly impacts RDMA's connection establishment and memory registration efficiency. Traditional RDMA connection mechanisms often require multiple round-trip times (RTTs), leading to inefficiency in high-latency environments.



Reducing the RTT count for connection establishment and optimizing memory registration effectively lowers connect time and improves initial transmission efficiency. In some sessions, connect time can be reduced by over 50%.

#### **Optimization measures:**

**Reducing handshake count:** By optimizing and sharing connection establishment methods, the connection process is divided into initialization and parallel data transmission, reducing total connect time.

**Fast memory registration:** Optimizing the memory registration process reduces the number of registrations and mappings. Pre-registering and caching frequently used memory regions minimize repeated registration overhead. During initial connection, memory region registration information is stored in caching for quick access in subsequent connections.

**Retransmission mechanism optimization:** Traditional Go-Back-N retransmission is inefficient in high packet-loss environments, as losing one packet requires retransmitting all subsequent packets. High packet loss leads to excessive retransmissions, wasting bandwidth and increasing latency. Selective retransmission reduces unnecessary packet retransmissions, lowering bandwidth consumption and significantly improving transmission efficiency in high-loss environments.

#### **Optimization measures:**

**Selective retransmission:** Introducing selective retransmission to replace Go-Back-N. Upon detecting packet loss, only the lost packets are retransmitted, reducing bandwidth waste. Combined with ACK (acknowledge) and NACK (repudiation) messages, the receiver notifies the sender of specific lost packets, ensuring only lost packets are retransmitted.

**Out-of-order reassembly:** Due to network complexity and varying transmission delays, packets may arrive out of order. During data transmission, each packet is assigned a unique sequence number, and a sequence bitmap is maintained to reorder packets at the receiver. In high-latency, high-throughput WAN scenarios, optimal out-of-order reassembly solutions include efficient sequence numbering, buffer management, and timeout/retransmission mechanisms.

Congestion control algorithm optimization: In WANs, congestion control is critical due to long delays, such as feedback lag caused by high latency. High WAN latency prevents congestion control algorithms from timely network state updates, affecting decision accuracy and timeliness. Fairness issues arise when multiple concurrent RDMA flows compete for bandwidth, leading to unfair allocation. Using proven advanced WAN congestion control algorithms combined with RDMA ensures fair traffic control, offering fast convergence and high throughput.

#### **FPGA-based high-performance network**

coprocessor: Leveraging FPGA's reprogrammable hardware (HW), FPGA-based NICs can offload and configure various data plane (DP) functions as needed. These reprogrammable DP functions enable real-time adjustments, providing highly flexible and efficient processing in high-speed WANs. All offloaded functions run at hardware speed, significantly improving data transmission efficiency and performance.

## 3. Data Express Application Scenarios and Outlook

As an efficient WAN data transmission system, Data Express not only addresses traditional issues like slow speed and low security but also demonstrates broad application value across industries.

In scientific research data exchange, it enables fast, safety ground transmission of experimental data and model files, improving exchange efficiency and fostering cross-institution, cross-region

collaboration to advance the pace of research and expedite the commercialization of results. In media content distribution, it rapidly transfers large files like HD videos and audio, enhancing production and distribution efficiency to meet the industry's demands for timeliness and stability. In intelligent computing WAN scenarios, it provides efficient data movement, significantly improving transmission efficiency and reducing latency.

Whether for massive data exchange in research institutions or efficient data flow in intelligent computing WANs, Data Express delivers high-performance, stable, and reliable transmission solutions, meeting the higher demands of the digitize era.

Looking ahead, with deepening digitization and explosive data growth, Data Express will continue innovating in technology, security, and applications, providing industries with more efficient, secure, and intelligent WAN transmission solutions, driving digital transformation and business innovation.

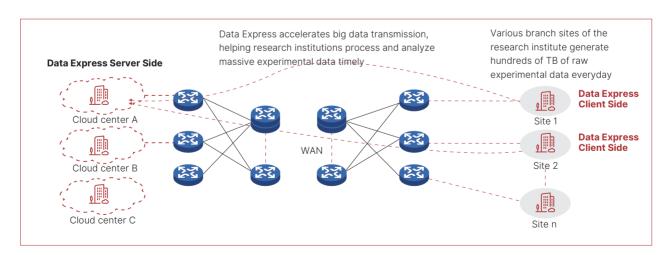
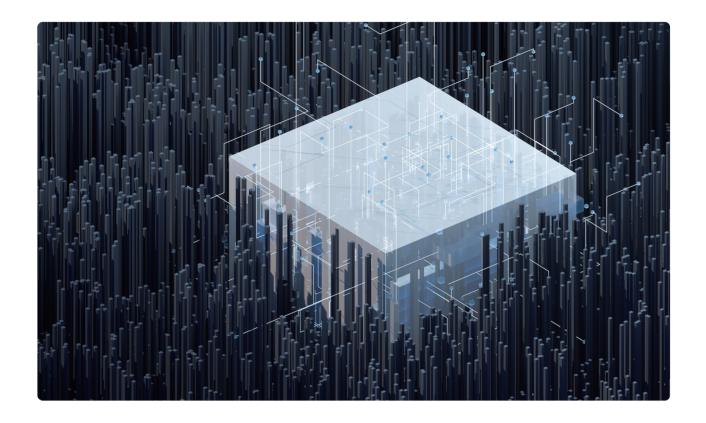


Figure 9. Example of Scientific Research Big Data Transmission Scenario





## Reflections on the Evolution of Computing Power Network Architectures

By Meng Dan, H3C Solutions Department

#### **Abstrac**

Computing power network deeply integrates previously fragmented cloud and network resources, representing a strategic choice for carriers in the continuous evolution of network architecture. This document introduces the development of cloud computing technology to derive the inevitability of computing power networks, compares centralized and distributed architectures, elaborates on key technologies for building such networks, and concludes with prospects for their future development.

#### **Keywords**

Computing power network; Centralized architecture; Distributed architecture; Computing-network brain; Computing power routing; Computing-network metrics

## 1. Background of Computing Power Network

In the field of computing power networks, standards preceded practice. The first technical document was initiated in CCSA in August 2019. After years of continuous technical discussions, industry consensus has formed, and computing power networks have become one of the strategic emerging industries in national research.

Computing power networks emerged from architectural shifts in cloud computing. Cloud computing offers enterprises and individuals a flexible, manageable, and cost-effective IT solution. Initially appearing in centralized form, thousands of servers were concentrated in a single geographic area, providing various laaS, PaaS, and SaaS services through highly multiplexed and shared pools. However, centralized cloud pools are not

suitable for all scenarios, such as users requiring low latency or high bandwidth, leading to the evolution of cloud computing into a multi-tiered centralized-edge collaborative architecture.

Different tiers provide varying quality of service (QoS): centralized clouds deliver high-throughput computing power, while edge clouds offer low-latency, high-bandwidth, lower-computing services. Computing power network technology provides an optimal solution for balancing computing capabilities and network metrics from access points to clouds.

Before the concept of computing power networks, cloud-network integration was proposed, focusing primarily on the cloud with networks merely providing reachability. Computing power networks, however, treat computing and networks as equally important.

Computing power networks are essentially an ecosystem encompassing users, applications, computing resources, and networks, interconnected through a scheduling platform. To be effective, this platform requires a rich application ecosystem, clear computing-network metrics, and dynamic scheduler

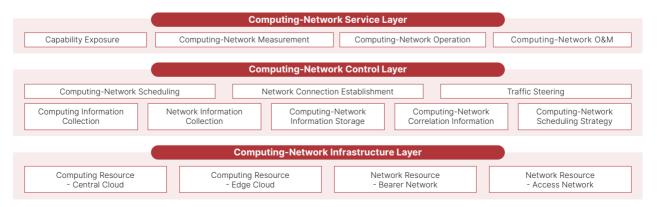


Figure 1. Computing Power Network Architecture



policies (SP). Without these, computing power networks remain impractical.

The architecture can be divided into three layers. The underlying layer consists of computing and network infrastructure. Computing initially supports proprietary systems before expanding to third-party resources. Networks extend from backbone/MAN to incorporate various access networks.

The middle layer is the control layer, bridging infrastructure and services by storing computing and network information and making integrated decisions to select appropriate pools and establish network paths.

The top layer is the service layer, providing open capabilities to various computing-network applications.

#### 2. Computing Power Network Architectures

Computing power networks include centralized and distributed architectures.

In centralized architectures, the scheduling platform acts as the hub, collecting computing usage and network topology/QoS data. For application requests, it centrally selects suitable pools and establishes network paths to direct traffic.

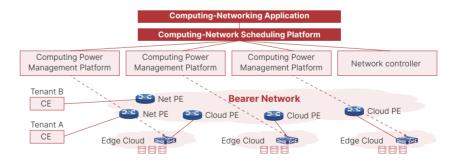


Figure 2. Centralized Computing Power Network Architecture

In distributed architectures, the scheduling platform's role is reduced, focusing more on operation and analysis. Edge nodes monitor local pool usage and advertise this information via routing protocols. Customer site



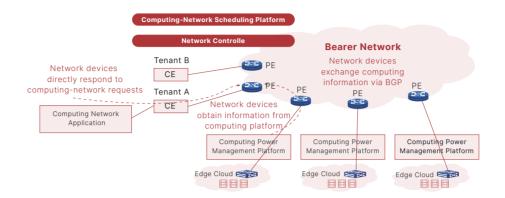


Figure 3. Distributed Computing Power Network Architecture

equipment can then make local decisions to select pools and establish paths.

The additional function on network equipment in distributed architectures is termed computing power routing, where devices integrate computing factors beyond mere traffic forwarding.

Centralized solutions are simpler, requiring no additional network device capabilities, but the scheduling platform becomes a performance bottleneck at scale, slowing response times.

Distributed solutions enable faster, real-time responses by embedding computing awareness into networks, though this increases technical complexity for network equipment.

Currently, centralized architectures are more mature, with carriers developing proprietary products.

Distributed architectures remain exploratory, with progress in standards and prototypes.

## 3. Key Technologies of Computing Power Networks

Key technologies include computing-network metrics (the foundational "language" between supply

and demand), awareness (collecting computing/network data for decision-making), and scheduling (fulfilling application requests).

Computing-Network Metrics: As a platform connecting supply and demand, standardized metrics are essential. Network factors include bandwidth, latency, jitter, packet loss, and reliability. Computing metrics vary by service tier: laaS uses CPU cores, memory/storage usable capacity; PaaS employs class type (CT)-specific metrics like QPS/TPS for databases or messages processed per second for queues; SaaS metrics include audio/video decoding capacity, session counts, or busy states.

Both architectures use identical metrics.

**Computing-Network Awareness:** The network must perceive and record computing/network information for decision-making.

In centralized architectures, the platform acquires computing data via subscription/polling and network topology from the domain. Large-scale deployments use elastic expansion or hierarchical platforms for flexibility.

In distributed architectures, network devices interfacing with pools gather computing data and flood it via routing protocols, extending traditional topology/reachability information with computing metrics.

Computing-Network Requests: Computing-network applications with demands for both computing power and networks initiate computing-network requests. Not all applications require computing-power networks. Applications insensitive to latency and bandwidth can simply follow the traditional model and be built solely on cloud computing, where the network merely serves as reachability support. When expressing computing-network requirements, computing-network applications should adhere to computing-network metrics and align with the supply side.

Under a centralized architecture, computing-network requests are initiated by computing-network applications to the computing-network scheduling platform, typically interacting via RestAPI interfaces. These interfaces specify the computing and network information of concern to the application. Fundamentally, only the application itself understands its demands for computing power and networks. Therefore, computing-power networks require deep involvement from applications, which entails certain modification efforts.

Under a distributed architecture, computing-network requests are directly embedded in the application's traffic. There are various ways to carry computing-network requests. If the application can be modified, it can convey computing and network demands via extension headers, though such

modifications require operating system support and pose technical challenges. Alternatively, predefined computing-network service templates can be offered externally with different service IDs, allowing applications to express their computing-network requests by carrying different service IDs.

Computing-network scheduling: Upon receiving a computing-network request from an application, the computing-power network combines the computing and network information collected during the awareness phase with dynamically configured scheduler policies (SP) to select the most suitable cloud resource pool.

Under a centralized architecture, the computing-network scheduling platform resolves the computing-network request, makes decisions, selects the most suitable cloud resource pool, establishes the transport path between the application's access point (AP) and the cloud resource pool, and directs the application's traffic into the transport network using various traffic steering techniques.

Under a distributed architecture, after an application initiates a computing-network request, the edge node at the customer site resolves the computing-network request carried in the application's traffic. Internally, it first maps the service ID to the detailed requirements of the computing-network request, then makes decisions, selects the most suitable cloud resource pool, establishes the transport path between the application's access point (AP) and the cloud resource pool, and directs the application's traffic into the transport network using various traffic steering techniques.

#### 4. Prospects of Computing-Power Networks

Current research on computing-power networks is expanding into deeper and broader directions.

On the computing side, from centralized cloud computing to cloud-edge collaborative edge computing, cloud resource pools at different hierarchies jointly provide computing-network services. In the future, computing power may extend further to the device side, where various ubiquitous computing resources can supplement existing cloud-edge resources. Computing-power networks can correspondingly extend downward, but the trustworthiness, measurement, and authentication of device-side computing power require ongoing analysis to avoid introducing security risks. Additionally, while consensus on computing-power metrics is easier to achieve at the laaS layer, the diversity of service types and characteristics at the PaaS and SaaS layers means metrics need to be jointly defined by the computing-network scheduling platform and applications, which will take a longer cycle.

On the network side, current development focuses on networks above the metro aggregation core. In the future, computing-power networks will extend further into the access network. While publishing computing-power information via routing protocols is a basic consensus, achieving interoperability between multi-vendor network equipment requires carriers to lead the definition of low-level (LL) protocol message formats, which will need thorough industry discussion.

On the ecosystem side, computing-power networks serve as a support platform for applications. Without applications, computing-power networks lack vitality. However, motivating applications to customize development based on computing-power networks requires sufficient incentives. Ensuring all parties continue to benefit from the evolution of computing-power networks is a direction requiring long-term promotion and focused attention.

The IETF, as the most authoritative standards organization in the field of IP networking, established the CATS Working Group in March 2023, which focuses on research related to computing-aware traffic steering. H3C actively participates in the drafting and advancement of these standards and aims to contribute more significantly in the standards domain.









#### Follow H3C for More Digital and Al Solutions

https://www.h3c.com/en/

Copyright @ New H3C Technologies Co., Ltd. All Rights Reserved.

#### **Disclaimer**

Although H3C attempts to provide accurate information in this document, it does not guarantee that the content is free of technical inaccuracies or typographical errors. H3C reserves the right to make changes to the content of this document without notice or notification to improve product quality.

This document contains, but is not limited to, predictive product information and new technologies, as well as information changes brought about by new application updates. Uncertainties exist in actual product applications, which may lead to slight differences between actual results and predictive information. The content of this document is for reference only and does not constitute any binding commitment. If there is a complete disclosure of the above information, H3C will notify separately.

The images in this document are for reference only. Please refer to the actual products.