FUJITSU

FUJITSU Server PRIMEQUEST 2800B

~Fujitsu's Server Optimized for In-Memory Computing~

Issue No.3
July
2014

CONTENTS

Featuring research from Gartner

shaping tomorrow with you

Predicts 2014: In-Memory Computing Will Be Adopted to Deliver High-Impact Business Value

Favored by maturing enabling software and main-memory rich commodity hardware, in-memory computing will deliver real business value to IT leaders by drastically improving core business processes. However, adoption will be incremental and will focus primarily on the most business-critical use cases.

Published: 25 November 2013
Analyst(s): Massimo Pezzini, Valdis Filks, Nigel Rayner, Errol Rasit, Andrew Norwood, Joseph Unsworth

Key Findings

>>In-memory computing (IMC) will finally enable business users to experience the ERP promises of more-efficient allocation of resources and more-effective decision making through real-time advanced analytics coupled with a new wave of process innovation.

>>Most IMC-enabled applications will be deployed on low-cost, commodity x86 hardware platforms. Less than one in five IMC applications will require adoption of specialized hardware platforms designed to support IMC-enabled software.

>>Only a minimal fraction of the data managed in a typical data center will be permanently stored in memory, but will be used by the most business-valuable applications.

>>Less than one-third of IMC-enabled applications will be the outcome of new IMC-native application development projects. In most cases, IMC technology will be leveraged to improve performance, scalability and usability of established applications.

Recommendations

During the next three years, IT leaders should:

>>Work with business users to identify and prioritize the potential benefits of accelerating ERP processes via the use of IMC technologies.

>>Plan for the predominant use of standard x86-based hardware platforms to support IMC initiatives, and look at specialized systems only for selected applications.

>>Adopt IMC technology for high-value applications that cannot be cost-effectively implemented using conventional, noninvasive accelerators, such as solid-state drives (SSDs).

>>Assess which applications would benefit from use of IMC technologies, which could support business transformation via a native IMC design and which would not benefit from IMC at all.

Table of Contents

Strategic Planning Assumptions

By 2018, at least 50% of Global 2000 companies will use in-memory computing to deliver significant additional benefits from investments in ERP.

By 2016, more than 80% of IMC-enabled applications will be deployed on industry-standard x86 platforms, thus favoring IMC adoption by mainstream organizations.

By 2017, less than 1% of data center application data and storage capacity will be processed and permanently stored within an in-memory data management platform.

Through 2016, more than 65% of production IMC applications will be the result of re-engineering projects, thus mostly delivering incremental-only benefits.

Analysis

What You Need to Know

During the past 12 months, the business benefits of IMC have started to be expressed in more concrete terms as adoption has expanded beyond leading-edge users to more mainstream organizations. For the most part, user organizations have leveraged IMC technologies to provide incremental performance, scalability and usability of well-established applications, especially in business intelligence and analytics, planning and forecasting, e-commerce, and other mainstream processes. Some organizations have experienced the profound impact that IMC can have on enabling transformational business initiatives.

Curiosity around IMC has grown exponentially, because a significant number of packaged application vendors are strategically using IMC in their software and SaaS offerings. Although the number of production deployments of IMC-enabled versions of core business processes — such as ERP, CRM, human capital management (HCM) and supply chain management — is still relatively low, many user organizations and industry players have a better understanding of when IMC can deliver tangible business benefits and when IMC is only nice to have.

Other factors favoring IMC adoption across geographies and vertical sectors are:

>>IMC-enabling application infrastructure software continues to mature.

>>Software megavendors (IBM, Microsoft, Oracle and SAP) have released, or at least announced, notable new IMC-enabling products, thus validating the IMC concept.

>>Enterprises are realizing that, in most cases, IMC can be deployed on inexpensive and powerful commodity hardware.

However, adoption continues to be primarily driven by pragmatic benefits analysis and well-thought-out business cases.

The predictions in this research will help CIOs, CTOs, enterprise technology architects, application managers, information managers and other IT leaders understand what impacts IMC could have on their organizations' businesses, and how to incorporate this new approach into their IT strategies.

Strategic Planning Assumptions

Strategic Planning Assumption:
By 2018, at least 50% of Global 2000 companies will use in-memory computing to deliver significant additional benefits from investments in ERP.

Analysis by:
Nigel Rayner

Key Findings:
Gartner estimates that organizations have invested approximately $1 trillion in ERP software and related services since the inception of this approach. The ERP vision promised three main benefits:

>>Process and data integration

>>Better allocation and use of enterprise resources: production, people and financial

>>Better visibility for executives into business performance, which would enable faster and more effective decision making

Despite massive investments in ERP over the last 25 years, this approach has only delivered on some of these promises. ERP does a good job (if implemented correctly) when it comes to using process and data integration to consolidate and standardize business operations, thus improving operational effectiveness and generating cost savings. But it has done a poor job in optimizing resource usage and providing a better understanding of business performance.

ERP architectures based on a traditional relational DBMS work well for processing high volumes of transactions, but have significant limitations:

>>Complex analytic and resource allocation functions (such as profitability analysis or material requirements planning) typically take many hours to execute. This means they effectively operate as a batch process, needing to be run in quiet periods when system usage is low (often overnight) and do not enable business users to respond in real time to changing business conditions.

>>Running sophisticated planning and simulation models is not practical due to the impact on online system performance. Consequently, for the last 20 years, organizations have had to spend more money on stand-alone business intelligence and analytics capabilities to perform the resource planning part of ERP, which creates latency and data reconciliation challenges plus additional running costs and IT landscape complexity.

IMC can remove these limitations. Long-running, complex processes can be performed in minutes instead of hours in ERP systems that are architected (or rearchitected from a traditional architecture) to use IMC data management technologies to store and process the complete database of record. Alternatively, some ERP systems may use IMC in a hybrid mode where "sidecar" applications (that is, applications complementing ERP systems with additional, noncore functionality) are used to accelerate certain application processes or provide additional analytic capabilities. In hybrid deployments, a subset of the database of record will be replicated from the on-disk database to the IMC data store. IMC can potentially also allow ERP systems to employ a hybrid transaction/analytic processing (HTAP) architecture to deliver advanced analytics against the in-memory data. Users will be able to run complex planning and simulation models in real time against the live transaction data.

ERP systems that leverage IMC will enable organizations to finally realize the full potential of their ERP investments through better analytics and a new wave of process innovation. With the combination of performance, scalability and real-time analytics, organizations will be able to use ERP systems in ways that were not previously possible. Organizations will be able to deliver a new wave of process innovation enabled by ERP, but will have to rethink and change the established ways of working to realize this potential.

Initially, IMC will deliver the most significant benefits to large organizations, because the size of their ERP systems and the complexity of their business environments significantly expose them to the limitations of traditional ERP architectures, more than small or midsize businesses (SMBs), for which these limitations usually have a lesser impact.

Market Implications:
Several vendors have used IMC to varying degrees in packaged analytic applications such as supply chain management (for example, demand planning and replenishment planning) and corporate performance management (see "In-Memory and Mobile Computing Facilitate More Agile Strategic Financial Planning"). However, IMC is not a capability that can simply be plugged into current ERP systems. Workday is one of the few vendors to have incorporated IMC in its architecture from the outset. Most other vendors are unlikely to create a new ERP application specifically to leverage IMC. Instead, they will need to retrofit IMC to an existing application or develop IMC sidecar applications for specific process areas. Consequently, vendor R&D programs and product road maps will need to accommodate IMC. Vendors that do not have a vision for IMC, especially those targeting Global 2000 organizations, will find themselves at a competitive disadvantage.

Early deployments of IMC in ERP applications often are using IMC to accelerate computing-intensive queries and existing application processes. However, they increasingly will exploit the benefits of speed and HTAP together to reimagine innovative new functionality. This will require the development of new advanced analytic applications, which are likely to come from ERP vendors and their partners.

Recommendations:
For application leaders:
- Identify ERP process areas that are performance bottlenecks and could benefit from IMC; work with business users to prioritize the potential benefits.
- Identify planning functions in which ERP falls short and is supported by specialist applications (for example, sales and operations planning, financial planning and budgeting, profitability modeling and optimization). Work with business leaders to identify where complex analytics and

simulation performed in real time on ERP transaction data could deliver significant business value.
- Review vendors' road maps to assess their plans for IMC. Identify what your ERP vendors offer now and what they plan to offer in the short term and midterm, and clarify whether this functionality will be covered under current software maintenance or subscription agreements.
- Consider becoming an early adopter if you identify short-term business value from IMC.

For CIOs and ERP competency center managers:
- Identify the impact on internal ERP support skills. IMC will shift the focus away from traditional database administration skills to analytics.

Related Research:
"In-Memory Computing Will Significantly Impact ERP Architectures and Investment Plans"

"Taxonomy, Definitions and Vendor Landscape for In-Memory Computing Technologies"

"In-Memory and Mobile Computing Facilitate More Agile Strategic Financial Planning"

Strategic Planning Assumption:
By 2016, more than 80% of IMC-enabled applications will be deployed on industry-standard x86 platforms, thus favoring IMC adoption by mainstream organizations.

Analysis by:
Errol Rasit

Key Findings:
The growth in IMC has been supported by the advances in x86 memory addressability, the ongoing decline in cost of DRAM, and user organizations' recognition of x86 architecture as a credible, mission-critical platform.

IMC-enabling system platforms come in a wide variety of styles that include single-node, scale-out and scale-up architectures. Current platforms use a variety of x86 and RISC-based CPU architectures.

The x86 architectures remain a de facto platform of choice for most IMC software providers due to their popularity, strong ecosystem and large installed base. Consequently, x86 architectures are likely to remain the most popular deployment option for IMC workloads.

Nevertheless, IMC system platforms based on non-x86 technology will still represent a viable platform that can offer advantages to IT organizations and will benefit from ongoing vendor investment. For example, advantages of non-x86 platforms for IMC include availability of larger RISC-based symmetric multiprocessor systems, acceleration via graphics processing unit (GPU), availability of larger out-of-the-box main memory and hardware/software optimization.

Market Implications:
A wide variety of IMC-enabling system platforms is available, and the breadth of offerings is likely to increase as new IMC system vendors come to market. Enterprises need to navigate through a broad range of IMC system platform styles, such as scale-out and scale-up architectures, x86, RISC and GPU-based platforms, and platforms packaged as appliances with preintegrated IMC software or as systems purchased independently of IMC software.

To contextualize today's purchase decisions, IT enterprises need to understand potential shifts in IMC-enabling system platforms. Ultimately, the breadth of the IMC-enabling system platform choice is driven by three major forces: customer pull, technology relevance and IMC software vendors' system platform certification.

There are a number of reasons why IMC customers will embrace industry-standard x86-based system platforms:

>>The x86 ecosystem has never been stronger, due to the largest number of supporting vendors, software development community, and related industry and enterprises' understanding and skills. In 1H12, x86 systems accounted for 77% of all server spending globally.

>>Hardware-neutral IMC software vendors usually support x86 as standard.

>>System and technology development for the x86 platform continues to improve at a sustained pace. Moore's Law regarding transistor density equating to performance improvements is still in effect.

>>Industry-standard building blocks of two- and four-socket racks and blade-based servers are combined into customized IMC system platforms and appliances.

>>Many customers want to streamline their IMC investments to existing major x86 system platform investments in the data center.

However, there are a number of reasons why IMC customers, at times, will embrace non-x86 system platforms:

>>There are significant legacy and proof points of SPARC and Power-based architectures for mission-critical application deployment.

>>Oracle and IBM's plans for their respective IMC software and IMC software acceleration include SPARC and Power support.

>>There is still significant investment and a scheduled road map of product releases for RISC-based symmetric multiprocessing (SMP) platforms, despite a waning broader market demand. IMC investment represents an almost singular positive element of the broader RISC market.

>>Emerging non-x86 extreme-low-energy system on a chip architectures are growing in acceptance and the community is showing growth.

>>The nascent demand for GPU acceleration capability for IMC is not based on x86.

Despite credible drivers for both x86 and non-x86 system platforms to deploy IMC architectures, due to the ongoing strength of the x86 ecosystem, a significant displacement factor across the server industry would need to occur to shift the x86 platform's dominance. Thus, we expect more than 80% of IMC deployments will be based on x86 industry-standard system architecture. However, the IMC market does not preclude non-x86 system platform investment, as the business value of the IMC platform to the enterprise can bypass considerations about underlying hardware architecture. As long as there are IMC solutions based on non-x86 IMC hardware, IT enterprises will invest in them, provided they can yield business value.

Recommendations:
For application leaders and other IT leaders focusing on developing and deploying IMC-enabled applications:

- Prioritize IMC purchase decisions on the business value of the solution, which should be driven primarily by the IMC software's capability. Considerations regarding the underlying server technology are of secondary relevance.
- Consider the underlying system architecture, but focus on the cost of an upgrade and capacity expansion as key measures of TCO, which includes software, labor, professional services and other factors.
- Plan for IMC investments that will be predominantly x86-based, where the IMC software vendor is hardware-neutral; but expect non-x86 architectures to still play a part, such as for non-x86 GPU acceleration, or when an IMC software vendor works closely with a non-x86 hardware provider.

Related Research:
"IT Market Clock for Server Technology, 2013"

"Taxonomy, Definitions and Vendor Landscape for In-Memory Computing Technologies"

"Hype Cycle for In-Memory Computing Technology, 2013"

Strategic Planning Assumption:
By 2017, less than 1% of data center application data and storage capacity will be processed and permanently stored within an in-memory data management platform.

Analysis by:
Valdis Filks

Key Findings:
The capacity of IMC application systems cannot be larger than the amount of server RAM produced within the market.

The ratio of worldwide enterprise class hard-disk drive (HDD) storage to RAM capacity deployed will remain approximately 99-to-1.

Even assuming the most optimistic forecast, increases in SSD adoption will only change the planning assumption of IMC to HDD system ratios by 1% to 2%.

Any data reduction techniques that enable more data to be stored in RAM or IMC systems are negated by similar and proportional improvements in HDD data reduction ratios. These are expected to move in step with each other.

Market Implications:
Do not expect or plan that IMC will become a ubiquitous standard architecture within the next five years, although this may happen in selected verticals or for specific use cases (for example, financial trading, supply chain planning, e-commerce and online entertainment).

Any increase in IMC application demand will increase the cost of RAM and, therefore, increase the cost of IMC system platforms, creating a self-balancing system that maintains IMC as a high-cost

and high-value market. A relatively small proportion of applications and services require the performance increases provided by IMC. All users want faster performance, but many do not require it. By moving data from HDD to SSD, you can achieve significant latency reduction performance improvements at a far less cost than implementing IMC. The storage latency performance improvement of 100 times achieved by moving data to SSD from HDD is often good enough, compared with the 100,000 times improvement of moving to IMC from HDD.

Recommendations:
For application architects and project leaders:
- Use IMC for high-value applications that cannot be cost-effectively implemented using conventional technologies.
- Qualify the payoff of all proposed IMC applications, and manage organizational or business expectations.
- Consider using SSD instead of IMC technology if performance improvement is the main concern, IMC costs cannot be justified and applications are storage bound or limited.

Related Research:
"Solid-State Drives Will Complement, Not Replace, Hard-Disk Drives in Data Centers"

"Forecast: Servers by Form Factor, Worldwide, 2Q13 Update"

"Forecast: DRAM Supply and Demand, Worldwide, 1Q12-4Q14, 2Q13 Update"

Strategic Planning Assumption:
Through 2016, more than 65% of production IMC applications will be the result of re-engineering projects, thus mostly delivering incremental-only benefits.

Analysis by:
Massimo Pezzini

Key Findings:
IMC technologies and architectures can be applied to application development projects in multiple ways, each characterized by different design and implementation efforts, but also driving to potential differentiated business and IT benefits. A so called native IMC application is based on IMC design principles; keeps all the application data in-memory; and is implemented by leveraging IMC technologies such as in-memory DBMS (IMDBMS), in-memory data grids (IMDGs), in-memory application servers (IMAS) or complex-event processing (CEP).

Developing native IMC applications is a major design and development effort that only a few leading-edge enterprises dare to tackle, given the lack of industry experience, skills and best practices. However, they can deliver the full IMC benefits in terms of faster performance, greater business scalability and deeper business insights.

Some of these benefits can be achieved by re-engineering traditional applications to take advantage of IMC technologies. For example, applications developed on top of traditional relational DBMSs can be re-hosted on top of an SQL IMDBMS, or they can be partially redesigned to adopt IMC design principles and technologies. For example, they can improve performance and/or scalability by storing and processing in-memory for only the most performance-sensitive application data or by adding new business functions implemented on top of IMC technologies.

The effort required to re-engineer traditional applications is lower than required for implementing native IMC, but only native IMC applications are likely to have a transformational impact. In most cases, re-engineering for IMC can improve the overall performance and scalability of established applications, but the benefits will vary. Some resource-intensive business processes may require native development on IMC to fully leverage the technology. For example, it will be difficult to implement advanced analytics and create process innovation through re-engineering an application. Therefore, re-engineering for IMC will provide various benefits – for example, obtain business intelligence reports earlier, calculate profitability faster, process more transactions per second, or reduce the cost of supporting mobile applications. In some industries, these benefits may enable business transformation; but, in most cases, these outcomes will only deliver incremental business value.

Nevertheless, mainstream application developers will adopt a re-engineering-for-IMC approach faster and more frequently than IMC design models. IMC re-engineering of established applications not only is less complex than redesigning it according to a native IMC approach, but also can deliver business results in few months, as opposed to the many months or years it may take to implement an IMC-native application. Re-engineering is relatively low risk: If the project fails, it is possible to go back to the original application design and look for some incremental performance benefit by utilizing SSD technology for the storage layer.

The majority of established packaged application vendors will follow this strategy to reinvigorate their products and establish competitive differentiation by delivering minimally disruptive innovation to their installed base. They will try to reduce the appeal of radically innovative, native-IMC products from competitors by re-hosting their products on IMC platforms. At the same time, they will have an opportunity to incrementally deliver native IMC add-on, side car applications. Vendors will leverage these applications to increase revenue, as well as to "test the waters" and build skills in sight of a potential major redevelopment of their core products according to a native IMC model.

For the same reasons, application infrastructure vendors providing analytics platforms, business process management (BPM), enterprise service bus suites, portal products, application servers and other products will prefer to re-engineer their established offerings for IMC, rather than redesigning them according to a native model.

The net result of these trends will be that during the next two to three years the majority of IMC applications that user organizations deploy into production will result from re-engineering established applications or from the adoption of re-engineered-for-IMC versions of the underlying application infrastructure (e.g., application servers, portal products, business intelligence tools or BPM tools).

User organizations and technology providers that have the courage to tackle the challenges of native IMC development projects will have higher chances of establishing a sustainable competitive advantage than those trying to limit risks by adopting only a re-engineering approach.

Market Implications:
To support and favor a re-engineering approach, IMC technology providers will strive to provide compatibility layers on top of their products by backing popular APIs and languages (for example, memcached, Ehcache, SQL, Java Message Service [JMS], REST and OData) to make these products' adoption by mainstream developers as smooth as possible in the context of established application environments.

Vendors of traditional DBMS technologies will incrementally add IMC capabilities to their products. They will try to maintain as close to 100% compatibility as possible with their established programming interfaces and development tools to prevent pure-play IMDBMS vendors from eroding their installed base.

Providers of application infrastructure products will increasingly update their technologies to leverage IMC. For example, vendors of business intelligence and analytics tools, big data platforms, development environments, data integration platforms and data warehousing products will certify their products for the most popular IMDBMS technologies. This approach will also be followed by vendors of BPM, enterprise service buses, application servers, portal products and other middleware technologies, including their cloud renditions in the form of platform as a service (PaaS) offerings. These vendors will re-engineer their products to leverage combinations of IMDBMs, IMDGs, high-performance messaging infrastructures and CEP platforms to improve scalability and performance, to support public/private cloud deployment models, and to deliver greater technical and business operational insights to applications built on these products.

Packaged application vendors will increasingly certify their products on top of IMC-enabled versions of the DBMSs they support. They will leverage these and other IMC technologies to provide valuable add-on components to their core products (for example, real-time predictive analytics and advanced planning tools). These innovations may be delivered to user organizations as cloud services (for example, SaaS).

The availability of such a wide array of IMC-enabled products and cloud services will accelerate IMC penetration, even in risk-averse mainstream enterprises. It will catalyze system integrators' investments and foster the building of IMC skills, ultimately helping reduce the costs and risks of IMC adoption.

Some of these re-engineering efforts may be driven by the erroneous assumption that IMC technologies can deliver dramatic benefits independently from the application architecture. Independent software vendors (ISVs) and user organizations' practical experience instead shows that, in general, some reworking of the application is required to fully experience IMC performance and scalability benefits. Organizations may find that failing to include a significant code optimization effort in the re-engineering project may lead to disappointing results in terms of tangible benefits for users, thus risking to undermine IMC credibility among business leaders.

Recommendations:
For application leaders:
- Assess which parts of the application portfolio would benefit from an IMC re-engineering

approach, which could deliver transformational business outcomes through a native IMC architecture redesign, and which would not benefit from IMC at all.
- Identify a relatively easy-to-implement opportunity to provide value to the business (for example, providing important business intelligence reports earlier) to run a proof-of-concept project. This will help them assess the efforts, costs and benefits of re-engineering for IMC their organizations' application portfolio and application infrastructure.
- Understand the IMC-adoption road map of their packaged business application and application infrastructure vendors, and assess which elements of these strategies can support an IMC re-engineering approach and which, instead, are aimed at native IMC developments.
- Define a midterm (two to three years) plan to leverage IMC re-engineering opportunities that can deliver short-term business or technical benefits (e.g., faster execution of key business processes, greater business scalability, deeper business insights, more widespread user acceptance of analytics applications, back-end system offloading and cost containment) either by re-engineering custom applications (for example, adopting an IMC-enabled application infrastructure) or by adopting IMC-enabled versions of traditional packaged applications.
- Use these IMC re-engineering projects as opportunities to understand IMC technologies and to build skills as needed to address with greater confidence future game-changing native IMC projects.

Related Research:
"SAP's Business Suite on Hana Will Significantly Impact SAP Users"

"Application Architecture Next Practices: Lessons Learned From Avanza Bank's Visionary In-Memory Architecture"

"In-Memory Data Grids Enable Global-Class Web and Mobile Applications (While Not Bankrupting Your Company)"

A Look Back

In response to your requests, we are taking a look back at some key predictions from previous years. We have intentionally selected predictions from opposite ends of the scale – one where we were wholly or largely on target, as well as one we missed. While we have noted a Prediction from last year in which we are largely on target, we were unable to uncover a missed prediction.

On Target: 2012 Prediction
By 2015, lowering memory costs will halve current in-memory database capacity price/performance, and enable a three-times increase in IMC adoption.

Analysis by:
Joseph Unsworth, Andrew Norwood

Previously Published:
"Predicts 2013: In-Memory Computing: Growing Gains, but Also Growing Pains"

Gartner preliminary estimates forecast a doubling in total IMC software revenue from 2012 to 2015. Due to several factors that will pressure software pricing, including intensifying competition among software megavendors, this doubling in revenue will correspond to a much faster growth in terms of adoption. Other factors that will drive adoption, while having minimal impact in terms of IMC software vendors' revenue, include IMC software being embedded in other software products and cloud services, and the maturation and growing popularity of open-source IMC software technology. All these trends confirm the industry is on track to tripling IMC adoption between 2012 and 2015.

While in 2013 the DRAM market has seen a pricing spike as we had forecast, the price trend was exacerbated by a fire at one of SK Hynix's facilities, which knocked out production for several months and resulted in taking roughly 7% of worldwide DRAM production capacity offline. The additional surge in prices will be ephemeral as production normalizes in 2014. Gartner still maintains DRAM average selling price (ASP) per GB falling roughly 45% between 2012 and 2015.

NAND Flash memory, which is used in SSDs, has also witnessed expected supply constraints that, when combined with limited cost reductions in 2013, have caused limited price declines for SSD. However, the expectation that Flash memory pricing, and thus SSD prices, will halve by 2015 is still on track. As new production capacity mounts, technological advancements promote the adoption of advanced flash technology to be used in SSDs for IMC environments. Competition has become especially fierce, with more than a dozen acquisitions in the SSD component and SSD appliance industries, which will foster a particularly intense battleground and result in price erosion.

Datasheet
FUJITSU Server PRIMEQUEST 2800B Business Critical

FUJITSU Server PRIMEQUEST 2800B provides high-end server functionality using superior Fujitsu technology, long cultivated and refined over generations of computer system development.

The most cost-efficient enterprise server

FUJITSU Server PRIMEQUEST 2800B is a high-end data center system focused on the needs of the growing enterprise. Based on technologies and innovations Fujitsu has refined over generations of highly reliable mainframe and UNIX servers, it provides high-performance, excellent service availability and the openness of x86 servers.

Global standard Linux® and Microsoft® Windows® operating systems, with highly advanced Fujitsu reliability, stability and manageability technologies, make PRIMEQUEST 2800B a highly cost effective platform.

Customer investment is fully secured by:

- Ability to operate a 24-hour, 365-day business.
- Outstanding performance and reliability.
- Excellent flexibility and scalability in an open server.
- A radically improved cost/performance profile.
- Use optimization and scalability for the future.

Main features	Benefits
Complete redundancy and online recovery function	
■ Almost all components can be redundantly configured	■ Systems on PRIMEQUEST 2800B can continue operation even if a component fails
Highest data throughput in mission critical arena	
■ 2.5 times the performance of predecessor model PRIMERGY RX900 S2 ■ 4 times the performance of PRIMERGY RX600 S6 ■ Highest level of database performance ■ High performance for computation, memory access, and I/O access ■ High data throughput even for very large data	■ PRIMEQUEST 2800B with much superior performance scalability can accommodate higher workloads ■ PRIMEQUEST 2800B can satisfy customers demand to expand business platform without disruption
Rich variety of I/O products	
■ Solid State Drive products for disk drives and flash memory connected through PCIe are available to accomodate customers demands such as cost efficiency and high performance ■ InfiniBand helps scale out database system without sacrificing performance	■ For demands for cost-efficiency and performance upgrade, SSD disk drives can be the best solution. For demands for high performance upgrade, PCIe SSD can be the best solution ■ PRIMQUEST with Infiniband can scale up Oracle Database unlimitedly
Protect confidential data	
■ Even if confidential data is stolen, the stolen data cannot be used maliciously	■ Intel® Xeon®'s on-chip encryption greatly reduces encryption/decryption time making database encryption a practical solution for protecting intellectual property
Much lower operational costs	
■ PRIMEQUEST 2800B can halve reduce power consumption compared to predecessor model at similar performance ■ With high performance Intel® Xeon® E7-8800 v2 product family and super fast I/O technology, PRIMEQUEST provides high performance with small number processor core	■ As server resource utilization changes, PRIMEQUEST 2800B power consumption stays low. In particular it only consumes power according to its application workload ■ PRIMEQUEST help reduce costs for datacenter space ■ PRIMEQUEST can save costs for Oracle DB license and support charge compared to other servers at the similar performance.
Simplified server lifecycle management	
■ An integrated suite of tools takes care of servers and their component products in your datacenter over the entire life of the server	■ Human resource costs for server management, including: installation, integration, monitoring, maintenance, and upgrading are reduced

Platform of standards and high availability

With outstanding redundancy, PRIMEQUEST 2800B provides the high uptime required from true business platforms. Under the strict quality assurance standards in Fujitsu's production processes - from server design to manufacturing and support – PRIMEQUEST can dramatically reduce failure rates.

Using Intel® Xeon® E7-8800 v2 product family and Fujitsu's cutting-edge I/O technologies, PRIMEQUEST 2800B has hit highest performance scores in the large systems server arena. This is further evidence that Fujitsu can deliver best-matched products, conforming to open systems standards that handle business workloads. For instance, PRIMEQUEST 2800B combined with variety of Solid State Drive products can leap data throughput from what existing servers could provide.
Based on state-of-the art technologies such as smart cooling, Intel® processors and other components PRIMEQUEST 2800B provides the best cost-efficiency in the enterprise server arena. It also reduces data center costs by lowering power consumption and minimizing the server footprint.

PRIMEQUEST 2800B with its high availability, performance scalability, cost-efficiency, and risk-free virtualization can lift the Return on Investment (ROI) of your IT system.

Topics

Almost every component redundant

With PRIMEQUEST 2800B, your most important business operations can be strictly protected from errors:

- Memory can be mirrored. This means that even if a memory failure occurs the application will continue using the correct data. With Double Data Device Correction (DDDC), even if two memory chips fail, the system can continue operation without interruption
- System Interconnect, called Quick Path Interconnect (QPI), provides multiple access routes. This ensures continued operation even if one route fails
- Fans, PCI Express switches, PCI Express cards, and Ethernet ports, are redundant as standard
- HDD can be configured redundant using hardware or software RAID
- Management Boards (MMB) and Power Supply Units (PSU) can be optionally configured as redundant

Only grow as you need

With up to 120 cores and maximum 12TB of memory, PRIMEQUEST 2800B has the resources to accommodate hundreds of workloads. Performance has doubled ensuring PRIMEQUEST 2800B capability as an enterprise workload platform continues to expand.

Performance improvement of PRIMEQUEST 2800B

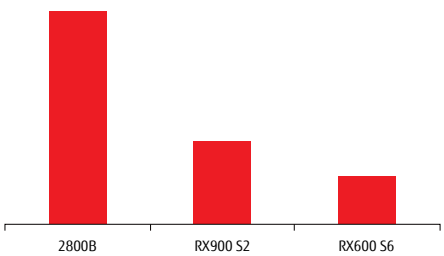

Protect confidential data

PRIMEQUEST 2800B embeds security measures into its hardware and OS. Due to swift encryption by Advanced Encryption Standard New Instructions (AES-NI) with the Intel® Xeon® E7-8870 v2 processor family, a hacker will not maliciously use any stolen data. As an example, Oracle Database 12 and Xeon® processors shrink the encryption time to one-tenth.

Much lower operational costs

PRIMEQUEST 2800B can reduce operational costs : power consumption, datacenter space, and Oracle license and support changes,

- This halves the power consumption charge compared to predecessor model at the similar performance.
- Oracle database license and support charges are halved compared to predecessor model PRIMERGY RX900 S2.

Simplified server lifecycle management

During a server's life cycle you must undertake a variety of actions including installation, integration, monitoring, maintenance, and upgrading. To do this you have to use different tools for different actions. It can be a nightmare. Fully integrated tools through the lifecycle are what our customers are demanding.

Fujitsu provides an integrated suite of tools that take care of server products at your datacenter for the entire life of the server. ServerView Suite, a bundled product with PRIMEQUEST and PRIMERGY, can help ease the pain in dealing with servers. This includes:

- Automated OS installation on multiple servers
- Automated RAID configuration
- Automated driver updates, hot fixes and security patch applications
- Integrated monitoring of multiple Fujitsu PRIMEQUEST and Fujitsu PRIMERGY servers
- Simplified setup and monitoring of disk array controllers, HDD, and logical drives

The suite also enables early problem detection and resolution via intuitive diagnostics, look-and-feel operation and pro-active error alerts.

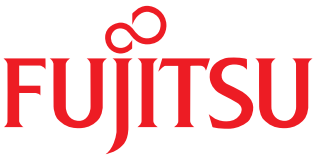

Contact