

DX2000 White Paper

Scalable Modular Platform for Data-Intensive Applications

Contents

1. High density server ideal for big data analysis infrastructure Scalable Modular Server DX2000	2
2. Key highlights of DX2000	3
3. DX2000 use cases.....	4
4. Summary	8



1. High density server ideal for big data analysis infrastructure

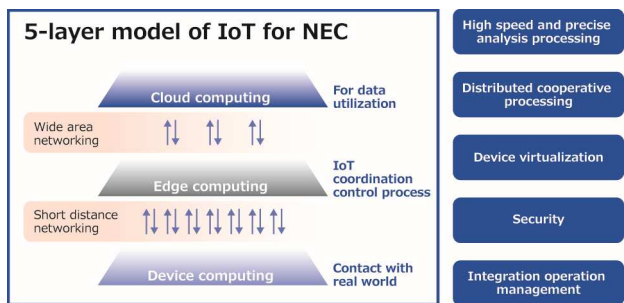
Scalable Modular Server DX2000

1.1 What is required of ICT platforms for IoT computing?

With the evolution of technology, we are approaching the age of IoT, in which all things connect to the Internet. It will be necessary to collect real time not just structured data, such as conventional business system data, but also explosively increasing data including unstructured data such as data generated from IoT devices such as sensors, and character data generated from notification service. ICT platforms are required to produce opportunities for creating new social value by converting the data into information and knowledge through advanced analysis, and providing real-time feedback.

In the 5 layer model of IoT which NEC has imagined, security and real-time capabilities are secured, and a system is provided to seamlessly connect IoT devices, edge computing, and cloud computing. We are promoting the development and sales expansion of products with the following points as a guideline to enhance ICT platforms.

- "High-speed and high-accuracy analytical processing" to process and effectively use in real-time large amounts of data collected from countless devices
- "Distributed cooperative processing" runs applications at the optimal level according to fluctuations in load, to create an efficient system
- "Device virtualization" securely and efficiently processes data collected from countless devices
- "Integrated operation management" includes ensuring 5-layer security, as well as devices and networks




Furthermore, NEC is developing server that utilizes "Real-time Capability", "Efficiency" and "Flexibility" as keystones which is necessary for platforms that rapidly process large volume, large-scale, and diverse data collected from IoT devices.

"Real-time capability" requires the analytical capabilities of processing explosively growing data in a shorter time. "Efficiency" requires superior scalability to maximize return on investment while facing demands to enhance the computing power of hardware such as CPU and memory. "Flexibility" requires the simplicity of providing

services according to client business and investment plans that can be scaled up smoothly.

NEC provides the Scalable Modular Server DX2000 as a highly integrated platform optimal for the analysis of large and diverse data collected from IoT devices.

There are three requirements for an ICT platform that generates new value.



1.2 The big data rapid analysis platform DX2000

DX2000 is a high density server that can mount 44 server modules with independent OS running in a 3U size chassis. Powered by the latest Intel® Xeon® D processor, DX2000 offers high speed, low power consumption and hence better TCO. Therefore, you can easily implement a high-efficiency distributed system with several servers in one chassis. With the maximum configuration, i.e 176 cores and 2.8TB memory in one 3U chassis, DX2000 can offer a data transfer of 1502GBs, 4.9 times faster than a conventional 4-core rack server. By leveraging this broad memory bandwidth to process large-scale data on the memory in real time, it can process analyses that would take a few hours on a conventional system in just a few seconds to a few minutes.

In addition, since it can accommodate 572 servers in one rack, setup space is greatly reduced compared to the conventional server implementation. Moreover, the chassis power supply unit and fan are present in redundant configuration, which helps avoid system crashes of the entire chassis due to failures of shared components. The DX2000 platform also contributes to reduce OPEX through advanced power saving technologies such as power-saving CPUs, high-efficiency power sources, and optimized cooling structure.

* if the same amount of 1U rack servers (1CPU configuration) were installed



2. Key Highlights of DX2000

2.1 Architecture that combines high density with great performance

DX2000 was developed to create higher performance in a small footprint. Consequently, the total performance per 1U rack space has been greatly improved. Moreover, it uses an architecture designed for network performance needed for in-memory distributed processing. Furthermore, since only industry standard components are used, it easily achieves high performance without using special technology requirements for the OS or software.

If you compare these features with a typical 1U size rack server, the integer computing performance is roughly 1.7 times better, the total bandwidth (GB/s) between CPU and memory is 4.9 times better, and the bandwidth per memory capacity is 4 times better. (Compared with a 3U size rack server with equivalent memory capacity and storage capacity)

Comparison of performance specifications

	3 conventional rack servers	1 DX2000
Integer arithmetic capability*	1	1.7
CPU-memory bandwidth [GB/s]	307GB/s	1,502 GB/s
Bandwidth [BW/GB] per memory capacity	0.133	0.533

* ExpressS800/R120F-1M (2 built-in Xeon E5-2699 v3 [18-core]) x 3 vs DX2000 (Xeon D-1527 [4-core]). Equivalent total memory, comparison with storage capacity

Take a look at the performance of Apache Hadoop/Apache Spark for a distributed processing network used in big data analysis, in a benchmark test run with DX2000 (8-core model), the price/performance compared to a 1U rack server with the above configuration is roughly 3 to 6 times better.

2.2 Scalability from small scale to large scale

DX2000 can be implemented with as few as 14 server modules in one chassis, and you can scale up the system according to data and processing requirements. Easy scale out with the module servers offers great advantage for data intensive IoT environment.

Nevertheless, concern of running out on network bandwidth with addition of new server module is reasonable. To that, DX2000 comes with two in-built switches offering two high speed 10G ports per server module with compliance to network standards for low latency backplane. We provide an architecture that allows you to connect server modules without cables, and achieve high-speed communication and low latency. Consequently, we found in benchmark tests on Hadoop that doubling the server modules increases performance by 1.8 times, which means scaling up almost linearly*. This sort of scalability is important for smooth scalability.

Moreover, securing network bandwidth between chassis is important while scaling in and out chassis configuration. By connecting to up to 16 external high-speed networks of 40GbE, DX2000 achieves 640Gbps uplink, and sufficient network bandwidth. It is also ideal for applications demanding for parallel distribution such as Hadoop/Spark as scaling out to several thousand nodes is simple.

Kmeans *performance




* Kmeans: One of the Hibench benchmark tests for determining the Hadoop/Spark performance index. With this performance index for processes used in areas such as machine learning, randomly allocated data is classified based on the measure of similarity between data.

2.3 Smooth Scalability with ease of installation

Vast number of cables becomes a bottleneck when adding servers by chassis. DX2000 chassis has in-built switch, so the number of cables is reduced by sharing the power source. All you need to do is connect 3 power cables and several network cables to use it. This means a reduction of up to 90% in cables compared to a rack server with equivalent configuration. Further using in-built switch as access instead of the ToR switch, can save extra money.

* if you consolidate 22, 1U rack servers into 1 DX2000

Set up and deploy a configuration with equivalent performance using less cables and without difficulty



Add servers simply by inserting server modules into the chassis. Since the platform supports hot-swap, you can add servers and replace defective servers without shutting down the entire system. There are 3 power sources built in to the chassis, and fan modules are configured redundantly, so there is no need to stop operation while recovering from failures.

When adding a large amount of servers, one concern is an increased management burden. To that end, DX2000 does not have general chassis level management modules on blade servers,

but rather uses the same management system as conventional rack servers, allowing for operation with the same image.



3. DX2000 Use Cases

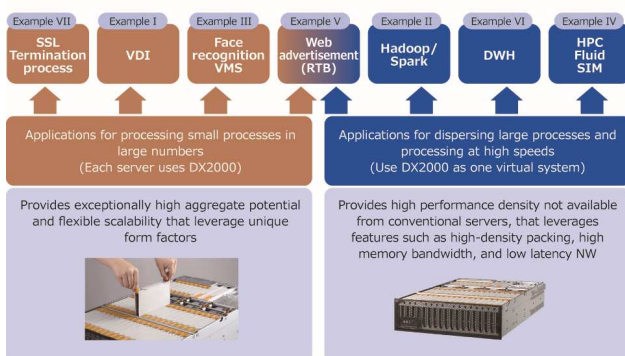
Having achieved efficient system management, the highly integrated and readily scalable DX2000 solves the problems of IT systems in the IoT era. There are two types of conceivable applications.

The first is applications for processing large numbers of small processes. Using the server modules of DX2000 as independent servers, superior consolidation and flexible scalability work effectively on large-scale systems. For example, thin client servers, image and face recognition servers, and SSL terminal servers. With DX2000, hardware configuration work is greatly reduced, and new installations are simple. Therefore, aside from initial installation, you can also scale up the system according to increases in processing quantity.

The second application is to distribute large, lumped calculations to multiple servers and run parallel processing. By using the entire DX2000 as one virtual system, properties such as high density implementation, high memory bandwidth, and a high-speed network show their effectiveness. Examples include big data processing with Hadoop/Spark, data warehouse (DWH), and fluid analysis processing in the HPC field.

From the next section, we have explained seven use cases leveraging the aforementioned applications:

Examples of leveraging DX2000 features



3.1 Consolidation of Citrix XenApp system

• Issues when running a thin client

Companies and public institutions are increasingly introducing thin clients which leave no data on client terminals from the viewpoint of making client PC management more efficient and strengthening the security level. However, managing thin client servers is an issue there. The more employees and staff there are, the number of servers roll into a snowball, and the burden on installations, settings, and operation management increases heavily.

• Resolving issues with DX2000

DX2000 with its high density servers excels at saving space and power, and is ideal for consolidating thin client servers. Even when adding new users, or during a server failure, you can add or replace server modules with hot-swap, so there is no need to shut down the system.

An advantage is that impact from system failure is minimized to the greatest extent. For example, when configuring a session with Citrix XenApp, the number of extension sessions per node is 250 for a 1U rack server or blade server, but 70 with DX2000 (8-core

Comparison at time of XenApp 1500 session configuration

	1Urack	Blade	DX2000
Aggregate potential	8U	6U	3U
Scalable unit Scalable sessions per node	roughly 250	roughly 250	roughly 70
Power consumed* [W]	When on standby, roughly 1.3K Under a high load, roughly 3.4K	When on standby, roughly 1.3K Under a high load, roughly 3.4K	When on standby, roughly 1.3K. Under a high load, roughly 2.8K


* Estimate from catalog specs (reference values).

model). Therefore, even during a system shutdown of a server, including software failures, the impact is limited to at most 70 sessions, so the damage is minimized to the greatest extent compared to rack servers and blade servers for which 250 sessions are impacted.

• Superior compatibility with Citrix XenApp

Citrix XenApp is equipped with a function for using the NetBoot type Citrix Provisioning Services (PVS) to transmit a clean OS image to XenApp servers at boot. With PVS, you can share an OS image even if there are many XenApp servers, simplifying maintenance duties such as applying patches and adding applications. Since you can use the OS image of a PVS server as is when you add a new user or add/replace a server due to a failure, adding hardware and troubleshooting is comparatively easy. In addition, even in case of a virus infection or the OS or application suddenly malfunctioning, you can resolve the issue by returning the OS image on the PVS server to a recover point.

Thin client image of XenApp/PVS



3.2 Infrastructure for Big Data Analytics (Hadoop/Spark)

• Issues with big data analysis processing

Open source-based Hadoop/Spark attracts attention as a distributed processing framework for big data. With the emergence of the in-memory based Spark, there is growing hope not only for large capacity data and batch processing, but also for social analysis that processes analysis of memory-stored data in real time, and for applications to machine learning.

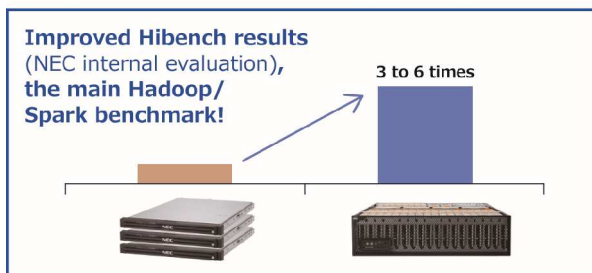
The Hadoop/Spark environment can be built using typical servers, but the number of servers will be huge, making it troublesome to scale out. Moreover, it is not easy to introduce new technology at a moment's notice.

• Resolving issues with DX2000

Since DX2000 can incorporate a large amount of server modules, and connect modules with a high-speed network, it is ideal as an operation base for Hadoop/Spark as the environment with distributed parallel processing is easy to create.

The run performance of Hadoop/Spark shows in measurements a price performance of 3 to 6 times improved in the Hibench comparison with a rack server that runs with equivalent installation size and memory capacity.

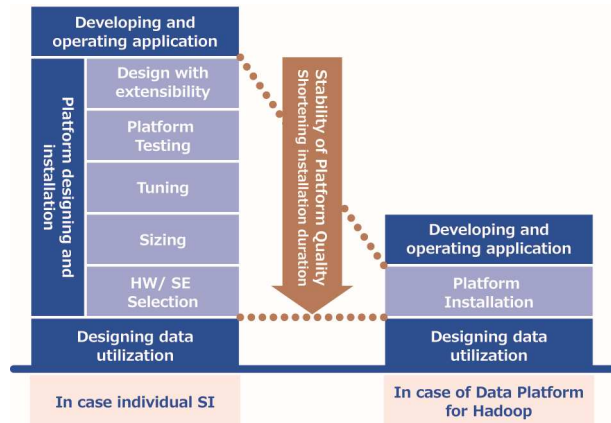
Hadoop/Spark performance 3 to 6 times* the price performance of conventional servers for many processes



* Express5800/R120F-1M (2 built-in Xeon E5-2699 v3 [18-core]) x 3 vs DX2000 (Xeon D-1541 [8-core]). Equivalent total memory, comparison with storage capacity

• Easy to implement

You need individual SI and operation to implement Hadoop/Spark, so related know-how is required. We at NEC offers NEC Smart Data Analytics which provides Hadoop/Spark infrastructure verified within NEC as a simple implementation solution for reducing such burdens. This solution has been verified and optimized the hardware and software for Hadoop/Spark that will reduce implementation period, provide stable quality and contribute on total cost reduction.



3.3 Infrastructure for Image Recognition (Face Recognition) System

• Issues with image recognition system

The face recognition system which uses NEC face recognition technology known for its world class* precision and speed has been adopted for a wide variety of applications from PC login authentication to security management applications such as access control to offices and facilities, entrance based on recognition in the entertainment sector, stadium surveillance and building monitoring, factory monitoring

If you are refining recognition when implementing these systems, more cameras are beneficial, but more servers for managing the cameras will be required. If camera resolution improves, and you use high-definition cameras or 4K cameras, the number of servers will further increase. Consequently, the root of the problem is that burdens of operation grow vastly such as installation and management of servers.

* Face recognition technology benchmark test (FRVT2013) of the National Institute of Standards and Technology (NIST)


• Resolving issues with DX2000

With the high density server DX2000, you can manage twice as many cameras or more compared with 3 typical 1U size rack servers.


The number of cameras managed by one server module with DX2000 is less than the amount of cameras managed by one rack server, which minimizes the impact on monitoring systems during

server failures including those involving software.

System configuration image



**High aggregate potential:
Aggregate twice as much
than conventional servers or more**



* 3 Express5800/R120g-1M (with 2 built-in Xeon E5-2699 v4 [22-core]) vs DX2000 (Xeon D-1571 [16-core]).

● Working together with Hadoop/Spark

If you use DX2000, you can send face image data acquired with the image recognition system to a system using Hadoop/Spark running on DX2000 to verify the face, cross-check images and identify suspicious persons as soon as possible, averting a crisis. Since the image recognition (face recognition) system and big data analysis can all be run on DX2000, there are also benefits from consolidation.

3.4 HPC Fluid Simulation Analysis system

● Issues with fluid simulation analysis

Fluid simulations are widely used in the research institutions of science and technology departments, as well as for product development in the manufacturing industry. Systems that connect super computers and general rack servers with high-speed interconnect (Infiniband) are used for these computations, but installation space requirements, power consumption, and hardware implementation costs are the concerns for securing the necessary processing performance.

● Resolving issues with DX2000


The high density server DX2000 in a limited space has a wide memory bandwidth with parallel processing provides an optimal case for performance requirements in fluid computations and simulation. With that, you can achieve a superior performance at better cost without using an expensive interface such as Infiniband.

Moreover, the floating-point arithmetic per unit of power can be improved to roughly twice with DX2000 than with a typical rack server.

In addition, it is fully compatible with Intel Xeon processors, so there is no need for special tuning in order to use applications, and open source applications can also be used directly.

* Express5800/R120f-2M (Xeon E5-2667 v3 [8-core]) vs DX2000 (Xeon D-1541 [8-core])

Proposal image of fluid calculation OSS (OpenFOAM) platform



● Applicable even outside of fluids analysis

The DX2000 is ideal for distributed parallel processing, and can be used for simulations in areas of fluids simulation analysis, genome analysis, drug trials, finance, and marketing.

3.5 Online, Real-time Bidding system for Web advertisements


● Issues with bidding system for Web ads

Real Time Bidding (RTB) involves bidding for online advertisements in auction format, and requires low latency. With RTB, each time a website is viewed; competitive bidding is run for that ad spot. It only takes milliseconds for the whole process of the website being viewed, bidding being carried out, the advertiser being determined, and the ad being displayed. It takes enormous processing to run RTB for each web visitor. Therefore, there is a need for a highly automated processing with very low latency to handle enormous data near the Internet Exchange where Internet connection points and traffic are concentrated.

● Resolving issues with DX2000

DX2000 has excellent processing capabilities ideal for low latency RTB, and can thus rapidly process bidding that occurs on a large scale in a short time. With the hyper-dense implementation, DX2000 offers great capacity in a small footprint and hence more computation resources near the highly responsive Internet Exchange.

RTB system image



3.6 Data Warehouse (DWH) System

• DWH system issues

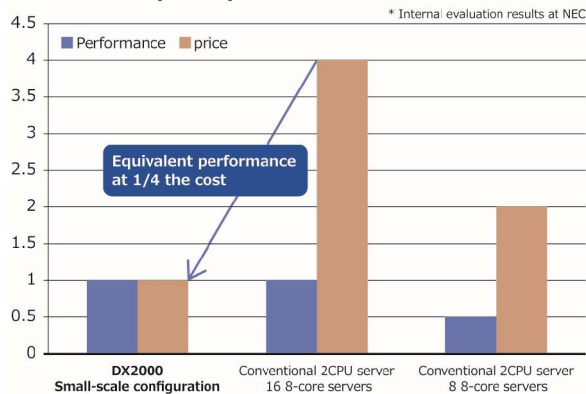
The DWH system processes and analyzes a large volume of operational data and external data accumulated from day-to-day corporate operations such as corporate sales, transaction data spewed out from production systems and sales systems, etc. Now converting that data into information requires high-speed analysis capabilities. At the same time, since the DWH system is large-scale, cost of implementation and space is also an issue.

• Resolving issues with DX2000

The DX2000 with its superior capabilities and high density offer high performance in a compact size. According to NEC's internal analysis, 14 DX2000 servers provide the performance of a conventional DWH set model (16 rack server configuration), saving hardware costs to 1/4.

Roughly 80% of DWH's typical system configuration is built with data capacity of 20TB or less, which is provided by DX2000 in one chassis. Excluding more special cases, since almost the whole system is built with data capacity of 40TB or less, two 3U chassis can fit the bill.

DWH performance: 4 times the price performance*



* Express5800/R120F-1E (2 built-in Xeon E5-2640 v3 [8-core]) x 16 vs DX2000 (Xeon D-1541 [8-core]) x 14 server modules

3.7 Consolidation of SSL Terminal Processing Servers

• Issues with SSL terminal processing servers

From the viewpoint of strengthening Internet security and compliance, the need to convert all internal website communication to SSL and increase the service level is to all the corners. Aside from that, there are also cases of constantly considering switching to SSL in order to raise the rankings of their own web page in search results, cases of implementing simple wire design by converting all Internet connections to HTTPS, and cases of considering switching to SSL as a security function "App Transport Security" measure introduced from iOS9.

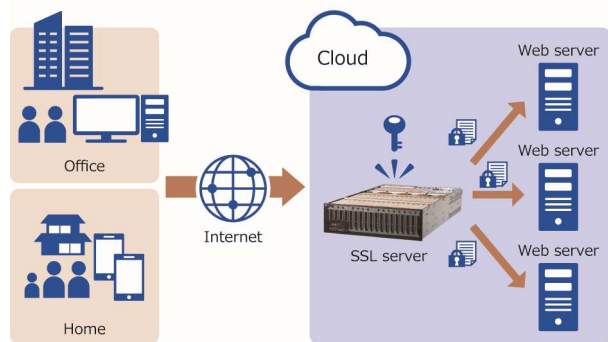
However, if companies and service providers switch to SSL on a steady basis, a large number of SSL terminal servers are needed, so securing installation space for servers and increasing management burdens is an issue.

• Resolving issues with DX2000

DX2000 packed with high performance can serve the requirement through efficient consolidation of SSL terminal servers. A report stated that in the proof of concept demonstration of a large web media provider, installation space was reduced by 80% and power consumption by 60%.

Moreover, scaling up server modules is possible, so gradual and smooth scalability is possible.

SSL terminal processing server system image



• Easy installation and maintenance

By using DX2000 configuration, everything from installation to maintenance is easier in an environment that uses many uniform servers such as SSL terminal servers. Specifically, the following three benefits are achieved with DX2000 architecture:

(1) Easy installation


One-touch setup from pre-arranged USB memory

(2) Easy to Scale

The setup for adding new equipment and integrating it into the system is automatically completed by simply adding server modules and chassis

(3) Easy patch application

Just apply the patch to one master server, and it can be replicated to all servers, increasing operational management efficiency



4. Summary

NEC's Scalable Modular Server DX2000 offers low latency processing which achieves high-speed analysis processing in order to provide real-time capability, efficiency, and flexibility as an integrated platform required for the big data and IoT era. Features of the DX2000 such as being both a high density and high performance architecture, scalability from small to large scales, and easy implementation that lends flexibility to scaling up are expected to be used in various applications along with the seven use cases introduced here.

NEC continues to strengthen scalable ICT platforms required in the data intensive era, and provide solutions that meet a variety of demands.

For enquiries, contact:

NEC Corporation

URL:

<http://www.nec.com/en/global/prod/dxseries/index.html>

Any product name or company mentioned herein is the brand name or registered trademark of that company. The contents mentioned in this white paper is as of September 2016.

Disclaimer

※NEC does not guarantee the correctness, usefulness, certainty, or guarantee of any other nature regarding the content provided in this white paper. The company shall not be held responsible for any damages owing to the use of the content provided in this white paper.