



Intel® Select Solutions for GBase 8a MPP Cluster*



Data driven digital transformation are ushering in a new era, driving innovative approaches for finance, telecommunications, energy, government and other industries, making it faster and more reliable. Mining data value from these massive data becomes a major direction of digitalization. But meanwhile, massive data put huge pressure on the data processing performance and call for continuous innovation in database technology.

GBase 8a MPP Cluster, developed by General Data Technology Co., Ltd. (hereinafter referred to as "GBase"), could be used to support structured big data processing. With its unique flat architecture, high availability and dynamic scalability, GBase 8a MPP Cluster provides a cost-effective large-scale distributed parallel database management solution for ultra-large data management.

The Intel® Select Solutions for GBase 8a MPP Cluster combine Intel® Xeon® Scalable processor, Intel® SSD D3-S4510 Series, and Intel® Ethernet Adapter X710-DA2 so as to significantly improve the performance of GBase 8a MPP Cluster and fully meet the needs of big data analysis for data-intensive industries.

GBase 8a MPP Cluster Supporting Big Data in Industries

In the big data era, massive structured data have been accumulated in industries during the long-term business development, and the amount is still growing rapidly. Typical structured data include transaction data from bank and telecommunication; transaction data, management data, financial data, and basic information data from traditional data centers; data from emerging industries like IoT and Internet, etc.. For these rapidly expanding databases, improving their processing performance is of vital importance for safeguarding critical business and mining data value.

For the processing of structured data, GBase 8a MPP Cluster has outstanding performance advantages over traditional databases. It supports data sets ranging from TBs to dozens of PBs and concurrency of more than 300 users. IO latency is as low as 10% of that of traditional databases, while the data analysis is more than 10 times faster than that of ordinary databases. In terms of scalability, each server of GBase 8a MPP Cluster uses local resources and has good horizontal scalability based on peer node flat architecture. The cluster computing performance and storage capacity can increase near-linearly as the cluster scales out, during which users do not need to suspend business.

GBase 8a MPP Cluster also has many big data processing technologies and features such as new column-based storage engine, extremely high data storage compression ratio, maintenance-free coarse-grained index, etc. Combined with MPP's efficient distributed computing mode and cost-based distributed intelligent optimizer, GBase 8a MPP Cluster can support PB-level structured data analysis applications. Meanwhile, it can also ensure high availability of clusters and multi-cluster support for big data and realize the deployment of multiple data centers through intra-cluster replica synchronization technology, dual active cluster technology, cross-domain cluster data transmission synchronization technology and virtual cluster technology.

GBase 8a MPP Cluster can meet the demand for structured data processing in data-intensive industries such as finance, telecommunications, energy, government and IoT, and support applications such as data query, data statistics, data analysis, data mining and data backup. Besides, it can be host database for data warehouse systems, BI systems and decision support systems.

Intel® Select Solutions for GBase 8a MPP Cluster

Intel® Select Solutions for GBase 8a MPP Cluster can improve data analysis performance with significantly shorter application response time for users in data-intensive industries. The solutions realize high performance, scalability and availability of database for industry users by leveraging advantages of Intel® Xeon® Scalable processor, Intel® SSD D3-S4510 Series, and Intel® Ethernet Adapter X710-DA2. Specifically, the solutions have the following characteristics:

- Selected Intel® Xeon® Scalable processor will give full play to the performance potential of GBase 8a MPP Cluster and deliver high Total Cost of Ownership (TCO). Industry users can choose to build GBase 8a MPP Cluster solution based on tested and verified Intel® Xeon® Scalable processor.
- Intel® Select Solutions and Intel® Xeon® Scalable processors provide compatibility, stability combinations of hardware and software, optimized for customers' workloads. They can significantly reduce industry users' cost in type selection and testing during earlier stage of project, and speed up system on-line.
- The Solutions provide detailed reference architecture to improve the flexibility of large-scale distributed database systems for enterprise customers. Customers can select freely based on current configuration architecture even when hardware of different brands are configured in the cluster or manufacturers needs to be changed in future expansion.

Hardware Selection

The Intel® Select Solutions for GBase 8a MPP Cluster integrate Intel® Xeon® Scalable processor, Intel® SSD D3-S4510 Series, and Intel® Ethernet Adapter X710-DA2, allowing industry users to deploy high-performance GBase 8a MPP Cluster distributed analytical database system fast on this basis.

Intel® Xeon® Scalable processor

With new features including Intel® Advanced Vector Extensions 512 (Intel® AVX-512), it provides workload optimization for advanced analysis, High Performance Computing (HPC) application and data compression, so as to improve throughput. These breakthrough technologies can provide Intel's data-centered computing product portfolio with economical high-capacity persistent memory capabilities, and significantly enhance the high-speed data processing capabilities of key computing applications such as databases and data warehouses.

Intel® SSD D3-S4510 Series

Based on Intel® 64-layer 3D NAND TLC media, Intel® SSD D3-S4510

What Are Intel® Select Solutions?

Intel Select Solutions are pre-defined, workload-optimized solutions designed to minimize the challenges of infrastructure evaluation and deployment. Solutions are validated by OEMs/ODMs, certified by ISVs, and verified by Intel. Intel develops these solutions in extensive collaboration with hardware, software and operating system vendor partners and with the world's leading data center and service providers. Every Intel Select Solution is a tailored combination of Intel® data center compute, memory, storage, and network technologies that delivers predictable, trusted, and compelling performance.

To qualify as an Intel® Select Solution, solution providers must:

1. Meet the software and hardware stack requirements outlined by the solution's reference-design specifications;
2. Replicate or exceed established reference-benchmark test results;
3. Publish a solution brief and a detailed implementation guide to facilitate customer deployment.

Solution providers can develop their own optimizations in order to give end customers a simpler, more consistent deployment experience.

performs significantly better than traditional mechanical disks. For GBase 8a MPP Cluster application, it can be used as a storage medium to meet users' performance requirements while controlling costs. In addition, the flexible workload function enables it to cover more workloads and deliver flexible capacity, durability and energy efficiency.

Intel® Ethernet Adapter X710-DA2

Intel® Ethernet 710 Series can accelerate performance of Intel® Select Solutions for GBase 8a MPP Cluster. 710 Series, with professional 10/25GbE performance, supports 1/10/25GbE single- or dual-port connections, functions well in PCI Express v3.0 x8 slot, and supports advanced functions such as Virtual Machine Device queues (VMDq) and Single Root I/O Virtualization (SR-IOV). In addition, compared with 40G network adapter, 710 Series can effectively reduce costs while meeting the bandwidth requirements, highlighting economy and practicality.

Verified Performance through Benchmark Testing

Intel® Select Solutions for GBase 8a MPP Cluster put forward uniform suggestions on software but different suggestions on hardware configuration, i.e. "Base" and "Plus", targeting at customers who prefer minimum cost and optimal performance respectively (the configuration list is shown in Appendix 1). We recommend that users of Intel® Select Solution take these configuration suggestions or use higher configurations.

Intel® Select Solutions for GBase 8a MPP Cluster are verified to meet a specified minimum level of workload-optimized performance capabilities. The testing is completed with the business intelligence computing test Decision Analysis* and the decision support system test benchmark Decision Support System*. Decision Analysis test is widely used to evaluate the application performance of decision support technologies, and provides comprehensive evaluation on the overall business computing capabilities of the system. Decision

Support System test measures query response time under single user model, query throughput under multi-user model, and data maintenance performance in more complex multi-user decisions.

In order to test the performance difference between "Base" and "Plus" configurations in Intel® Select Solutions for GBase 8a MPP Cluster, the query and input time (the shorter the better) are tested using Decision Analysis* and Decision Support System* respectively, and the test results are shown in Table 1. ¹

Table 1: "Base" configuration and "Plus" configuration testing

GBase	Decision Analysis at 10TB Total time (s)	Decision Analysis at 10TB Load time (s)	Decision Support System Total time (s)	Decision Support System Load time (s)
Base	31400	17900	67500	11900
Plus	9100	5400	26000	4400

Test results show that in the Decision Analysis test, the "Plus" Performance is more than 200% higher than the "Base" in Decision Support. In the System test, this increase is greater than 150%. "Base" shows the basic configuration of Intel® Select Solutions based on GBase 8a MPP Cluster, while "Plus" supports higher performance in the deployment of GBase 8a MPP Cluster by system integrators, solution providers and industry users. Industry users can choose more suitable configuration for your business according to actual workload and requirements.

- **Intel® Turbo Boost Technology:** For peak demand, Intel® Turbo Boost Technology enables the processor to run faster than the rated operating frequency, accelerates processor and graphics performance for peak loads.
- **Intel® Smart Speed Shift Technology:** Allows the processor to quickly select its best operating frequency and voltage for optimal performance and power efficiency without intervention from the operating system.

Technology Selections for Intel® Select Solutions for GBase 8a MPP Cluster

In addition to the Intel hardware foundation for Intel® Select Solutions for GBase 8a MPP Cluster, Intel technologies integrated into Intel® Xeon® Scalable processor provide further performance and reliability:

- **Intel® Platform Trust Technology (Intel® PTT):** This technology can be directly integrated into the Trusted Platform Module (TPM) in the chipset to protect encryption keys. The new Intel® PTT function is provided as an option rather than as a stand-alone chip, simplifying integration and activation.
- **Intel® Hyper-Threading (HT) Technology:** Enables more efficient use of processor resources and multiple threads to run on each core. Intel HT Technology also increases processor throughput, improving overall performance on threaded software, and keeps system responsiveness while running multiple demanding applications.

Intel® Xeon® Scalable Processors

2nd Generation Intel® Xeon® Scalable processors:

- Offer high scalability that is cost-efficient and flexible, from the multi-cloud to the intelligent edge
- Establish a seamless performance foundation to help accelerate data's transformative impact
- Support breakthrough Intel® Optane™ DC persistent memory technology
- Accelerate artificial-intelligence (AI) performance and help deliver AI readiness across the data center
- Provide hardware-enhanced platform protection and threat monitoring

Intel® Select Solutions for GBase 8a MPP Cluster feature Intel® Xeon® Gold processors.



Intel® Select Solutions for GBase 8a MPP Cluster Satisfying the Demand of Industry Users for Massive Data Processing

The verified and optimized Intel® Select Solutions for GBase 8a MPP Cluster provide optimized workload configuration for Intel® Xeon® Scalable processor. When industry users need to deploy GBase 8a MPP Cluster solutions, they can obtain configurations optimized and tested in advance through Intel® Select Solutions,

saving a lot of time on debugging and performance optimization, thus meeting the needs of massive data storage and processing, safeguarding critical businesses and preparing for digital transformation.

Learn More

Intel® Select Solutions: intel.com/selectsolutions

Intel® Xeon® Scalable processors: intel.com/xeonscalable

Intel® Select Solutions are supported by Intel® Builder: <http://builders.intel.com>

Appendix 1: The Base and Plus Configurations for Intel® Select Solutions for GBase 8a MPP Cluster

GBase/Min 4 nodes	Base	Plus
Processor	2 x Intel® Xeon® Gold 5118 @ 2.30GHz 12C or Intel® Xeon® Gold 5218 @ 2.30GHz 16C or higher	2 x Intel® Xeon® Gold 6140 2.30GHz 18C or Intel® Xeon® Gold 6240 2.60GHz 18C or higher
Memory	192 GB or higher (12 x 16 GB DDR4-2400)	192 GB or higher (12 x 16 GB DDR4-2400)
Boot Drive	480 GB or larger Intel® SSD DC S4510	480 GB or larger Intel® SSD DC S4510
Storage Drive	12 x 1 TB SAS HDD 10000RPM or larger, RAID50	12 x 960 GB Intel® SSD DC S4510 or larger, RAID50
Data Network	1 x 10 GB Dual-Port Intel® Ethernet Converged Network Adapter X710-DA2 SFP+ or better	1 x 10 GB Dual-Port Intel® Ethernet Converged Network Adapter X710-DA2 SFP+ or better
Management Network	Integrated 1 GbE or better	Integrated 1 GbE or better



¹ CONFIG1 - Base: Test by Intel & GBase as of 5/30/2019. 4-node, 2x Intel® Xeon® Gold 5118 Processor, 12 cores HT On Turbo ON Total Memory 192 GB (12 slots/ 16GB/ 2400 MHz), 1x 480 GB Intel® SSD DC S4500, 12x 1TB SATA 7200RPM HDD, Microcode:0x200005e, Centos 7.4, 3.10.0-957.12.2.el7.x86_64, GBase 8a MPP Cluster 8.6.2.

CONFIG2 - Plus: Test by Intel & GBase as of 5/30/2019. 4-node, 2x Intel® Xeon® Gold 6140 Processor, 18 cores HT On Turbo ON Total Memory 192 GB (12 slots/ 16GB/ 2400 MHz), 1x 480 GB Intel® SSD DC S4500, 12x 960 GB Intel® SSD DC S4510, Microcode:0x200005e, Centos 7.4, 3.10.0-957.12.2.el7.x86_64, GBase 8a MPP Cluster 8.6.2.

Component performance tests are measured using specific computer systems. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit intel.com/benchmarks

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark® and MobileMark®, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit intel.com/benchmarks

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com

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, the Intel logo, and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2019 Intel Corporation. All rights reserved.