

Inspur InCloud Rail HCI All-in-One Based on Intel Technologies

"The next generation of hyperconvergence goes far beyond the integration of storage architecture. It now involves private cloud solution, heterogeneous computing power, cloud-native, and cloud-data-intelligence integration. Through our partnership with Intel and the application of Intel® SPDK in the InCloud Rail solution, we have effectively eliminated performance bottlenecks such as distributed storage in kernel context switching and CPU interruptions. The optimized solution takes full advantage of the performance potential of storage media, meets the performance requirements of industries such as finance and healthcare in a variety of application scenarios, and accelerates the intelligent transformation of companies."

– **Zhang Dong**
Chairman

Inspur Data Technology Co., Ltd.

"Hyperconvergence represents the integration of storage, computing, and networking into a single unified system. Offering simplicity, flexibility, and scalability, hyperconvergence has gained recognition for its powerful solution advantages. Through our partnership with Inspur to drive the optimization of HCI all-in-one, Intel continues to meet the needs of customers and accelerates the smart transformation of the market."

– **Carl Li**

Managing Director, China Enterprise and Government, Global OEM
Intel Cloud and Enterprise Solution Group

Overview

As a form of infrastructure that integrates software and hardware, hyperconvergence offers advantages in terms of simplified deployment and easy operation and maintenance. Hyperconvergence has seen wide use across a variety of application scenarios and industries. Focusing on cloud-native, 5G, big data, cloud-edge-end, and other application scenarios, Inspur has continued to drive the innovation and evolution of the InCloud Rail solution over recent years. The solution supports cross-cloud platform convergence and centralized management of heterogeneous virtual environments through the InCloud DataCenter cloud management platform. The solution also comes equipped with a rapid basic environment deployment tool that is much faster than those of traditional solutions.

When deployed in finance, healthcare, and other industries, HCI all-in-one for database and other scenarios face huge challenges in terms of storage performance. Users need to maximize data throughput and reduce latency to support the efficient operation of key services. To improve the performance of storage subsystems, Inspur used Intel® Storage Performance Development Kit (SPDK), with acceleration through the NVMe over Fabrics (NVMe-oF) protocol. As an alternative to the iSCSI protocol, the NVMe-oF protocol allows the host to access distributed storage via the native NVMe protocol. This allows for low-latency, high-throughput block storage device and solves the issues of kernel context switching and CPU interruptions that often occur when kernel-mode drivers attempt to read and write NVMe disks. The optimized solution provides high-performance distributed storage for virtual machines, reduces overall delay and latency, and meets the needs of finance and healthcare customers in scenarios such as databases.

Challenge: The Finance and Healthcare Industries Face Storage Performance Bottlenecks

Due to the large scale of users and the high value of businesses and data, finance and healthcare institutions have always been forerunners in informatization and technical applications,

with business development highly dependent on IT systems. As finance and healthcare institutions have started the process of digital transformation over recent years, the optimization and reconstruction of IT infrastructure along with the development of a flexible and reliable IT service platform for upper-layer applications and business innovation have become priorities of digital strategies.

As a result, more and more finance and healthcare institutions are trying to embrace hyperconvergence to reduce costs by reconfiguring hardware deployment through all-in-one for software-defined, flexible, and efficient IT infrastructure. This allows for reduce pressure on the maintenance and expansion of IT infrastructure, greater business resiliency, and unlocks more resources for operational growth.

In embracing hyperconvergence, finance and healthcare institutions are particularly concerned with its ability to provide powerful storage capabilities. In tandem with user growth and business innovation, scenario-based, mobile, and fragmented data is growing rapidly in the finance and healthcare industries, placing enormous pressure on storage systems. Using financial databases as an example, Internet transactions, data risk control, real-time marketing, and other applications require highly responsive databases. In high concurrency operations, storage systems can easily become performance bottlenecks that affect overall responsiveness.

One of the primary causes of storage system performance bottlenecks is storage media. The past few years have seen rapid developments in storage media. In modern day, the IOPS of NVMe SSDs is far higher than that of HDDs. As latency has been reduced from milliseconds to microseconds, system performance bottlenecks have gradually transferred from storage hardware to networks and processors. Methods such as traditional file systems and schedulers have become new bottlenecks in storage systems as they cannot fully

utilize the performance capabilities of modern storage media. These bottlenecks include:

- When performing read/write operations on NVMe disks with regular NVMe kernel drivers, issues such as kernel context switches and CPU interrupts will occur. In high-performance all-flash array storage, interrupts result in latency uncertainties, which in turn lead to significant latency and performance overheads.
- In the traditional I/O model, applications sleep after submitting read and write requests. After I/O is completed, the interrupt will wake, meaning that the interrupt overhead plays a key role in overall I/O time.

Solution: Inspur InCloud Rail HCI All-in-One Based on Intel Technologies

Representing the next generation of HCI solutions, the Inspur InCloud Rail HCI all-in-one realizes server resource pooling through software-defined compute, storage, and network technologies. This provides the entire IT environment with higher availability, security, and scalability when compared with physical hardware alone. The InCloud Rail solution can effectively meet enterprise demands in terms of cost reduction, simplified management, and improved security and stability, facilitating the migration of core businesses to the cloud and the construction of enterprise cloud data centers.

The Inspur InCloud Rail solution realizes the pooling and unified management of storage resources, while supporting the integration of heterogeneous computing power through SmartONE software-defined storage and all-flash array architecture. SmartONE uses etcd components for cluster management, covering the relationship between distributed storage nodes, metadata transfer and




Figure 1. InCloud Rail HCI All-in-One




Figure 2. Inspur SmartONE software architecture

consistency between nodes, and other aspects. In terms of resource management, SmartONE provides service portals for qemu, iSCSI, and NVMe-oF, connects with VDI volume management portals, and provides external storage resources. In terms of object storage, SmartONE primarily receives I/O requests from VDI, dispatches the requests to relevant nodes in accordance with the data distribution algorithm, then calls Gateway to distribute the requests to distributed storage nodes in consideration of the data replication strategy and EC policies.

In order to meet the demanding requirements for storage performance in scenarios such as databases, real-time read/write, and random access to ultra-large-scale datasets, the InCloud Rail solution has adopted Intel® Xeon® Scalable processors and the Intel® Ethernet Network Adapter E810 to resolve performance bottlenecks at the storage engine level.

- **Intel Xeon Scalable Processors:** Designed for modern data centers, these processors improve the operating efficiency of a variety of infrastructure, enterprise applications, and technical computing applications to improve total cost of ownership (TCO) and increase use productivity. With enhanced single-thread performance, the processor offers superior performance and scalability for compute-intensive workloads in compute, storage, and networking applications.
- **Intel Ethernet Network Adapter E810:** Features 100/25 GbE single or dual port connection, high PCIe 4.0 x16 performance,

and support for various advanced functions such as Application Device Queues (ADQ), Dynamic Device Personalization (DDP), RDMA iWARP, and RoCEv2 to improve workload predictability, throughput and to lower application latency.

Inspur Uses Intel SPDK to Optimize Storage Performance

The Intel SPDK provides a suite of tools, libraries, and solutions for writing performant and scalable user mode storage applications. The Intel SPDK realizes high performance and scalability through techniques such as moving drivers to user space, avoiding system calls, and allowing zero-copy access from applications. The Intel SPDK provides a high-performance application framework through non-blocking, message passing, and asynchronous programming, along with a unified user-mode general block device that can efficiently manage a variety of backend storage devices.

After using the Intel SPDK, the user mode driver is completed by polling the hardware instead of relying on interrupts. In addition to effectively reducing the total delay and delay difference, this also provides clear performance advantages in IOPS per CPU core when compared with kernel drivers. The Intel SPDK also features a non-blocking and performant I/O path mode that avoids all blocks in key I/O paths by relying on message passing to share resources across multiple threads to increase parallelism.

Inspur has partnered with Intel to use the Intel SPDK to optimize modules such as the standalone storage engine and NVMe-oF storage service of SmartONE distributed storage.




Figure 3. Architecture of the Intel SPDK

Standalone Storage Engine

Inspur SmartONE distributed storage implements a performant foundation based on Intel SPDK NVMe drivers, supports direct interaction with NVMe disk devices, and adopts a non-blocking design to realize the parallel processing of I/O commands. At the upper layer, NVMe implements a dedicated standalone storage engine to effectively avoid dual writes that are common in traditional file systems through a memory-based metadata and log management system.




Figure 4. NVMe dedicated standalone storage engine

The next generation of standalone storage engines based on Intel technologies deliver significant improvements in terms of performance. Test data shows that the standalone storage engine of Inspur SmartONE distributed storage has almost reached the maximum performance supported by NVMe hard disks.

NVMe-oF Distributed Storage Service

The NVMe-oF distributed storage service provides two block service solutions for HCI platforms. The first uses the vhost-user protocol for shorter virtual machine I/O paths. The second operates as a storage service that provides NVMe-oF TCP/RDMA block services for servers. SmartONE uses SPDK vhost-user technology to eliminate the need for a guest virtual machine to accessing NVMe devices through PCIE, avoiding kernel updates of PCI configuration space. QEMU virtual IO is directly captured in user mode, with data transferred to the storage system in a zero-copy manner.

SmartONE supports block storage services through the NVMe-oF storage protocol, providing external access through both TCP and RDMA. As an alternative to the iSCSI protocol, the NVMe-oF storage protocol allows hosts to access distributed storage through the native NVMe protocol, providing low-latency high-throughput block storage devices.




Figure 5. SmartONE NVMe-oF Distributed Storage Service

In order to verify the performance of the SmartONE NVMe-oF distributed storage service, Inspur compared the performance of the iSCSI storage protocol on Inspur M5 series servers¹ and Inspur M6 series servers², along with the performance of the NVMe-oF storage protocol on Inspur M6 series servers (NVMe clone mode, 3 clones).




Figure 6. Comparison of 4K read/write performance between iSCSI and NVMe-oF storage protocols on M5/M6 platforms³




Figure 7. Comparison of 1MB read/write bandwidth between iSCSI and NVMe-oF storage protocols on M5/M6 platforms⁴

Test data shows that the 4K read/write performance and 1MB read/write bandwidth of the NVMe-oF storage protocol have been significantly improved to meet the demanding storage performance requirements of critical healthcare and finance applications.

Result: Eliminated Performance Bottlenecks and Realization of High-Performance Storage

The application of Intel software and hardware solutions to the Inspur SmartONE distributed storage system unlocks the performance potential of NVMe SSDs and supports database application scenarios with high IOPS and latency requirements.

- Significantly improves storage system performance and reduces data latency. It is capable of meeting the requirements of demanding database application scenarios while effectively controlling the TCO of distributed storage systems.
- Takes advantage of the performance advantages of Intel hardware and benefits from innovations in next-gen Intel hardware.
- Integrated into a HCI all-in-one, with the ability to significantly simplify the complexity of deployment, management, and O&M through a pre-built virtualization platform, cloud management platform, and rapid deployment tools and related workflows.

Outlook

Focusing on cloud, data, and intelligence application environments, Inspur will continue to drive innovations in HCI systems (including storage). Utilizing next-gen Intel Xeon Scalable processors, Intel® Optane™ persistent memory, Intel® Ethernet adapters, and other hardware with leading software solutions, Inspur will improve core competitiveness in performance, reliability, simple O&M, and scalability to meet the needs of all industries and scenarios.

Inspur also plans to utilize Intel Optane persistent memory as the storage medium for the cache layer of its standalone storage engine. Intel Optane persistent memory offers solutions that provide high speed, cost-effectiveness, high capacity, persistent data protection, advanced encryption, and more. When used in conjunction with the Intel PMDK, Intel Optane persistent memory allows applications to directly access persistent memory without going through the file system page cache system, system calls, and drivers to reduce the overhead of I/O processes and significantly reducing data latency.

Looking forward, Inspur will continue to stand by its user-oriented principles and closely integrate cutting-edge development trends and market demands. Inspur will also further innovate, lead, and iterate to develop HCI all-in-one that integrated simplicity, extraordinary performance, robustness and reliability, and an open ecosystem to contribute to cloud, digital, and business transformation.

¹ Testing on M5 series servers was based on 2nd Gen Intel Xeon Scalable processors. For more details, refer to <https://www.inspur.com/lcjtww/2509917/2509945/2510106/index.html>. Intel does not control or audit third-party data. Please review the content, consult other sources, and independently confirm if the data provided is accurate.

² Testing on M6 series servers was based on 3rd Gen Intel Xeon Scalable processors. For more details, refer to <https://www.inspur.com/lcjtww/2509917/m6/index.html>. Intel does not control or audit third-party data. Please review the content, consult other sources, and independently confirm if the data provided is accurate.

^{3,4} Data quoted from tests conducted by Inspur in June 2022. Intel does not control or audit third-party data. Please review the content, consult other sources, and independently confirm if the data provided is accurate.

About Inspur

Inspur Group is a leading Chinese cloud computing and big data service provider that consists of three companies, namely Inspur Information, Inspur Software, and Inspur International. Its primary business includes cloud computing, big data, IIoT, next-gen communications, and a number of application scenarios. The company provides IT products and services to over 120 countries and regions around the world. Inspur was one of the first IT brands established in China. Dedicated to innovation, the company has made several key contributions to the development of the Chinese IT industry. Through the provision of leading technologies and enhancing competitiveness, the company has become a leader in next-gen information technologies, comprehensively serving the digital transformation and high-quality development of society and the economy.

About Intel

Intel (Nasdaq: INTC) is an industry leader, creating world-changing technology that enables global progress and enriches lives. Inspired by Moore's Law, we continuously work to advance the design and manufacturing of semiconductors to help address our customers' greatest challenges. By embedding intelligence in the cloud, network, edge and every kind of computing device, we unleash the potential of data to transform business and society for the better. To learn more about Intel's innovations, go to newsroom.intel.com and intel.com.



Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

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.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.