WHITE PAPER

SoftBank Tests Next-Generation IT Platform Based on Intel® Rack Scale Design in Pursuit of Information Revolution

Intel® Rack Scale Design for Data Center Efficiency

"Intel® Rack Scale Design is an integral piece of architecture for real-time computing at SoftBank as we accelerate our global strategy including the IoT, artificial intelligence, and smart robots."

– Mr. Amane Kito
Senior Vice President, CISO and
Head of Information Systems Unit
SoftBank Corp.

Mr. Amane Kito
Senior Vice President, CISO and
Head of Information Systems Unit

Mr. Masahiro Sekiya
Senior Director
System Infrastructure Division
Information Systems Division

Mr. Toshiyuki Sugiyama
System Infrastructure Development Section
System Infrastructure Division
Information Systems Division

Executive Summary

SoftBank Corp. is pursuing an "information revolution" that will spur the latest innovations in technology, including the Internet of Things (IoT), artificial intelligence (AI), and smart robots. The company needs the IT platforms which support its rapidly growing businesses to have agility to optimize real-time performance and high efficiency to utilize all available resources. The consolidation of its Japanese telecommunications companies is what led SoftBank to consider the Intel® Rack Scale Design, which pools resources at the rack level and allocates them via software. During workload-based testing, a development test demonstrated a potential reduction in the number of servers, a great reduction in staff workload, and a significant electrical conservation. The test result also showed that both server and network delivery times and fault response times could be significantly reduced. Having concluded that Intel® Rack Scale Design is suitable as a next-generation platform for real-time computing, SoftBank is now planning to test expansion functions and functional improvements in preparation for the launch of Intel® Rack Scale Design.

Consolidation of Japanese Telecommunications Companies Poses Challenge for IT Infrastructure Consolidation

• Testing Background

In April 2015, SoftBank Mobile, SoftBank BB, SoftBank Telecom, and Y!Mobile were consolidated into a single new company, SoftBank Corp., that handles both mobile and fixed line telecommunications. This created a new organizational structure under which the bulk of SoftBank Group's Japanese telecommunication businesses are consolidated.

One issue that comes with the structural reorganization is how to consolidate the IT platforms. While the business activities of the four old companies have merged, they retain the same IT infrastructure as before. Furthermore, the gradual acquisition of a variety of other companies has created a silo structure. Among the server, storage, and network hardware from different vendors and operating under different standards were a range of general-purpose systems. In order to give them the competitive advantage and differentiate themselves, they must adopt a common IT platform to achieve synergies between the four merged companies and improve efficiency.

Another factor behind SoftBank's desire to consolidate its IT platforms is to create an "information revolution", a policy they pursue in recognition of the important role that the IoT, AI, and smart robots will play in future growth. In the near future, network connections will be commonplace not only for mobile devices but also other devices such as industrial equipment, robots, and different type of sensors, with a forecast of 50 billion online devices in 2020. This could lead to the emergence of a steady stream of new IoT-based businesses, and in turn, the potential to change the fundamental role of these Internet "things." If AI-integrated smart robots enter widespread use, the need to enable robots to interpret human emotions by using large amounts of data will rise. Accordingly, there will be a demand for faster and more scalable computing resources to support the real-time processing and vi-

sualization of the high volume of transactions arising from this proliferation of devices.

These are some of the considerations that piqued SoftBank's interest in Intel® Rack Scale Design, a vertically integrated architecture for consolidating server, storage, and network hardware at the rack level and interconnecting them via high bandwidth Ethernet Fabric. Senior Vice President, CISO and Head of Information Systems Unit of SoftBank, Mr. Amane Kito, put it as follows.

"Intel® Rack Scale Design has the potential to dramatically improve the utilization of computing resources that in the past have not been efficiently utilized. SoftBank Group is currently working on a unified platform for real-time computing called "Chronos." We have been conducting tests on an Intel® Rack Scale Design development system because we expect it to play an important role in implementing this new platform.

"One of SoftBank Group's strengths is that we are quick to adopt advanced Information and Communication Technology(ICT), and ever since our formation we have consistently been among the first to test promising new technologies. This attitude has allowed us to take the first steps toward implementation of advanced ICT by using the only evaluation system in Japan, therefore learning about Intel® Rack Scale Design before anyone else. One of our future goals is global deployment of this system as a common global platform for SoftBank Group."

Confirming Functions Provided by Intel® Rack Scale Design and Operation under Actual Workloads

- **Testing Objectives**

The aim of testing was to affirm that the Intel® Rack Scale Design's functions can optimize resources in real time. The first objective was to confirm that the processors, memory, storage, and network could be configured on the hardware platform. The second was to verify the ability to execute the actual workload on top of the platform OS and middleware used by SoftBank. This work also included testing network coordination and the optimization of compute, storage, and other resources to assess the effectiveness of the automation essential for the next-generation platform. Mr. Masahiro Sekiya, Senior Director, System Infrastructure Division, Information Systems Division, commented that, "The overriding objective of the testing was to confirm whether or not Intel® Rack Scale Design, as a next-generation architecture, was suitable in practice for deployment as a platform, and could be put to beneficial use."

Intel® Rack Scale Design is one of the architectures used to implement the concept of software-defined infrastructure (SDI) being promoted by Intel, which abstracts the various hardware used at a data center, including processors, memory, storage, and networking, and uses software to modify its setup. Use of Intel® Rack Scale Design enables hardware configurations to be modified as necessary. Whereas resources in the past were only available after servers, storage, or other devices were added individually, with Intel® Rack Scale Design, they can be made available by software setting in much the same way as selecting a "component" from a pool of available resources organized on the rack level. Mr. Sekiya commented, "Because an architecture like Intel® Rack Scale Design is such a groundbreaking development, this testing will have considerable impact on our future IT infrastructure strategy."

Prototype Evaluation System Tested over an Effective 10-Week Period

- **Testing environment**

  The system configuration used for testing is listed below. The evaluation system was a prototype used by Intel for development, with hardware put together from off-the-shelf components.

1. Intel® Rack Scale Design rack
   - 26U rack
2. Rack was fitted with a common power supply
3. Rack was fitted with a common cooling system
4. Pod Manager (also incorporating rack management functions)
   - A small desktop PC kit (Intel® NUC) mounted at the bottom of the rack
5. Top-of-rack (ToR) Ethernet switch
6. Computing resources drawer
   - Each drawer was fitted with an Intel® Xeon or Intel® Atom™ processor module
   - Fabric switch with Intel® Ethernet Switch FM5224 silicon base
   - Each drawer and ToR switch connected via 40GbE optical cable
7. Intel® Xeon® E5 v3 processor modules and Intel® Atom™ C2000 SoC modules (each module has 12 Intel® Atom™ C2000 SoCs)
8. JBOD storage(object storage)

- **Testing dates**

From May to September 2015 (effective testing time of 10 weeks)

- **Tests**

SoftBank tested its basic performance and functions by using the web UI for resource monitoring; composing servers from the compute pool; assigning storage and networking, installing and booting the OS on the server; installing and executing the application; and testing the Intel® Rack Scale Design API. The following tests were conducted with an emphasis on verifying whether the basic concept of Intel® Rack Scale Design had been achieved.

[Test 1] Intel® Rack Scale Design pod management function using Pod Manager

Testing confirmed that Pod Manager, the control center for Intel® Rack Scale Design could detect key components, such as computing node and storage on the rack and manage them centrally. It also confirmed when computing resources were added or removed, the change was automatically detected by the system and the resource was recognized as available or unavailable.

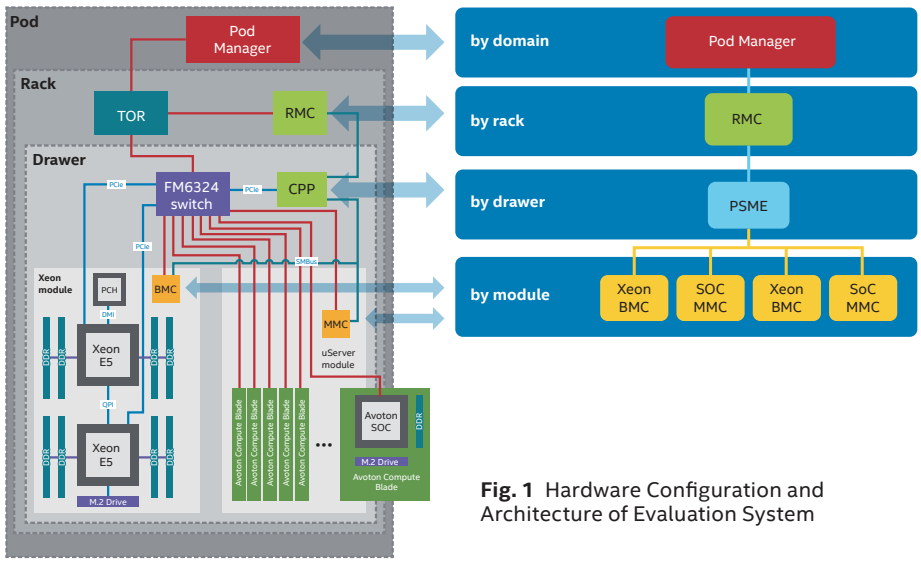

Fig. 1 Hardware Configuration and Architecture of Evaluation System

[Test 2] Assignment from resource pool using Pooled System Management Engine

Pooled System Management Engine is a management tool for drawers. Testing confirmed that it could be used to allocate resources appropriately.

[Test 3] Intra-pod storage management function and allocation through integration with distributed block storage (Ceph)

Testing confirmed that storage disk capacity could be allocated appropriately for the distributed storage (Ceph) used for testing.

[Test 4] Division and allocation of network resources

Testing confirmed that the pool of network resources on the rack could be correctly divided up and allocated to the selected nodes.

[Test 5] Efficiency of cabling using optical cable connection from networking integrated into drawer

In a typical Intel® Rack Scale Design Platform, the compute nodes within a drawer are connected to a disaggregated switch which then connects to a TOR. This allows fewer cables and prevents the tangle of cables that might otherwise be required. (Note: Based on disaggregated switch cost and performance requirements, some systems may decide to connect from compute node to TOR directly.) Testing confirmed that when a server was setup for testing, a network connection and IP address were automatically provisioned.

[Test 6] Run and test actual workloads

The same OS and middleware as used on the actual IT platforms were installed and applications run to conduct workload testing to verify that Intel® Rack Scale Design could be used in SoftBank's commercial systems. The installed OS was an Ubuntu* Linux* distribution and the middleware was a scalable key-value store (KVS) for column-based data storage. Testing confirmed the ability to run the KVS on Ubuntu* and that the workload could be executed on the KVS.

[Test 7] Intel® Rack Scale Design API testing

This involved evaluating the standard API provided with Intel® Rack Scale Design. The testing provided feedback on APIs that SoftBank would like added for use in commercial systems.

Confirmation of Basic Operation in Preparation for Fully Automated Operation in Future

- **Testing results**

1. Testing of basic functions for software-defined infrastructure (SDI)

The testing demonstrated the ability to perform provisioning using the configuration they set up, including networking, with resources in the rack being detected automatically, and the required computing resources and storage assigned by software alone using the Web UI. The testing represents a major step forward toward SDI, as Mr. Sekiya noted by saying, "What is significant is that this is not virtualization, but an actual physical, bare-metal environment."

That trouble-free execution of the workload was achieved on a configuration in which the drawers had a mix of both Intel® Xeon® and Intel® Atom™ processor modules. This demonstrated Intel® Rack Scale Design's responsiveness in permitting the reallocation of processor resources during use, even in mixed configurations that, in the future, may also include even faster programmable LSIs (FPGAs) and coprocessors (such as the Intel® Xeon Phi™ product family).

The current testing only went as far as the manual allocation of resources. In the future, the intention is to test a fully automatic mechanism for maintaining capacity without any human intervention whatsoever. The software should identify system overloading automatically and respond by adding just the right amount of additional resources, and likewise, conserve resources when they are no longer needed. This should lead to higher resource utilization and further optimization of IT infrastructure.

2. Improved operational efficiency

Testing demonstrated that Intel® Rack Scale Design significantly reduced administration workloads by pooling resources in a rack and managing them centrally using Pod Manager. Whereas adding processors or memory under previous administrative environments, for example, required servers to be purchased and added one at a time, this can be accomplished with Intel® Rack Scale Design by adding the required processor modules, memory, or storage to a computing resources drawer in the rack. This was found to be a much more efficient means of operation, and also eliminated the need to physically reconnect network cables.

Routine resource monitoring has previously involved individual monitoring of physical server resources, but resource optimization is impossible when physical servers are independent of each other, and can be influenced by workloads. Changing to the rack-based Intel® Rack Scale Design makes it possible to centrally manage the physical environment. It was also found that maintenance could be streamlined by conducting it at the component rather than the server level.

Intel® Rack Scale Design

Intel® Rack Scale Design is a reference architecture devised as part of Intel's work on its future data center strategy (software-defined infrastructure). A logical architecture that subdivides processors, memory, storage, and networking to provide highly efficient resource pools, it is designed to allocate resources automatically and dynamically in response to application demand. Use of Intel® Rack Scale Design enables the implementation of a software-controlled data center without a large investment. Specifically, the processors, memory, storage, and networking in a rack are designed for flexible allocation and pooled via a fabric switch. Rack drawers can hold multiple boards on which processors, memory, and networking are mounted. Processors are connected to the switch in the drawer, while the top-of-rack (ToR) switch (so called because it is located at the top of the rack) connects to these drawer switches allowing fewer cables.

3. Review of Testing

Testing of Intel® Rack Scale Design was a series of new discoveries. Looking back at the work, Mr. Toshiyuki Sugiyama of System Infrastructure Development Section, System Infrastructure Division, Information Systems Division commented that, "This was the first time I had ever been involved in building IT infrastructure from scratch, and while it was a big challenge to test actual operation in parallel with getting the prototype up and running, it left me with a sense of achievement." For SoftBank, this work on a new architecture also represented an opportunity to gain a major advantage.

Mr. Sekiya commented that, "The arrival of a new technology constitutes a change in the world, and success or failure is determined by how quickly you can gain experience from it. We were able to advance our IT strategy by being the first to work with a technology that has the potential to become a standard some years in the future, and to find out what it is capable of. This testing should represent a major turning point for SoftBank. It seems significant to me that this new technology was offered not by a manufacturer or system integrator but by Intel, a device vendor."

Confirmation of Cost and Workload Reductions

- **Remarks**

The test results demonstrated benefits in terms of both cost and administration. Results indicated a potential reduction of servers compared to past architectures. Along with this, the company believes it will be possible to expect a significant conservation of electric power and great reduction in operational staff workload. In terms of administration, tests indicated that required resources could be provided more quickly, with greatly shortened delivery times of servers and networking, while dealing with faults promptly.

Mr. Kito commented on these results by saying, "As the scale of SoftBank's business is expanding to encompass the IoT, AI, and smart robots, the ability of Intel® Rack Scale Design to implement automation with high performance and achieve a major reduction in IT costs not only facilitates faster decision making, it also has the potential to become an important platform for growth." Furthermore, IT is seen as providing the driving force for accelerating SoftBank Group's strategy of halving non-personnel costs and achieving a two-fold increase in productivity.

Testing of Expansion Functions and Functional Improvements in Preparation for Commercial Deployment of Intel® Rack Scale Design Systems

- **Future Activities**

As the testing was conducted using a pre-release prototype, it did not extend to considerations of actual operation. However, based on the test results, the intention is to use a commercial Intel® Rack Scale Design system, once available, to conduct further testing on expansion functions for use in commercial operation, integration with OpenStack, and testing of functional improvements. The aim of testing on a commercial Intel® Rack Scale Design system will be to execute the applications used by the business, not just the workload execution testing conducted this time. Furthermore, with a view to utilizing future new technologies released by Intel, Mr. Kito made the point that, "A world in which machine learning, image recognition, and voice recognition are performed in real time will need processors, I/O, and other technology with higher performance. Accordingly, we anticipate the integration of Intel® Rack Scale Design with new types of processors, such as FPGAs and co-processors, and also advanced memory and storage technologies that use next-generation non-volatile memory. We also recognize its potential as a next-generation IT platform that will support the global strategy of the SoftBank Group, including in operational terms and through the testing of the latest technology."

"For SoftBank, Intel® Rack Scale Design is a potential opportunity for changing the world."

– Mr. Masahiro Sekiya
Senior Director,
System Infrastructure Division
Information Systems Division

Expectation for Intel® Rack Scale Design

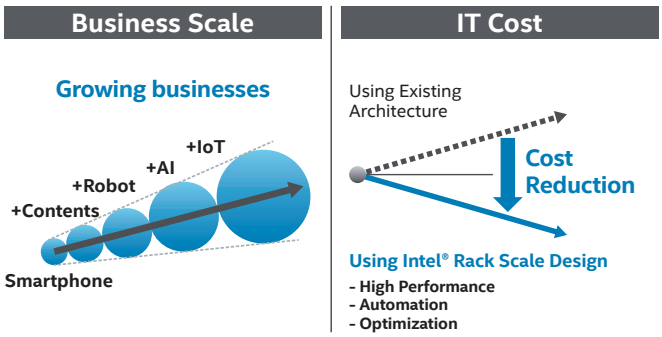

Fig. 2 (Anticipated) Benefits of Adoption
Source: SoftBank Group Corp.

Mr. Amane Kito
Senior Vice President, CISO and
Head of Information Systems Unit

Mr. Masahiro Sekiya
Senior Director,
System Infrastructure Division
Information Systems Division

Mr. Toshiyuki Sugiyama
System Infrastructure Development Section
System Infrastructure Division
Information Systems Division

Please note that the attribution and position are confirmed at the time of customer interview.

JPN/1609/PDF/SE/BDG/YS 334307-002US