


CASE STUDY

2nd Gen Intel® Xeon® Scalable Processors
Intel® Omni-Path Architecture (Intel® OPA)
High-throughput Computing
Materials Genome Initiative (MGI)



The Chinese Academy of Sciences Accelerates Materials Genome Research with Data-Augmented Supercomputing

Intel has partnered with Dell to help the Institute of Physics, Chinese Academy of Sciences, further optimize the compute and storage capabilities of its materials genome computation and data processing platform based on high-throughput computing technology for greater efficiency in new materials research and achievements sharing



“One of the key options for today’s materials research is to embrace new technologies and trends in high-performance computing (HPC) and leverage high-throughput computing (HTC) to accelerate the materials genome computation for greater efficiency of new materials R&D. With the introduction of advanced products and solutions such as 2nd Gen Intel Xeon Scalable processors and Intel OPA, Dell EMC’s PowerEdge servers, PowerVault ME4 Series and Isilon Series Storage, our new platform has earned praise from front-line materials researchers for its performance, accessibility, reliability, and serviceability that brings them greater momentum in new materials research.”

Miao Liu
Distinguished Research Fellow
Institute of Physics, Chinese
Academy of Sciences

Traditional materials research methods are gradually becoming a major bottleneck in the development of innovative industrial products and quality improvements due to the long R&D cycle and unpredictable success rates. To accelerate new materials research and strengthen nation-level industrial competitiveness, the Institute of Physics, Chinese Academy of Sciences (IOP-CAS), starts to rely more heavily on high-throughput computing (HTC), an emerging branch of high-performance computing (HPC). They have built a high-throughput materials genome computation and data processing platform to support the work on the Materials Genome Initiative (MGI) project, and provided data sharing services for more materials researchers across the country with the materials genome database and the cloud-based resources platform.

Different from conventional HPC, where the pursuit of faster speed and higher efficiency of individual computation tasks made “faster calculation” the key metric, HTC, a new type of data-augmented HPC technology, is focused on the ability to perform “more calculations”. In other words, HTC focuses on improving task parallelism and throughput for massive data processing. For this reason, HTC platforms tend to have higher or more demanding requirements for parallel processing performance, average response time, throughput, scalability, and cost.

To meet these practical requirements, IOP-CAS works with Intel, an innovator in core computing power, and Dell, a provider of servers, storage devices and solutions, to build a next-generation materials genome computation and data processing platform (the New Platform). The New Platform introduced several key computing and high-speed interconnect technologies such as 2nd Gen Intel Xeon Scalable processors, Intel OPA, as well as Dell EMC PowerEdge servers, Dell EMC PowerVault ME4 Series and Isilon Series storage devices and solutions, which incorporate Dell’s unique innovative capabilities and features. In addition to higher materials genome computational efficiency, the New Platform also provides a comprehensive shared materials genome database for a broad community of materials researchers.

Application Advantages for IOP-CAS:

- The deployment of 2nd Gen Intel Xeon Scalable processors allowed IOP-CAS to access outstanding high-throughput materials genome computing power at a reasonable thread cost, and confidently handle the tens of thousands of materials computation tasks every year, significantly accelerating the research of new materials in China;
- Dell’s proven computing/storage products and advanced hyper-converged devices together with 2nd Gen Intel Xeon Scalable processors and Intel OPA help the New Platform deliver a high-performance data closed loop for its high-throughput materials genome computation. This allows IOP-CAS to build a world-class materials genome database and helps it evolve into a source of power driving materials genome research in China.

If crude oil is the lifeblood of modern industry, then materials are its bones and muscles. Any attempt to make industrial products more competitive requires a continued process of “transmutation”. This means continuous R&D and introduction of new materials are inevitable. In the past, materials researchers mainly used the trial and error approach to study the structural characterization of different materials in order to discover the nature of new materials. However, as the composition of modern materials becomes more and more complex, especially with the emergence of composite materials, conventional experimental techniques no longer work in many circumstances. Issues such as long R&D cycles, high costs, and low success rates have become increasingly pronounced.

The launch of the Materials Genome Initiative (MGI)¹ on a global scale and the growing maturity of theoretical simulation methods such as Density Functional Theory (DFT) and the Monte Carlo Methods brings materials researchers breaking points. They are now more often leveraging strong parallel computing power from computer methods, HTC in particular, to screen materials in the way of simulating calculation.

Take high-profile semiconductor materials research as an example, the application and performance of each new material, such as the growth behavior of thin films on a substrate and whether there were deactivation or surface reconstruction, all required repeated testing and verification in the past. Today, with the DFT theory, researchers can use partially known density functional calculations to fit the total energy of the system and use methods of quantum mechanics to precisely calculate electron-electron interaction between atoms in a material to obtain data on the electronic structure, magnetism, and structural stability of more compounds, theoretically guiding material design and accelerating the application of more new materials to semiconductors.

As one of the top institutes for materials science research and application in China, IOP-CAS is staying abreast of the latest developments and is committed to applying HTC to the research of new materials. The results have been very encouraging with many breakthroughs in superconductor, nano, and power battery materials. IOP-CAS believes that one of key missions to accelerate new materials research in China is to continue incorporating the latest developments in Information Technology into materials research and build a comprehensive materials genome data platform with greater and more efficient computing power. This will help researchers focus more on the materials research itself instead of dealing with sophisticated procurement for and construction of such a computing environment.

IOP-CAS has worked with Intel and Dell to build a solid IT infrastructure foundation for the next-generation high-throughput materials computation and data processing platform. By introducing advanced computing and interconnect technologies, along with mature data processing/storage products and solutions, the platform is a powerful engine in China's materials genome computation and research.

Breakdown of the High-throughput Materials Genome Computation Solution

As mentioned above, the HTC solution plays a key role in IOP-CAS' New Platform. To put it simply, it is a computation system with the ability to process large volumes of independent tasks in parallel. The following features distinguish it from conventional HPC solutions:

- **Processing:** Tasks are generally processed with thread-level parallelism and workloads constantly change to the task requirements;
- **Efficiency:** It is more focused on the integration of data with computation and the read/write performance has a direct impact on system efficiency;
- **Performance:** It is more focused on increasing concurrent computation and data processing requests within a unit of time;
- **Cost:** There is positive correlation between system processing capability and thread counts, and prioritization of single-thread cost improves overall processing performance.

In general, HTC that can process large volumes of workloads is well suited to scenarios that involve the screening of numerous samples, for example, the R&D process in biology and pharmaceuticals etc. HTC can also be used in simulating calculation and screening massive materials genomes during materials genome research.

As shown in Figure 1, materials genome computation can often be simply divided into the following steps:

- 1 The system selects data from an external materials genome database and generates a file that can be called by specific simulation software (such as VASP²);
- 2 Then, with high-performance parallel computing, it obtains data about corresponding material properties such as energy density, electronic structure, and energy increase/decrease in synthesis;
- 3 The computation results are kept in the internal storage system for further analysis. A high-performance parallel distributed file system is usually needed;
- 4 Once the data has been post-processed and integrated, further in-depth analysis can be performed with the results being used for materials screening. The analysis results can also be imported to the materials genome database. Thus, a complete procedural computation process is formed.

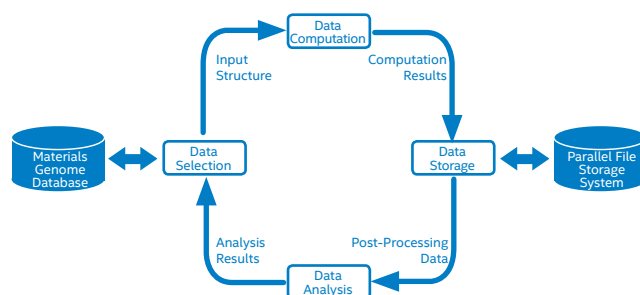


Figure 1. Materials genome computation workflow³

The generation of structure files, theoretical simulation and computation, and analysis of calculation results in the processes above can be converted into a set of independent instructions and tasks in a certain way. These instructions and tasks can be executed by specific software packages, which provides the perfect scenario for HTC applications.

In addition to executing high-throughput materials genome computation tasks, IOP-CAS is planning to build a data center for the New Platform to store calculated data for subsequent analysis. It is also going to provide a shared data platform in a private cloud to share data, codes, and computing tools used during materials genome computation. This will enable more materials researchers across the country to focus on their research instead of having to spend time, effort, and money in procuring the same type of equipment and set up a similar environment.

With all these ideas, IOP-CAS plans to deploy its high-throughput materials genome computation solution through four key software packages: the first is the high-throughput dispatching package based on Simple Linux Utility for Resource Management (Slurm). It will ensure optimal execution performance through the dispatch and monitoring of all computational and analytical tasks; the second is the data archiving package, which can perform high-speed archiving of calculation results in the Lustre parallel distributed file system; the third is the post-processing package that performs inference and integration of the raw calculation results for subsequent data analysis or for graphical displays/queries; and the fourth is the big data analysis package, which helps researchers utilize artificial intelligence (AI) and other technologies to screen new materials and predict synthesis outcomes more efficiently with a massive volume of calculated data and analysis results.

Excellent design ideas require the support of appropriate hardware infrastructure. This represents no small challenge to IOP-CAS that is more focused on materials research and lacks IT expertise and experience. To better unleash the potential of the New Platform and maximize its performance in future computation tasks, IOP-CAS decided to “leave it to the experts” by bringing in Intel and Dell as its innovation partners. They joined hands in developing a high-performance and high-availability compute and storage hardware architecture tailored to its needs based on the comprehensive software framework above. IOP-CAS raised several requirements for the New Platform's infrastructure before the three parties commenced their work:

- **Advanced Technology:** The compute and storage hardware architecture of the New Platform must be able to support continued increases in high-throughput materials genome computation tasks over the coming years while maintaining a performance advantage;
- **Operational Stability:** Once the New Platform has been built, it will need to handle massive high-throughput materials genome computation tasks. Any system failure will incur unpredictable losses, making the reliability and stability of the architecture incredibly important;

- **Simple Operations and Maintenance (O&M):** IOP-CAS' relative lack of human resources means that no dedicated O&M experts can be assigned to the platform, so the O&M of the compute and storage hardware architecture must be agile and easy.

Intel Xeon Processors Delivering More Computing Power and Optimal Parallelism at the Same Cost

Incorporating advanced computing technology into the New Platform was IOP-CAS' foremost requirement. That is, its requirement for computing power or performance must be met.

The computing power of the New Platform mainly comes from processors, so choosing an appropriate primary processor that could satisfy the high-throughput materials genome computation requirements of IOP-CAS was a top priority in the construction of the platform. It needs to be considered thoroughly and balanced across multiple dimensions. To be specific, as HTC is more focused on parallel processing which requires the computation system to simultaneously process as many tasks as possible, and features a positive correlation to thread counts, the number of cores and threads in the processor was the foremost consideration. Secondly, the L3 cache (last-level cache) is shared by all cores and threads in a processor, so a larger size of L3 cache needed to be considered to ensure more caches were allocated to each parallel task, increasing the cache hit ratio. Thirdly, a higher processor clock speed is essential to increase the processing speed of individual tasks. Finally, prices or costs also had to be considered. Processors with higher configurations may be better at meeting the above specifications and metrics, but they also increase the cost. A more appropriate choice would be to maximize the return on capital expenditure (CapEx) without compromising parallel processing performance. As such, the principle of prioritizing single-thread cost must be followed.

To help IOP-CAS select a higher-affordability processor, Intel worked with experts at the institute to conduct thorough tests and assessment of several Intel® processors. The Intel® Xeon® Gold 6230 processor was IOP-CAS' initial consideration. However, new plans that required the New Platform to undertake much more computations in the future led IOP-CAS to switch to the Intel® Xeon® Gold 6248 processor with a higher clock speed. But after helping IOP-CAS conduct thorough tests, Intel made the following suggestion: Though the higher clock speed of the Intel Xeon Gold 6248 processor could enhance performance in some areas, it had the same number of cores and threads as the Intel Xeon Gold 6230 processor, so this enhancement did not offer a breakthrough in performance. With careful evaluation of the upgraded SKUs of the 2nd Gen Intel Xeon Scalable processor family released at the start of 2020, Intel suggested IOP-CAS choose the Intel® Xeon® Gold 6230R processor.

The upgraded SKUs in 2020 are all enhanced in Turbo, core/thread count and cache size compared to the 2nd Gen Intel Xeon Scalable processors released in 2019. Take the Intel Xeon Gold 6230R processor as an example, its core counts, thread counts, and cache size are all 30% bigger than the previous model (Intel Xeon Gold 6230 processor), and its Turbo frequency is increased from 3.9 GHz to 4.0 GHz, as shown in Table 1, while the cost remains roughly the same.⁴ It means this choice would significantly improve the CapEx returns of the New Platform.

	Clock Speed	Turbo	Cores	Threads	Cache
Intel® Xeon® Gold 6230R processor	2.10 GHz	4.00 GHz ↑	26 ↑	52 ↑	35.75 MB ↑
Intel® Xeon® Gold 6230 processor	2.10 GHz	3.90 GHz	20	40	27.5 MB

Table 1. Comparison of Intel® Xeon® Gold 6230R processor (2020) and Intel® Xeon® Gold 6230 processor (2019) in key product specifications

In the meantime, as IOP-CAS may require even higher compute performance in certain scenarios, Intel also recommended the Intel® Xeon® Platinum 9242 processor. With 48 cores and 96 threads for high-concurrency and 12-channel DDR4 memory support, the Intel Xeon Platinum 9242 processor offers even more computing power to enable IOP-CAS' further exploration in applications such as materials genome algorithm optimization and material property prediction based on machine learning and deep learning.



Hailed as the “Bedrock of the Modern Data Center”, the **Dell EMC PowerEdge server** is a critical component driving the computing power of the next-level high-throughput materials computation and data processing platform. Featuring a highly scalable system architecture and flexible internal storage, it can provide even higher performance for a wide variety of applications together with 2nd Generation Intel Xeon Scalable processors. Its features such as “pre-configured VMware virtualization software with cloud DNA”, “highly scalable business architecture”, and “smart automation that is easy for remote O&M” are invaluable to enterprises seeking to accelerate IT transformation and business innovation.⁵

Combination of Dell's Storage Products with Intel OPA for Seamless Computation and Data Integration

Having addressed the computing requirements of the New Platform, Intel, Dell, and IOP-CAS began looking into how to enhance the integration of computation and data. This is exactly what the HTC technology means for “data augmented”, and cares besides parallel computing power — it has more demanding requirements for throughput, latency, and bandwidth of storage modules.

As such, Intel and Dell chose to introduce advanced software/hardware products and technologies such as Dell EMC PowerVault ME4 Series Storage, Dell EMC Isilon Series Scalable NAS Storage, Dell EMC VxRail Hyper-Converged Infrastructure, and Intel OPA to the New Platform.

As Figure 2 shows, to meet IOP-CAS' requirements for materials genome computation, data storage, data analysis, and visualization, the hardware architecture of the New Platform is divided into two platforms — one for high-throughput materials genome computation, and one for materials genome data processing. The former is primarily responsible for HTC tasks and high-speed file storage in the calculation process, while the latter provides data storage, analysis, and cloud-based sharing. A 10 Gbps production network is used to connect the two platforms.

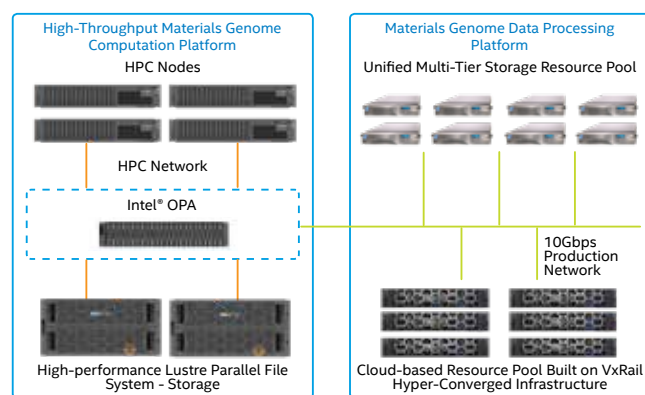


Figure 2. The primary hardware architecture for IOP-CAS' next-level high-throughput materials genome computation and data processing platform⁶

In the process of high-throughput materials genome computation, hundreds of thousands of parallel and independent computation tasks generate large amounts of process files. These files need to be stored at high speed for subsequent data inference and integration. To this end, the New Platform utilizes the high-performance Lustre parallel file system built on two Dell EMC PowerVault ME4084 High-Density Storage devices (2*84*8 TB, for a total capacity of 1344 TB) based on Intel processor architecture. The Lustre file system supports data storage in hundreds of PBs and concurrent bandwidth in TBs, and offers high scalability, so IOP-CAS can upgrade the ME4 Series Storage devices flexibly and easily as necessary.



Dell's advanced architecture-based **Dell EMC PowerVault ME4 Series Storage** is equipped with Intel Xeon processors and optimized for SAN/DAS to enable enterprise users to quickly build and integrate agile, high-performance, and cost-effective storage systems. The ME4 Series featuring extensive capabilities supports all driver types, a multitude of protocols, and turnkey software functionality. It provides many options of scalable storage modules, and is suitable for accelerating enterprise applications such as HPC, backup, and VDI.⁷

Materials genome computation produces large amounts of unstructured data. Take electronic structure (electron density) as an example, the electron density file of a single material could be a single image, up to 10 GB in size. The New Platform must therefore provide a high-performance and scalable storage system for massive outcome data. Dell's solution was the Dell EMC Isilon Series Scalable NAS Storage devices, known for outstanding efficiency and excellent scalability. Four Isilon H400, a hybrid scale-out NAS system, and eight A2000 Isilon, a scale-out NAS system for archive data storage, were used to develop a unified multi-tier storage resource pool for data analysis and secure data storage respectively. Both products offer excellent scale-out capability and can help the platform expand storage capacity flexibly. Their built-in OneFS operating system provides up to 80% of storage utilization, and a variety of data protection and security measures to ensure the safety and reliability of platform data.⁸ This also satisfies IOP-CAS' requirements for stability and reliability of the New Platform infrastructure.



Dell EMC Isilon Series Scalable NAS Storage based on the OneFS operating system is a hybrid NAS array with various capabilities. It offers a better balance between performance, capacity, and value. With multiple built-in protocols, it provides users with more agile and flexible interoperability and helps enterprises support workloads involving a broad range of unstructured data on a single platform for files/data consolidation, elimination of costly storage silos, and streamlined management.⁹

Meanwhile, Dell also provided the New Platform with six Dell EMC VxRail hyper-converged appliances to develop a cloud-based resource pool for resource sharing. As mentioned above, one of the objectives for the New Platform is to share outcome data, codes, and computation tools used during materials genome computation to help improve the efficiency of more researchers in the field of materials. VxRail Hyper-converged Infrastructure can be fully integrated into Dell's Software Defined Data Center (SDDC) software system. This makes it easy for the platform to deploy a cloud environment based on VMware Cloud Foundation using VxRail hyper-converged appliances. The advantage of such a fully-integrated platform is: it greatly simplifies the complex process of setting up a private cloud from planning, to construction and deployment. The complexity of subsequent O&M is also greatly reduced, providing IOP-CAS with an all-in-one cloud solution.



The **Dell EMC VxRail hyperconverged appliance**, a co-engineered hyperconverged platform with VMware, provides users with one-click-deployment, one-click-upgrade, and out-of-the-box cloud platform delivery capabilities, and enables full lifecycle management from hardware to software, and to the cloud platform. This simplifies O&M and management of the cloud platform, allowing users to scale on demand and allocate resources flexibly as they grow in size.¹⁰

To ensure that storage devices and compute nodes work together more efficiently, Intel OPA was used on the New Platform to improve interconnect efficiency and scalability. Intel OPA optimizes data stream management and enhances protection of packet integrity to further reduce latency through technological innovations at the link layer. More importantly, the Intel OPA-based switch chip supports up to 48 ports, allowing the New Platform to support up to 1152 ports with just one single switch. This exceptional scalability satisfies IOP-CAS' requirements for future upgrades and expansions of the New Platform.

Outcome and Prospects

Through its three-way partnership with Intel and Dell, IOP-CAS has now deployed 160 Intel Xeon Gold 6230R processor nodes and 20 high-performance Intel Xeon Platinum 9242 processor nodes on its new high-throughput materials genome computation platform. A complete, unified, and multi-tier storage resource pool and a cloud-based data platform for resource sharing have also been set up in the data center where the materials genome data processing platform is located.

Once all the above nodes are put in service, around 3,000 DFT-based tasks can be executed and 100-700 (depending on material complexity) calculations for inorganic crystalline materials genomes will be completed on a daily basis. All inorganic crystalline materials known to mankind (around 100,000 after deduplication) could be analyzed within a single year. In the next 3-5 years, IOP-CAS will use the computing power from the New Platform to continue exploring materials not yet known to mankind and incorporate the prediction data for 500,000 unknown materials into the platform's materials database.¹¹ This serves to transform the materials research approach from conventional "discover-calculate" to "calculate-discover", converting Information Technology into a driver of new materials research.

In addition to the acceleration of materials research, the New Platform will provide support for the integration of industry, universities, and scientific research institutes. The cloud-based shared resource pool will allow students, teachers, and researchers from universities, scientific research institutes, or businesses to utilize existing materials

genome data on the data platform to conduct new materials synthesis quickly, avoiding the waste of resources caused by duplicate calculation. Moreover, with rich data accumulated on the data platform, IOP-CAS is able to visualize the internal structures of various inorganic crystalline materials to teachers and students in a variety of ways. Virtual Reality (VR) technology could even be introduced to make the study and research of materials "tangible".

Looking forward, IOP-CAS also plans to work with its IT partners such as Intel and Dell on further optimization and expansion of the materials genome computation platform, so that HTC nodes will be increased to around 500, and computing power be doubled¹², making IOP-CAS' materials genome data platform one of the best in the world. The computing focus of the New Platform will also be expanded from inorganic crystalline materials to molecular materials. The scope for new materials research will continue to broaden. Intel and Dell will leverage their own strengths and help IOP-CAS make more achievements in materials genome computation with their evolving products, technologies and solutions.

Solution Providers



戴尔科技集团



¹ For more information about the Materials Genome Initiative, please visit: <http://www.mgi.gov>

² VASP (Vienna Ab-initio Simulation Package) is a commercial software package commonly used in modern materials research. It was developed by Hafner et al. at the University of Vienna for electronic structure calculations and quantum-mechanical molecular dynamics simulation. For more information, please visit the VASP website at: <https://cms.mpi.univie.ac.at/vasp/vasp/vasp.html>

³ The workflow partially references Jain A, Hautier G, Moore C J, et al. "A High-Throughput Infrastructure for Density Functional Theory Calculation" [J]. Computational Materials Science, 2011(50): 2295 - 2310, Link: <https://www.sciencedirect.com/science/article/abs/pii/S0927025611001133>

⁴ Based on the processor pricing published on the Intel website: <https://ark.intel.com/content/www/us/en/ark/compare.html?productids=192482,192446,192437,199346>

⁵ For more information about Dell EMC PowerEdge servers, please see the introduction on the Dell website at: <https://www.delltechnologies.com/zh-cn/servers/index.htm#section=rack-servers>

⁶ Management, service, and application nodes on the platform are not marked in the architecture diagram

⁷ For more information about Dell EMC PowerVault ME4 Series Storage, please see the introduction on the Dell website at: <https://www.delltechnologies.com/zh-cn/storage/powervaultme4.htm#scroll=off>

^{8,9} For more information about Dell EMC Isilon Series Scalable NAS Storage, please see the introduction on the Dell website at: <https://www.delltechnologies.com/zh-cn/storage/powerscale.htm#scroll=off>

¹⁰ For more information about Dell EMC VxRail hyper-converged appliances, please see the introduction on the Dell website at: <https://www.delltechnologies.com/zh-cn/converged-infrastructure/vxrail/index.htm#scroll=off>

^{11,12} Cited from IOP-CAS' internal data. For more information, please contact IOP-CAS at: <http://www.iop.cas.cn/>

For more information about IOP-CAS' materials database, please visit: <https://atomly.net>

Intel does not control or audit third-party data. You should review this content, consult other sources, and confirm whether referenced data are accurate.

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

No product or component can be absolutely secure.

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

© Intel Corporation