

Lawrence Berkeley National Laboratory

Recent Work

Title

Comparing Server Energy Use and Efficiency Using Small Sample Sizes

Permalink

<https://escholarship.org/uc/item/2zc334th>

Authors

Coles, Henry C.

Qin, Yong

Price, Phillip N.

Publication Date

2014-11-01



ERNEST ORLANDO LAWRENCE BERKELEY NATIONAL LABORATORY

Comparing Server Energy Use and Efficiency Using Small Sample Sizes

Coles, Henry C., Yong Qin, Phillip N. Price

Environmental Energy Technologies Division
November, 2014



DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

ACKNOWLEDGEMENTS

This work was supported by the California Energy Commission's Public Interest Energy Research Program and by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Program, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Thanks to key partners Intel, Super Micro Computer and Dell Inc. Intel loaned the project three complete servers. Specific appreciation goes to Henry Wong, Dave Hill, and David P. Larson at Intel and Tony Lai and Nihir Parikh at Supermicro. Tim Bock from Dell provided technical guidance.

This project would not have been possible without the tireless and generous support provided by colleagues working in Lawrence Berkeley National Laboratory's High Performance Computing Services group, headed by Gary Jung. A special thank you goes to Susan James for her support of IT equipment, network configuration, and testing during the project. Additional thanks to Karen Fernsler and Bernard Li, and many thanks to Mike Magstadt, who supported the effort by allowing us to use a Lawrence Berkeley National Laboratory data center for this project. In addition the team thanks Daniel Fuller from the Lawrence Berkeley National Laboratory Environmental Energy Technologies Division for his assistance.

The project team and the High Performance Computing Services group appreciate the support provided by Server Technology, which provided electrical power measurement equipment in the form of rack-mounted power distribution units.

PREFACE

The California Energy Commission Energy Research and Development Division supports public interest energy research and development that will help improve the quality of life in California by bringing environmentally safe, affordable, and reliable energy services and products to the marketplace.

The Energy Research and Development Division conducts public interest research, development, and demonstration (RD&D) projects to benefit California.

The Energy Research and Development Division strives to conduct the most promising public interest energy research by partnering with RD&D entities, including individuals, businesses, utilities, and public or private research institutions.

Energy Research and Development Division funding efforts are focused on the following RD&D program areas:

- Buildings End-Use Energy Efficiency
- Energy Innovations Small Grants
- Energy-Related Environmental Research
- Energy Systems Integration
- Environmentally Preferred Advanced Generation
- Industrial/Agricultural/Water End-Use Energy Efficiency
- Renewable Energy Technologies
- Transportation

Comparing Server Energy Use and Efficiency Using Small Sample Sizes is the final report for the [High Efficiency Server] project (contract number DE-AC02-05CH11231, work authorization number 500-10-052) conducted by Lawrence Berkeley National Laboratory. The information from this project contributes to Energy Research and Development Division's Buildings End-Use Energy Efficiency Program.

For more information about the Energy Research and Development Division, please visit the Energy Commission's website at www.energy.ca.gov/research/ or contact the Energy Commission at 916-327-1551.

ABSTRACT

This report documents a demonstration that compared the energy consumption and efficiency of a limited sample size of server-type IT equipment from different manufacturers by measuring power at the server power supply power cords. The results are specific to the equipment and methods used. However, it is hoped that those responsible for IT equipment selection can use the methods described to choose models that optimize energy use efficiency.

The demonstration was conducted in a data center at Lawrence Berkeley National Laboratory in Berkeley, California. It was performed with five servers of similar mechanical and electronic specifications; three from Intel and one each from Dell and Supermicro.

Server IT equipment is constructed using commodity components, server manufacturer-designed assemblies, and control systems. Server compute efficiency is constrained by the commodity component specifications and integration requirements. The design freedom, outside of the commodity component constraints, provides room for the manufacturer to offer a product with competitive efficiency that meets market needs at a compelling price.

A goal of the demonstration was to compare and quantify the server efficiency for three different brands. The *efficiency* is defined as the average compute rate (computations per unit of time) divided by the average energy consumption rate. The research team used an industry standard benchmark software package to provide a repeatable software load to obtain the compute rate and provide a variety of power consumption levels. Energy use when the servers were in an idle state (not providing computing work) were also measured.

At high server compute loads, all brands, using the same key components (processors and memory), had similar results; therefore, from these results, it could not be concluded that one brand is more efficient than the other brands. The test results show that the power consumption variability caused by the key components as a group is similar to all other components as a group. However, some differences were observed. The Supermicro server used 27 percent more power at idle compared to the other brands. The Intel server had a power supply control feature called *cold redundancy*, and the data suggest that cold redundancy can provide energy savings at low power levels.

Test and evaluation methods that might be used by others having limited resources for IT equipment evaluation are explained in the report.

Keywords: server power, compare server power, server efficiency, compare server efficiency, measuring server power, Monte Carlo simulation, server power simulation, cold redundancy, computation per watt, small sample size, idle power

Please use the following citation for this report:

Coles, Henry C., Yong Qin, Phillip N. Price (Lawrence Berkeley National Laboratory). 2014. *Comparing Server Energy Use and Efficiency Using Small Sample Sizes*. California Energy Commission.

TABLE OF CONTENTS

EXECUTIVE SUMMARY	1
Introduction	1
Project Purpose	1
Methods and Results	2
As-Received Equipment, Methods and Results	3
Components and Sub-Assemblies Swapped-Test Methods and Results	3
Conclusions	5
Benefits	5
CHAPTER 1: Introduction	8
1.1 Summary	8
1.2 Demonstration Goals	8
1.3 Participants	8
CHAPTER 2: Methods	10
2.1 Introduction	10
2.2 Goals	11
2.3 Demonstration Process	12
2.3.1 Introduction	12
2.3.2 Setup and Equipment	12
2.3.3 Data Collection	14
2.3.4 Test Description	15
2.3.5 Analysis	21
CHAPTER 3: Results	22
3.1 As-Received Power Consumption	22
3.2 Swap A: Swapping Intel Nodes	24
3.3 Swap B: Swap Key Components from One Intel Server to Another Intel Server	25
3.4 Swap C: All Brands - Using One Set of Key Components	26
3.5 Experiment Simulations	28
3.6 Other Observations	30
CHAPTER 4: Conclusions and Recommendations	33
4.1 Conclusions	33
4.2 Recommendations	34

Appendix A: BIOS Adjustment Procedure.....	A-1
Appendix B: Swap B Analysis	B-1
Appendix C: As-Received Measured Power Data	C-1
Appendix D: Test Results Data	D-1
Appendix E: Swap A Measured Power Data	E-1
Appendix F: Swap C Measured Power Data.....	F-1

LIST OF FIGURES

Figure 2-1: Nomenclature	11
Figure 2-2: Servers Installed In Demonstration Rack	13
Figure 2-3: Server Node Number (0–19) and Power Distribution Unit Plug Assignments	14
Figure 2-4: Block and Cyclic Distribution of Processes (Eight-Core Processor Used as Example)	16
Figure 2-5: Swap A: 4 Nodes from One Intel Server Tested into All Intel Servers.....	18
Figure 2-6: Swap B: Processors and Memory from One Intel Server into Another Intel Server ..	19
Figure 2-7: Swap C: Key Component Set Installed in All Brands	19
Figure 2-8: Simulation Input Data.....	21
Figure 3-1: As-Received Server Mean Power - All Brands	23
Figure 3-2: As-Received Efficiency, by Brand	24
Figure 3-3: Intel Server Power Using One Set of Nodes	25
Figure 3-4: Intel Server Power Using One Set of Nodes.....	26
Figure 3-5: Server Power Consumed for All Brands Using Same Key Components	27
Figure 3-6: Compute Efficiency for All Brands Using One Set of Key Components.....	28
Figure 3-7: Simulation Results.....	30
Figure 3-8: Intel Server Individual Power Supply Measurements	31
Figure 3-9: Intel Server Test with Different Inlet Air Temperatures	32

LIST OF TABLES

Table 2-1: As-Received Server Configurations.....	13
Table 2-2: Collected Data Used.....	15
Table D-1: As-Received Test Description and Results.....	D-1
Table D-2: Intel #1 Nodes Installed in Three Intel Server.....	D-1
Table D-3: Intel #3 Components Installed in Intel #1 Server and Nodes.....	D-2
Table D-4: Intel #1 “Gold” Components Installed in Dell and Supermicro	D-2
Table D-5: Intel #1 “Gold” Components Installed in Dell and Supermicro - Full Power Efficiency	D-2

EXECUTIVE SUMMARY

Introduction

Data centers in the United States currently consume approximately 2 percent of the nation's electrical energy (Kooimey 2011). A large part of the electrical energy consumed in data centers, often 50 percent or more, is consumed by the power distribution and cooling systems often referred to as the *infrastructure*. In recent years much of the attention on data center efficiency has focused on reducing the energy consumed by the infrastructure needed to support the IT equipment rather than reducing the IT equipment energy use. This demonstration was focused on the energy use of the IT equipment itself.

In the summer of 2013, a presentation at Lawrence Berkeley National Laboratory (LBNL) by a high-volume server manufacturer included claims that one of their server models used less energy compared to similar models made by competitors. This claim sparked interest within the High Tech and Industrial Systems group at LBNL. Following that, the California Energy Commission (Energy Commission) authorized the study.

The initial focus for this demonstration was to study the IT equipment energy use. Hereinafter we will refer to the IT equipment as a *server* or *servers*. Servers are distinguished from other IT equipment such as internet routers and storage modules. Reduced server energy use saves in two ways: through energy savings by the servers themselves and through reduced energy required to support the infrastructure for the servers. Given that the infrastructure's energy use is equal to the energy use of the IT equipment in many data centers, the total energy saved can often be twice the energy saved at the server.

While servers are offered in many shapes, sizes, and capabilities, mainstream models across different brands often have very similar mechanical and electronic architecture and configuration attributes. This similarity provided an opportunity for a fair comparison of power use and efficiency for servers from different brands. A particular set of server attributes were selected, and server models with the same general specifications from three brands were tested and compared.

This demonstration focused on determining if the energy use and efficiency was different as a function of the server brand. Other findings revealed themselves in the process. In addition methods used to test and quantify the comparison were developed and presented along with the results.

Project Purpose

The purpose of the project was to compare the performance of three models from different brands that appear to have the same capabilities and design. Mechanical and electronic architecture of commercial servers available at a point in time have many similarities. For example, models from different manufacturers may take up the same space in a server rack,

have motherboards designed for the same processor components, and have common memory slot specifications.

While results for a specific comparison may be interesting, an important goal was to develop and present methods that could be used by others for their efforts to select servers with high efficiency in mind.

To assist those interested in selecting equipment with the best efficiency, this demonstration attempted to answer the following questions using a limited sample quantity:

- How much variability in power use and efficiency was there across the three brands?
- How much variability was there from machine to machine, for machines that were supposedly “identical” (i.e., same brand and configuration)?
- What testing protocol could be used to determine the answers to the questions above?

Methods and Results

Servers from three manufacturers (Intel, Dell, and Supermicro) were tested for the purpose of comparing power consumption and compute efficiency.

When servers are not performing any work, energy use alone is a helpful metric. When computational output is expected, an efficiency metric is needed. A common industry metric—million floating point operations per second/watt (MFLOPS/W) (compute rate /electrical power)—was used in this demonstration. LINPACK tools (commonly used benchmark software), were used to load the servers at the 50 percent and 100 percent power levels and to measure performance. The computational efficiency was calculated from the LINPACK run results and power measurements. Tests at idle power level were performed with no application or test software running.

Recent processors and related components contain a number of energy-efficiency features that may be controlled by the basic input/output system (BIOS) settings. The BIOS settings used in the demonstration were selected to provide comparable results to the extent possible for all of the servers, as determined by the demonstration team in collaboration with each server manufacturer.

Two groups of tests were performed: (1) as-received tests, and (2) components and sub-assembly swap tests.

As-Received Tests: The as-received tests were performed to explore energy use and efficiency differences when servers were fitted with different components; for example, different processor models. The power use and efficiency differences found in these tests should not be used to compare brands but may provide some insight relating to variation caused by different model components.

Components and Sub-Assembly Swap Tests: By swapping components and sub-assemblies among different copies of the same server brand, and across all the server brands, more

definitive conclusions can be made regarding which components are responsible for the variation in power use and efficiency.

As-Received Equipment, Methods and Results

Equipment and Methods

- The Supermicro server was equipped with older Intel “Sandy Bridge” processors (E5-2670 v0), which have been reported to be less efficient (i.e., to perform fewer computations per unit energy or use more power at idle) than the more modern Intel “Ivy Bridge” processors.
- The Dell server used 10-core E5-2670 v2 processors, which are similar to the 8-core E5-2650 v2 processors used in the Intel servers in that both processor models use Ivy Bridge technology. Memory configurations also differed between brands.
- The servers were tested as-received from the loaning entity. Each server was a 2U server enclosure containing 4 independent nodes, with each node having 2 processors. All power measurements include all 4 nodes at idle or in operation at the same time.
- The power for each server was measured at both server power cords, while the servers were operating at three different computational load levels: idle, 50 percent, and 100 percent). The computational efficiency was calculated from the LINPACK run results and power measurements.

Results

- Idle power was about 50 percent higher for the Supermicro server than for the Intel and Dell servers. The Intel and Dell servers’ average idle power was 235 W.
- At 100 percent compute load, the Supermicro server consumed 1,382 W, or about 35 percent more power than the Intel and Dell servers. The Intel and Dell servers used 1,023 W average as a group.
- In the as-received condition, the three Intel servers used different amounts of power; Intel Server #1, #2, and #3 used 985 W, 1,037 W, and 1,019 W, respectively. This range of 53 W (985 W to 1,037 W) was approximately +/- 5 percent from the mean value and larger than expected.
- Compute efficiency was highest for the Dell server (1,365 MFLOPS/W), a bit lower for the Intel servers (1,131 MFLOPS/W), and considerably lower for the Supermicro server (871 MFLOPS/W). As noted above, the Supermicro server was equipped with older Sandy Bridge processors compared to the Dell and Intel servers, which were equipped with Ivy Bridge processors. Different processor models are the likely the key factor for the different performance levels, not the server vendor's design.

Components and Sub-Assemblies Swapped-Test Methods and Results

Methods

- **Swap A:** A complete set of four-node assemblies were swapped among the three Intel servers to explore how or if the energy use variance was due to the node assemblies or

the power supplies. A node assembly consists of a motherboard and all the components attached to it, including the processors and memory dual in-line memory module (DIMMs).

- **Swap B:** The key components (processors and memory) were swapped from one Intel server into another Intel server to determine if the energy use variance was dominated by the processors and memory or by the other components on the node board.
- **Swap C:** A single set of 8 processors and 16 DIMMs were installed in all three brands, to determine the energy use and efficiency when key components are fixed.
- Additionally, the research team developed and exercised the ability to simulate component swaps, using component power parameter estimates obtained from the experiments. These simulations can be used to gain insight into expected variability for different server configurations.

Results

- **Swap A:** The Intel servers had identically designed sub-assemblies (node boards or node assemblies), so it was possible to switch them from one Intel server to another. In the case of the Intel server, the node assembly includes fans. When one set of four-node assemblies were installed in all three Intel servers, the power varied by only 7 W. This shows that almost all of the variability between the servers was due to variable power consumption of the node assemblies (containing the processors and memory), and not variability in the power supplies (which were not switched between servers).
- **Swap B:** When a complete set of just the processors and memory from one Intel server was swapped into another Intel server, the resulting power difference was 44 W. This result was combined with the as-received Intel server results to conclude that the processors and memory components as a group compared to all other components were responsible for most of the variance in power use.
- **Swap C:** All brands used in this demonstration had mechanical architectures that allowed processors and memory to be switched between servers. When the same set of processors and memory was tested in each server brand at 100 percent compute load, the power difference had a range of 87 W. The average compute performance rate for all was within 1 percent of the mean of the three brands. We conclude that the processors and memory together have a similar contribution to the overall power variability as all the other components taken together when the other components are from the other brands.

In addition, when the same set of processors and memory was tested in each server at 0 percent compute load (i.e., at idle) the Supermicro server consumed 27 percent more power than the other brands.

Conclusions

We developed methods to compare server energy use and efficiency with limited sample sizes and measurement equipment. These methods can be used to assist in selecting server IT equipment having high energy-efficiency attributes.

The as-received results combined with the swapped-test results indicate that the different processor models (Sandy Bridge vs. Ivy Bridge) are likely responsible for the large as-received efficiency differences between the Dell and Supermicro server not the design of the server.

The three brands, equipped with a single set of processors and memory, had slightly different computational efficiencies at high computational loads. The efficiency variation was caused by power variation and not computational rate differences. However, considering the measured results and small sample size, it can not be concluded that one brand is more computationally efficient than either of the other two at a high computational load.

However, comparing the measurements at idle power levels, using the same single set of components, provided more interesting results. The Supermicro server used 27 percent more energy compared to the mean of the other two brands. This result was consistent with the as-received measurements at idle power levels. The research team concluded that the Supermicro server model in our demonstration, with the BIOS software used, was likely to use more energy at idle compared to the other two brands. This result, when considered in practice, may or may not be significant, depending upon how much time the servers are in an idle state.

Server efficiency variation caused by processors and memory power use differences as a group is significant, and appears to be on the order of +/- 3 percent at high server loads. However a similar fraction of server variance (when comparing the three brands) is attributable to all other components, sub-assemblies, and controls. Therefore, there are opportunities to improve server efficiency by changing the circuit design or controls (e.g., internal cooling) or by selecting components on the basis of efficiency.

Benefits

The energy use and performance measurement methods presented in this report can assist those specifying IT equipment to select models and configurations that have superior energy use efficiency.

When selecting server-type IT equipment, the purchaser will benefit by testing models from different brands or within the same brand to find those that meet the user's computational needs and that have low energy use at idle and good efficiency at high computational loads.

Energy use at low computational loads or at idle is significant. Purchasers will benefit from selecting servers that have optimum fan speed controls, efficient power supplies and power supply controls.

The methods presented apply to situations where very small sample sizes and limited instrumentation are available for performing a comparison. In addition, the methods presented can be performed using equipment commonly available in many data centers.

CHAPTER 1: Introduction

1.1 Summary

Data centers in the United States account for approximately 2 percent of the electrical energy consumed countrywide (Koomey 2011). Approximately one-half of this energy is used to provide cooling and other support (also known as *infrastructure energy*) for the electronic equipment. The cooling portion of this support historically and currently is the dominant contribution to the total infrastructure energy requirements.

In recent years, new products, improved controls, and expanded operational limits have been the focus of demonstrations that provide infrastructure energy savings. Examples include improving airflow management, computer room air handler (CRAH) efficiency, or computer room air conditioning (CRAC) efficiencies (Greenberg 2013; Coles 2012) or providing close-coupled cooling at the rack level (Coles 2014).

This demonstration focused on the energy use and efficiency of the IT equipment. Energy not consumed by the IT equipment also reduces the energy used by infrastructure systems, such as cooling and power distribution, that must be present to support the IT equipment. Considering that data centers use similar amounts of electrical power for IT equipment and infrastructure, saving energy going to the IT equipment can double the energy savings.

1.2 Demonstration Goals

The main demonstration goal was to quantify power consumption and efficiency of similar model servers from three different manufacturers using a very small sample size (in some cases, one) with limited measurement equipment. An additional goal was to develop comparison methods or testing protocols that might be used by others to compare server power consumption and efficiency.

In addition, other questions were explored during the process:

- How much variability (energy consumption and efficiency) was there from manufacturer to manufacturer for servers of a similar hardware generation?
- How much variability was there from server to server for machines that were supposedly “identical” (i.e., same manufacturer and same configuration)?
- Considering that energy-intensive commodity components must be incorporated in a typical server, what was the remaining energy savings and performance improvement potential?

1.3 Participants

In addition to colleagues working together at Lawrence Berkeley National Laboratory (LBNL), a number of other organizations were involved in this demonstration. Intel provided three

servers, along with much-appreciated and valuable technical guidance. Supermicro provided processor components and technical guidance. ServerTech provided power measurement equipment in the form of rack-mounted power distribution units along with technical support. On-site support was provided by engineers and technicians working in the High Performance Computing (HPC) Services group at Lawrence Berkeley National Laboratory.

CHAPTER 2: Methods

2.1 Introduction

Servers from three manufacturers (Intel, Dell, and Supermicro) were tested for the purpose of comparing energy and computational efficiency. Three servers from Intel were provided for the demonstration, along with one equivalent Dell and Supermicro (SM) model servers borrowed from the LBNL data center.

All the servers contained commodity components that are responsible for a large part of the electrical energy consumed. Examples of industry commodity components are processors, memory, and storage devices such as disk drives. Much of the server power and electronic design is dictated by the interface requirements and functionality of these commodity parts. However there are many design decisions the server manufacturer can make to reduce overall energy use and provide optimum performance while at the same time using commodity components. This demonstration attempted to quantify two separate power consumption categories within a server: (1) consumption by the key commodity components (Figure 2-1), and (2) all other server components.

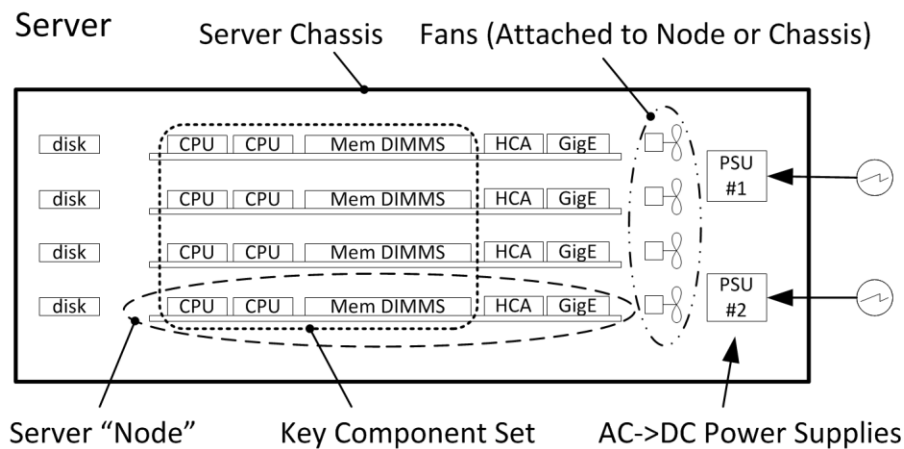
Servers are made from hundreds of electronic and passive components. Each of these component types has a variation in performance as received from the component manufacturer. Therefore a server manufactured one day is not likely to have the exact performance as one manufactured the following day. Quantifying power differences with appropriate statistical results using a small sample size is problematic because the power difference measured may appear to be significant but may be well within the values expected in the total population of the same model server. In addition, recent processors and related components contain a number of energy-efficiency features that may be controlled by the Basic Input/Output System (BIOS) settings. The BIOS configuration software varies from manufacturer to manufacture. Therefore not all efficiency modes and combinations provided by a given processor model are necessarily made available to the end user. The BIOS settings used in the demonstration were selected to be as close as possible for all of the servers, as determined by the demonstration team in collaboration with each server manufacturer.

Within the industry there are a variety of generic descriptive names assigned components or assemblies comprising a complete server or computer. In this paper the nomenclature will be as noted in Figure 2-1. The term *server* refers to the complete assembly typically received from an original equipment manufacturer (OEM). The server contains all the parts and assemblies needed as installed in a data center. In recent years servers contain a number of sub-assemblies referred to as *server nodes*, *node assemblies*, or *nodes*. These nodes have one or more motherboards or node boards each. The nodes are easily replaceable as a single assembly by the end user. Air-cooled servers contain fans mounted integrally attached to each node or installed as part of the chassis. In the case of the Intel server, the fans are included as part of the node assembly, which is not shown as such in Figure 2-1. The *server chassis* refers to the sheet metal structure assembly

and other parts that are thought of as permanently attached to each other. The chassis typically contains passive components. *Passive components* refer to materials, parts or assemblies containing a small amount of electronic functionality and energy use such as light-emitting diode (LED) lights, a power harness, and an on/off switch. In our demonstration, power supplies in a redundant configuration (quantity of two per server) slide in at the rear of the chassis, providing the power for all components and sub-assemblies in each server.

The servers used in this demonstration had similar enclosure dimensions and electronic configurations. The common attributes included: a two-rack unit (2U) form factor, four populated node bays with a motherboard for each node containing two sockets each capable of accepting E5-2600 series Intel processors, two power supplies and one disk drive per node. Each node motherboard had four or eight dual in-line memory module (DIMM) connectors per processor.

Figure 2-1: Nomenclature



Each server had a removable component set consisting of the processors and memory DIMMs installed in the four nodes (Figure 2-1). One component set (consisting of 8 processors and 16 memory DIMMs) from Intel server #1 was identified as the “Gold” set, or key component set, and is referred to later.

2.2 Goals

Attempting to quantify energy use and efficiency differences across three server brands using a small sample size (in some cases, one) was a primary goal. An additional primary goal was to develop comparison methods or testing protocols that might be used by others to compare server energy and computational efficiency. Detailed and possibly proprietary design information and specialized equipment would normally be employed for an effort similar to what was attempted. However, in our study, only electrical power measurements at the server inlet power cords were taken. Thus, the study had measurement constraints similar to those found in data centers with per-outlet power measurement capability.

Significant server design considerations are constrained by the integration requirements of commodity components such as processors and memory. In addition we attempted to quantify the

power consumption differences that can be attributed to particular manufacturer design decisions outside of the constraints of the integration of the commodity components.

This demonstration does not attempt to replace in-depth measurement and analysis that would be possible in a laboratory well equipped with specialized component power measurement instruments. An investigation by OEM manufacturers into the energy used in a server would include detailed component-level power measurement equipment and detailed design and component information. In comparison, the methods described in this report are meant to guide testing and analysis done by those having access to very limited sample quantities and measurement capability.

The demonstration describes and uses a specific computing efficiency measurement and calculation method. While the researchers feel the method is valid in its own right and may be similar to standard practice, the method used is not meant to redefine, displace, or compete with computational efficiency measurement methods or metrics described outside of this paper.

2.3 Demonstration Process


2.3.1 Introduction

The demonstration was performed in a data center on the LBNL campus in Berkeley, California. An empty rack was borrowed from the data center, and the servers used in the demonstration were installed (Figure 2-2). Details on the setup and equipment, test description and analysis follow.

2.3.2 Setup and Equipment

Each server was assigned a name: (1) Dell, (2) Supermicro, and (3) Intel #1, #2, and #3, starting at the bottom of the rack. Some nearby perforated floor tiles were replaced with non-perforated models to reduce fluctuations of air temperature entering the servers at their front bezels. Blanking panels (not shown in Figure 2-2) were added along the complete front of the rack.

Figure 2-2: Servers Installed In Demonstration Rack



Server Descriptions

The demonstration required borrowing server equipment from the LBNL HPC data center and Intel. Therefore it was not practical to require that all servers have the exact same key components, such as processors and memory. It was assumed that processors and memory would need to be changed or swapped between brands to provide a fair power consumption and efficiency comparison. The server model selected from each manufacturer had the same capabilities for accepting processor and memory types. Initial tests were completed using the “as-received” configurations, as shown in Table 2-1. Refer to Table 2-1 to identify how components or assemblies were configured for a test using a given server.

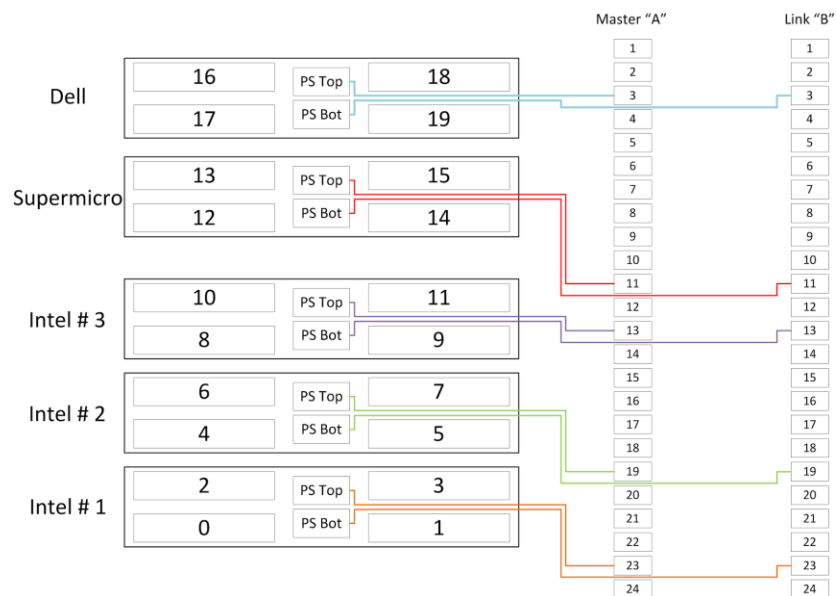
Table 2-1: As-Received Server Configurations

Server Manufacturer	Intel [®]	Dell	Supermicro
Model #	Bobcat Peak H2216XXKR	PowerEdge C6220II	6027TR-HTRF+
Processor Model (Intel [®] Xeon [®])	E5-2650 v2 8 core 2.6GHz	E5-2670 v2 10 core, 2.5GHz	E5-2650 v0 8 core, 2.6GHz
Processor Technology	Ivy Bridge	Ivy Bridge	Sandy Bridge
Node Quantity	4	4	4
Processors/Node	2	2	2
Memory Slots/Node	8	16	16
Memory DIMMs per Node	8	8	8
Memory DIMM Type	4 GB DIMM DDR3-1600	8 GB DIMM DDR3-1600	8 GB DIMM DDR3-1600
Server Power Supply Capability	2 x 1200 W	2 x 1200 W	2 x 1620 W

Demonstration Nomenclature

Each server node was assigned a unique number (Figure 2-3). Physical labels were placed on the nodes to allow tracking when the nodes or components moved. The node number assignments are referred to throughout this report. For example swapping nodes from Intel Server #1 to Intel Server #2 is referred to as “Intel Server #1 (nodes 0–3) in Intel Server #2.” When just the processors and memory components were swapped, an example description is “Intel Server #1 (0–3 node components, key components, or “Gold” components) in Intel Server #2 nodes.”

Figure 2-3: Server Node Number (0–19) and Power Distribution Unit Plug Assignments



Rack Power Distribution Unit (PDU)

A rack-mounted PDU set (master and link), model numbers CLG-24V2C415A1 and CSG-24V2C415A1, donated by ServerTech, was used to measure and provide a data collection path for the power consumed by all server power supplies. In addition, the PDU provided a convenient air temperature probe that was used to record the air temperature in front of the servers near the front bezels.

2.3.3 Data Collection

There were three sources of data: (1) PDU measurements using Simple Network Management Protocol (SNMP) at each server power supply inlet and front panel air inlet, (2) intelligent platform management interface (IPMI) data collected from each server, and (3) performance results from LINPACK software. The data collected that were used for calculating the results or other observations are listed in Table 2-2. The software tools used for data analysis were R and Microsoft Excel.

Table 2-2: Collected Data Used

Description	Source	Collection Method	Collection Interval
Server Power Supply Power (kW)	PDU	SNMP	20 seconds
Server Front Panel Air Temperature (°C)	PDU	SNMP	20 seconds
Processor Temperature (°C)	server	IPMI	20 seconds
Server Fan Speed (RPM)	server	IPMI	20 seconds
LINPACK Performance (GFLOPS)	LINPACK	software report	per run

GFLOPS = giga-floating-point operations per second

Data collection intervals of 1 and 5 seconds were attempted, but the variance in IPMI server response time varied across the different server manufacturers resulting in missing data. No missing data were observed when the polling interval was changed to 20 seconds. Therefore a data collection interval of 20 seconds was used for this demonstration.

2.3.4 Test Description

Test Software and Server Setup

Servers in a typical data center have a variable software load. For example, if servers are used for administrative purposes the load may be high during the work week and lower during weekend or holiday periods. Methods to provide a constant, sustained, and repeatable load with ability to measure performance were needed for power consumption and performance measurements.

The High-Performance LINPACK (HPL) software package was used to generate constant and repeatable computational loads, as well as to measure and report performance. Due to its high scalability and high efficiency, HPL is widely used in the HPC industry as a stand-alone benchmark to evaluate supercomputers.

Three target load levels were selected for comparison:

- Idle: no defined computing executing
- 50 percent: half of the central processing unit (CPU) cores are computationally loaded using cyclic process distribution
- 100 percent: all CPU cores are computationally loaded

Results of tests using the 50 percent power level were reported but not analyzed because the power test results appeared similar to those from the 100 percent power level tests.

HPL Use

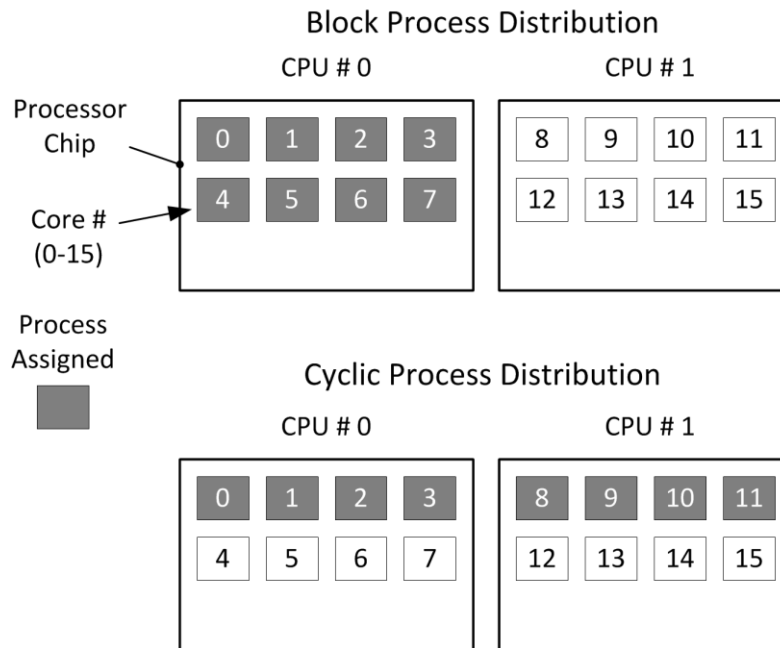
The High-Performance LINPACK software package was used to produce a computational load for two purposes: power measurement and performance measurement.

- Power Measurement.** Although HPL provides a way to run a series of problems in one batch to extend the total run time, loading of each new problem requires that the existing memory be depleted and a new random linear system to fit the memory be generated. In practice this process causes multiple several-second power drops during each power level run affecting the power measurement quality. To improve the quality of energy use measurements, HPL was modified to allow the calculation of the existing linear system to repeat and reduce power fluctuations during the test period for each power level. Each test run lasted approximately three hours—one hour at each power level: idle, 50 percent, and 100 percent.
- Performance Measurement.** Separate tests were completed for each hardware setup using HPL configured as it would be used for an industry benchmark. To achieve the optimal CPU performance and compute efficiency (MFLOPs per W), HPL was optimized with Intel compilers and Intel Math Kernel Library (MKL) for the Intel processor model being tested.

Cyclic and Block Process Assignment

Due to the multi-core (multiple processors) architecture of these servers, there are a number of ways to apply a 50 percent software load. Processes can be assigned to cores on a given node using a cyclic or block distribution scheme. Figure 2-4 depicts block and cyclic process distributions on a server with two octa-core CPUs and 50 percent of target load. As shown in Figure 2-4, with a block distribution of processes, CPU 0 is fully loaded, while CPU 1 remains idle. The cyclic scheme of process assignment was used for all tests.

Figure 2-4: Block and Cyclic Distribution of Processes (Eight-Core Processor Used as Example)



BIOS Settings

Another important factor that can influence the power consumption and performance are the BIOS version and BIOS settings applied. The BIOS settings can change the power consumption and performance significantly. For simplicity, in this study we picked the default settings across all three vendor platforms and made a few adjustments in an attempt to achieve configurations that fairly equalize power consumption and performance. The BIOS settings that can affect the power profile the most include: Turbo Boost, Hyper-Threading, SpeedStep, CPU C1E State, CPU C3 State, and CPU C6 State. Appendix A lists the BIOS version and setting procedure used for each server brand.

In our tests both Hyper-Threading and Turbo Boost were disabled, using the BIOS settings to reduce the power fluctuations that might be encountered due to these technologies. Normally users would leave these functions enabled for more efficient processing of variable workloads.

Test Plan

Five tests were performed as part of understanding, quantifying, and comparing the power use and performance across three brands of servers:

- Measure the power consumption and performance of the five servers as-received.
- Install a single set of four Intel nodes in each of the three Intel servers, and measure power consumption. This test is later referred to as *Swap A*.
- Install a set (8 processors and 16 memory DIMMs) from one Intel server node set into another Intel server node set and measure the power consumption. This test is later referred to as *Swap B*.
- Install a set (8 processors and 16 memory DIMMs) from one Intel server into the other two server brands to obtain power and performance measurements of all three brands using the same key component processor and memory set. This test is later referred to as *Swap C*.
- Simulate experiments using results obtained by measurements combined with component and subsystem estimates.
- Intel engineers recommended that the E5-2650 v2 model processor be selected for use in the key component set because of its 95W thermal design power (TDP) limit as opposed to the E5-2670 v2 model with 115W TDP limit. In general, the lower the TDP limit the less power consumption variation between parts. The E5-2650 v2 processor was therefore used in the key set of components (8 per server, 2 per node), along with the memory DIMMs (16 per server, 4 per node) found in Intel Server #1, which were 4 GB DDR3 1600 1.35v models from the Crucial memory brand.

As mentioned above, three power levels were measured for each test run. The lowest level was achieved by letting the servers idle without a software application or LINPACK program being activated. The medium power level referred to at 50 percent power was provided by assigning LINPACK to use half of the cores available, using cyclic assignment on each node. Full or

“100 percent” power was provided by configuring LINPACK to use all available cores in each node.

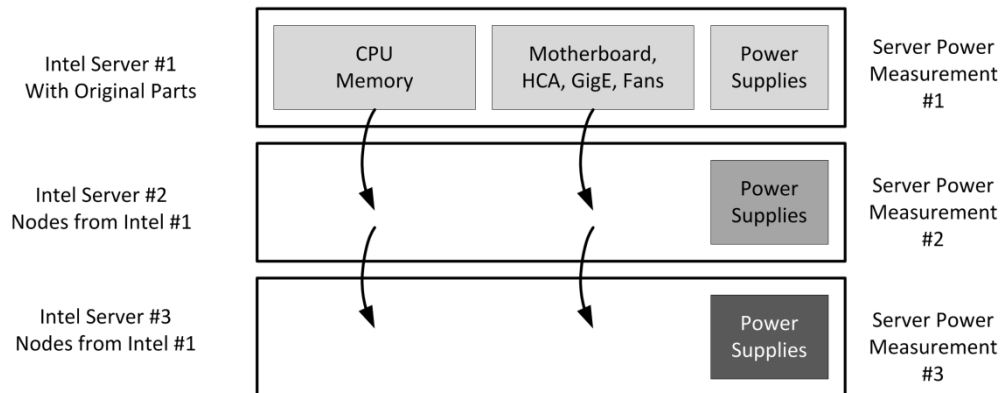
As-Received Tests

Each of the five servers was tested in the as-received configuration to help with the data collection system development and to get an idea of power and efficiency differences when the processors and memory are varied. The configurations of the as-received servers are listed in Table 2-1. There were a variety of DIMM manufacturers found in the five servers, except for Intel server #1 which was populated with a single model from one manufacturer. The processors (quantity 8) and memory (quantity 16) in the as-received Intel Server #1 became what was referred to as the *Gold* or *key* component set.

Swap A: Swap Intel Nodes Between Intel Servers

This test was performed to see how the power varied across the three Intel servers while keeping the node assemblies constant. The results should confirm to what extent the total server power is dominated by the components within the node, compared to the variance associated with the power supplies. Figure 2-5 shows a diagram of the test configurations.

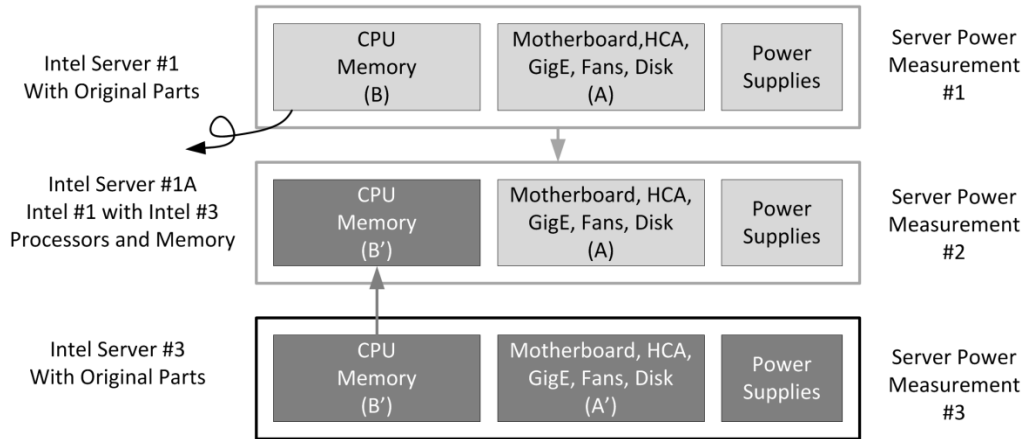
Figure 2-5: Swap A: 4 Nodes from One Intel Server Tested into All Intel Servers



Swap B: Swap Processors and Memory from One Intel Server to Another Intel Server

This test was performed to see how the power varied from one server to another while holding all components and assemblies constant, except for the processors and memory. The component swap is pictured in Figure 2-6. The power differences were calculated for component groups using the measurement results (Appendix B). This will compare the effect of changing the processors and memory (as a group) to changing all other components.

Figure 2-6: Swap B: Processors and Memory from One Intel Server into Another Intel Server



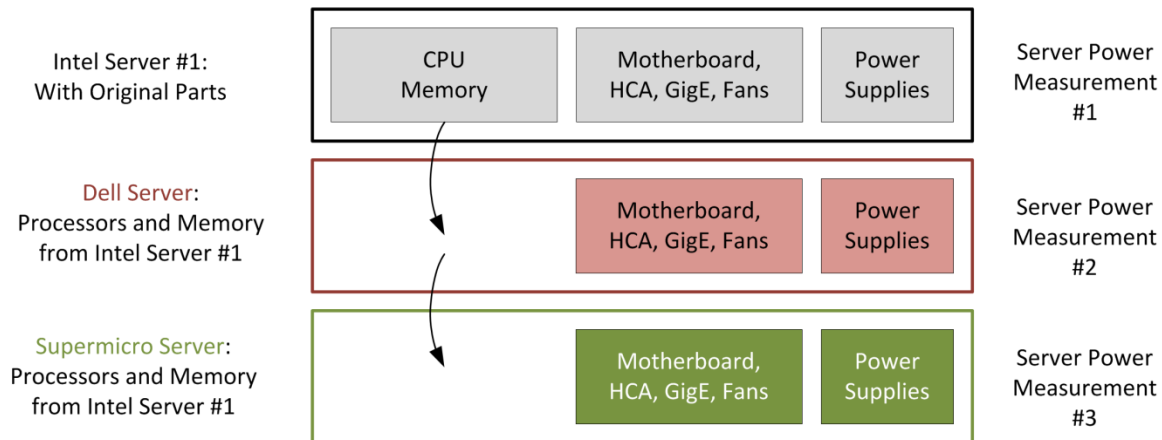
Swap C: Key Component Set Installed in All Brands

The best comparison that could be undertaken was to select a single set of processors (8 Intel E5-2650 v2) and memory (16 4 GB DIMMs), install these components on each node (4 per server) for all three brands of servers, and measure the power use and performance.

The final two tests were performed by measuring the power and performance for the Dell and Supermicro servers with the key component set from Intel #1 (Figure 2-7) using a 100 percent software load. These final tests provided data for the power use and performance needed to compare the servers from the three manufacturers.

These tests should indicate to what extent the server power level and/or performance is attributable to the processors and memory, compared to all other components, when incorporated by different manufacturers into a final product.

Figure 2-7: Swap C: Key Component Set Installed in All Brands



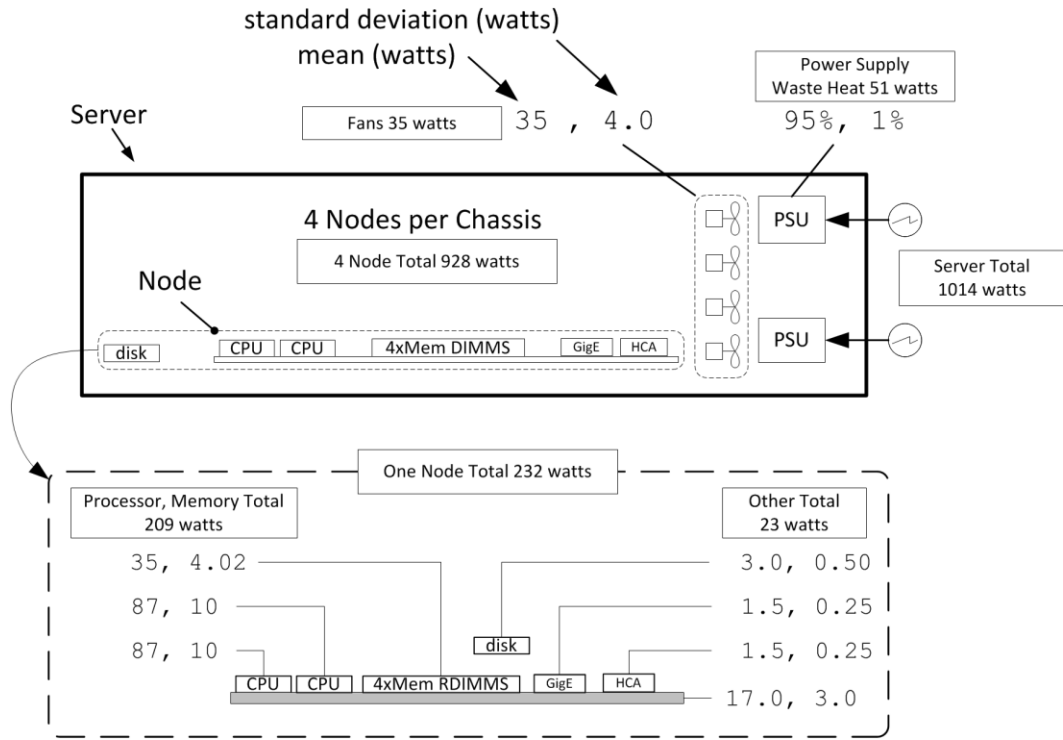
Simulations

Monte Carlo simulations were developed using the R statistics language, and this provided estimates of statistical distributions of total server load for three 100 percent power-level scenarios. Monte Carlo simulations were used to calculate the mean and standard deviation of server power for 2,000 independently constructed servers. The simulations were completed using the input values shown in Figure 2-8. The simulation mean input values were estimates, and were adjusted to match the mean of the power measured. While the total of the mean values was assumed to be correct per the measured data, the individual values were not measured, and were estimates. The individual mean values were arrived at by reviewing specifications found online and in consultation with Intel and Supermicro. Additional discussion regarding the individual component mean and standard deviation estimates are provided in the Results section.

Three cases were simulated:

- Simulation #1: All components were randomly selected from a normal distribution using the estimated mean and estimated standard deviation for components or sub-assemblies for each server.
- Simulation #2: The mean values for processors and memory were held constant. All other components were sampled. The result should provide an estimate for the expected server power population mean and standard deviation if the processors and memory are held constant.
- Simulation #3: This simulation used the mean values for all components except the power supplies. The result provided an estimate for the expected server power standard deviation, which can be compared to the magnitude of the changes in the Swap A test.

Figure 2-8: Simulation Input Data



2.3.5 Analysis

The mean power level and standard deviation were calculated for each one-hour test period for the above-described tests. The performance results from LINPACK for the 100 percent power-level tests were combined with results from the power measurement results. The resulting performance efficiency metric was calculated by using the LINPACK run results listed in GFLOPS (10^9 floating point operations per second) (performed until the performance measurement run finished) and the mean of the power measurement to obtain MFLOPS/W (10^6 floating point operations per watt).

CHAPTER 3: Results

The results of the demonstration are reviewed as follows:

- As-Received
- Swap A
- Swap B
- Swap C
- Simulations
- Other Findings

3.1 As-Received Power Consumption

The as-received power measurement test results are graphically presented in Figure 3-1, with the raw time-series data shown in Appendix C. The as-received Supermicro server was configured with a Sandy Bridge processor, while the other brands were fitted with Ivy Bridge processors. The Supermicro power consumption is larger by 49, 26, and 35 percent for the idle, 50 percent, and 100 percent power levels, respectively, compared to the mean of the other brands. The computational efficiency of the Supermicro as-received server was 27 percent lower compared to the mean of the other brands combined (Figure 3-2). Server brand comparisons regarding the power consumption or efficiency should not be made from the as-received test results, because of the processor model and other configuration differences. The processor difference is likely the main reason for the power use differences, not the vendor's server design.

The tests were performed to explore performance differences encountered when different model processors and other components were varied. The mean values and standard deviation calculations for the as-received tests and other tests are included in Appendix D.

Figure 3-1: As-Received Server Mean Power - All Brands





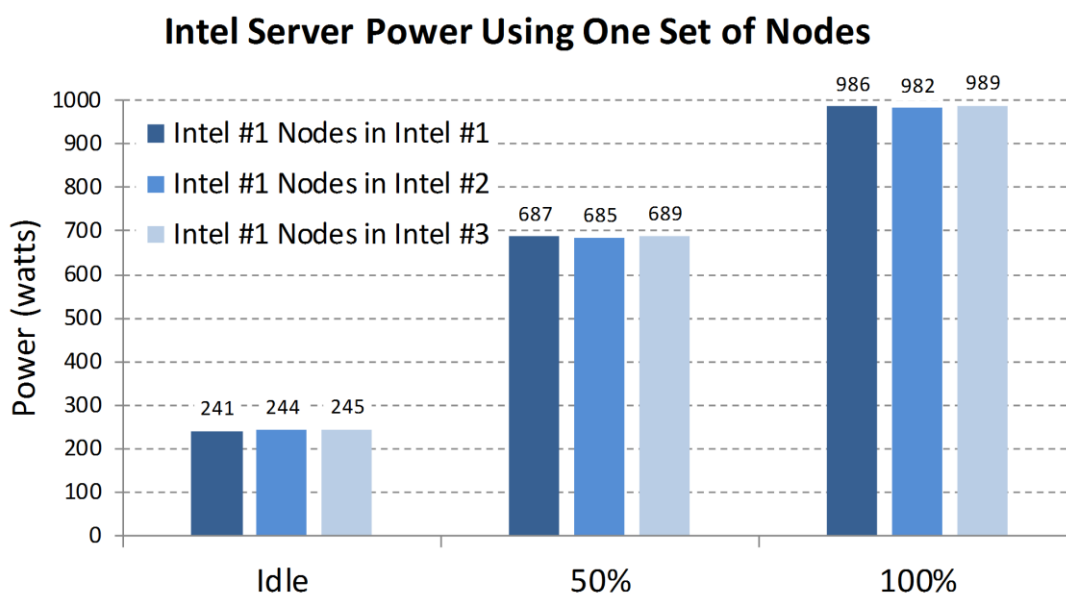
Figure 3-2: As-Received Efficiency, by Brand



3.2 Swap A: Swapping Intel Nodes

The server-level power was measured with all nodes from one Intel server installed in all three Intel servers. The raw time-series data is graphed in Appendix E. The Intel server as-received measured power range was 985 to 1037 W (a range of 52 W) (Figure 3-1) at a 100 percent software load. Server #1 consumed 985 W. When the nodes from Intel Server #1 (including the fans) were swapped into server #2 and server #3, the resulting power consumption was 982 W and 989 W, respectively, (a range of 7 W across three Intel servers) (Figure 3-3). This shows that all three Intel power supplies in this test have nearly identical efficiency, which also suggests that the variability in this model of power supply is very low. We conclude that most of the power variability between the Intel servers in their original configurations is due to variation in node power, not due to power supply efficiency variation.

Figure 3-3: Intel Server Power Using One Set of Nodes



3.3 Swap B: Swap Key Components from One Intel Server to Another Intel Server

From Swap A we learned that the Intel server power supply variation was small, and was not the cause of the variation in power consumption among as-received Intel servers. We then tried to determine how much of the variability was attributable to the processors and DIMMS (which we termed “key components”) and how much was attributable to the other component types (e.g., motherboard, host channel adapter (HCA), GiGabit Ethernet (GigE), fans, and disk drives).

Swap B, shown in Figure 3-4, was completed by moving key components (8 processors and 16 memory DIMMs) from Intel Server #3 to Intel Server #1. The resulting server, which we termed Intel Server #1A, then had key components from Server #3, and all other components from Server #1.


The power measurement for Server #1 was 986 W, and for Server #1A it was 1030 W, a difference of 44 W. If we assume the power supply efficiency is 100 percent, this implied that, taken as a group, key components (memory and CPUs) from Server #3 consumed 44 W more than those from Server #1. This implies (as shown in Appendix B) that the other components from Server #3 consume 11 W less than those from Server #1.

In fact, of course, power supply efficiency was not 100 percent. If we assume the power supply efficiency $E = 95$ percent (an average of 80 PLUS™ Platinum and Titanium ratings at a 50 percent load) and is the same for Server #1 and Server #3—which Swap A confirm—then the key components from Server #1 consume 42 W less than those from Server #3, and the other

components (not including the power supplies) from Server #1 consume 10 W *more* than those from Server #3. 80 Plus (trademarked 80 PLUS) is a voluntary certification program intended to promote efficient energy use in computer power supply units.

In short, the key components are responsible for most of the variability in power consumption in this experiment. Although the Swap B experiment was performed in only two servers, we believe this to be a general result: the total power consumed by the other (non-key) components was relatively small, so it seems unlikely that the server-to-server variability in those components could have accounted for 30 W or more. We concluded that most of the variability in power consumption between servers was due to variability in the key components.

Figure 3-4: Intel Server Power Using One Set of Nodes



3.4 Swap C: All Brands - Using One Set of Key Components

The Swap C results for measured power and compute efficiency are shown in Figure 3-5 and Figure 3-6, respectively. The raw time-series data for the Swap C test are graphed in Appendix F. The Dell server had the best compute efficiency, but it is unknown if this would be true of a larger sample. Supermicro had 4.4 percent lower compute efficiency, and Intel had 9 percent lower compute efficiency, compared to the Dell server. The 95 percent confidence intervals are shown on Figure 3-6.

A result that stands out is that the idle power level for the Supermicro (SM) server was 27 percent above the mean of the other two brands. However, this result was based on only a single Supermicro server, so it was only suggestive rather than definitive. Server idle power use is significant and contributes to overall inefficiency at a large number of data centers. An

Uptime Institute survey records that at many data centers, 10 percent or more of servers are idle at any point in time (Uptime 2013).

Figure 3-5: Server Power Consumed for All Brands Using Same Key Components


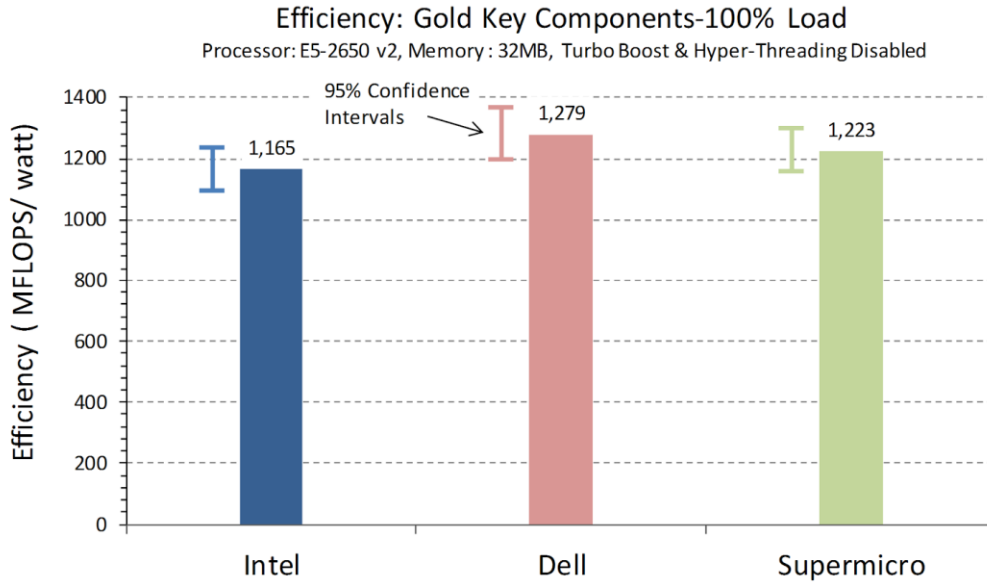


Figure 3-6: Compute Efficiency for All Brands Using One Set of Key Components



3.5 Experiment Simulations

An interesting question is the extent to which the results from our small sample of servers can be generalized to the wider population of servers in the real world. To address this question, we used Monte Carlo simulations, implemented with the R statistics language, to simulate, in part, the results of our component-swapping experiments. Specifically, we were trying to determine how much variability can be expected when components are swapped from one server to another, and how well our experimental results conformed to those expectations.

Assumptions about component power consumption, and variability in power consumption, are detailed in Figure 2-8. The mean and standard deviation of each component type were chosen to be consistent with (a) published power consumption figures, (b) with private communications with industry representatives, and (c) our experimental results. The component mean values were adjusted so that the total mean simulated server power matched that of the three Intel servers at a 100 percent software load.

Intel indicated the E5-2650 v2 processor should have more consistent power consumption for a given set of conditions. However, our results strongly suggest that, at least for the components in our servers, one set of eight of these processors consumed approximately 40 W–50 W more than another set, which is almost 5 percent of the server power. Although it is possible that we were sent an unusually variable set of processors, we think it is more likely that there is processor-to-processor variability in power consumption that is on the order of several percent.

We observed an approximately 5 percent variation in power from one group of eight processors to another. If this result is typical—if groups of eight processors vary in power by about 5 percent (45 W) on average—then this implies that individual processors could vary by about 16 percent (14 W). It is impossible to be certain of this result because (a) memory was switched along with the processors, and (b) the small sample size in these experiments limits the ability to draw firm conclusions about the larger population.

The simulation standard deviations for the components were adjusted until the overall standard deviation of the server power was near 30 W, which is consistent with our experimental observations.

Each of the simulation runs calculated power consumption estimates for the assembly and test of 2,000 servers.

Simulation #1: Simulating the As-Received Intel Equipment

The Sim #1 results are presented using a histogram showing the overall mean and standard deviation (Figure 3-7). The measurement results from the as-received Intel server tests are indicated by the “R” text above the corresponding value on the histogram axis.

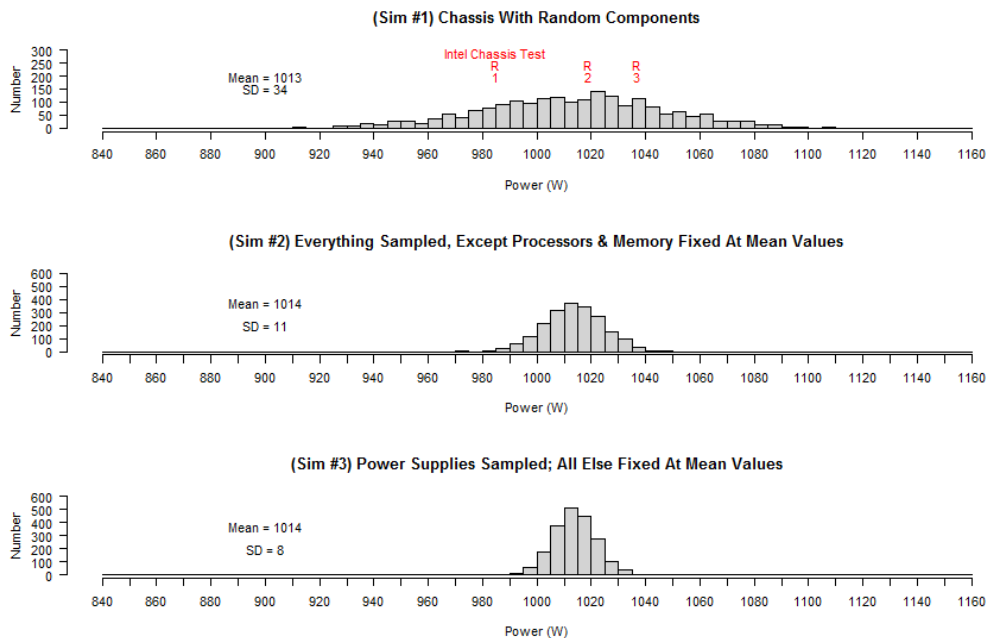
Simulation #2: Key Components Fixed at Mean Values, All Others Sampled

The Sim #2 histogram (Figure 3-7) shows the amount of variability among servers if key components did not vary at all. The resulting standard deviation of only 11 W was approximately 75 percent lower than the Sim #1 results, and was inconsistent with the variation we observed between our three as-received Intel servers. This supports the belief that there was substantial variation in power consumption among the CPUs.

Simulation #3: All Components Fixed at Mean Values, Except Power Supplies Are Sampled

The Sim #3 histogram (Figure 3-7) results contain an even lower value for overall standard deviation compared to the Sim #2 results. A considerably lower standard deviation value is expected because the standard deviation inputs for all but the power supplies do not vary and the power supply standard deviation was set at 1 percent in an attempt to be consistent with the Swap A variation results.

Figure 3-7: Simulation Results



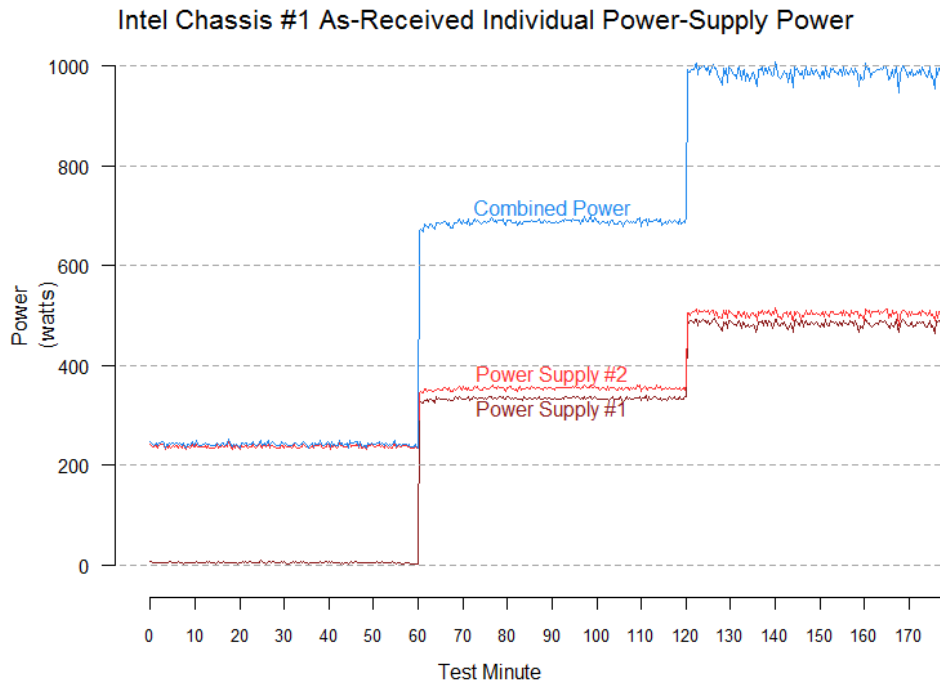
The simulation exercise results could be made more accurate by measuring individual component and assembly power consumption using instruments in a laboratory setting or by obtaining more detailed information from component manufacturers. A more accurate mean power value would be provided by more than three measurements, but the sample size was limited in this demonstration to three. The main goal of the simulations was to demonstrate that reasonable estimates of the range of expected power consumption can be made without detailed sub-measurements. Simulation results suggest that there may be a significant range of power consumption and efficiency (+/- 5 percent to +/-10 percent) for servers of the same make and model.

3.6 Other Observations

The power consumption for each power supply for a server (two power supplies per server) was measured and recorded separately for all server tests. Viewing the raw data for individual power supplies provided an interesting observation. The power consumed when comparing each power supply for a server was close to equal during all tests except for the idle-level Intel runs. During the Intel server idle tests, one of two power supplies provided approximately 98 percent of the total server power (227 W) required, while the other power supply provided the remaining amount of approximately 5 W (Figure 3-8). Discussions with Intel uncovered that the servers used in our tests were equipped with cold redundancy. Cold redundancy, part of the Intel Efficient Power Technology suite, reportedly improves efficiency by powering off any power supplies that are not needed to support the loading condition. In our idle power level

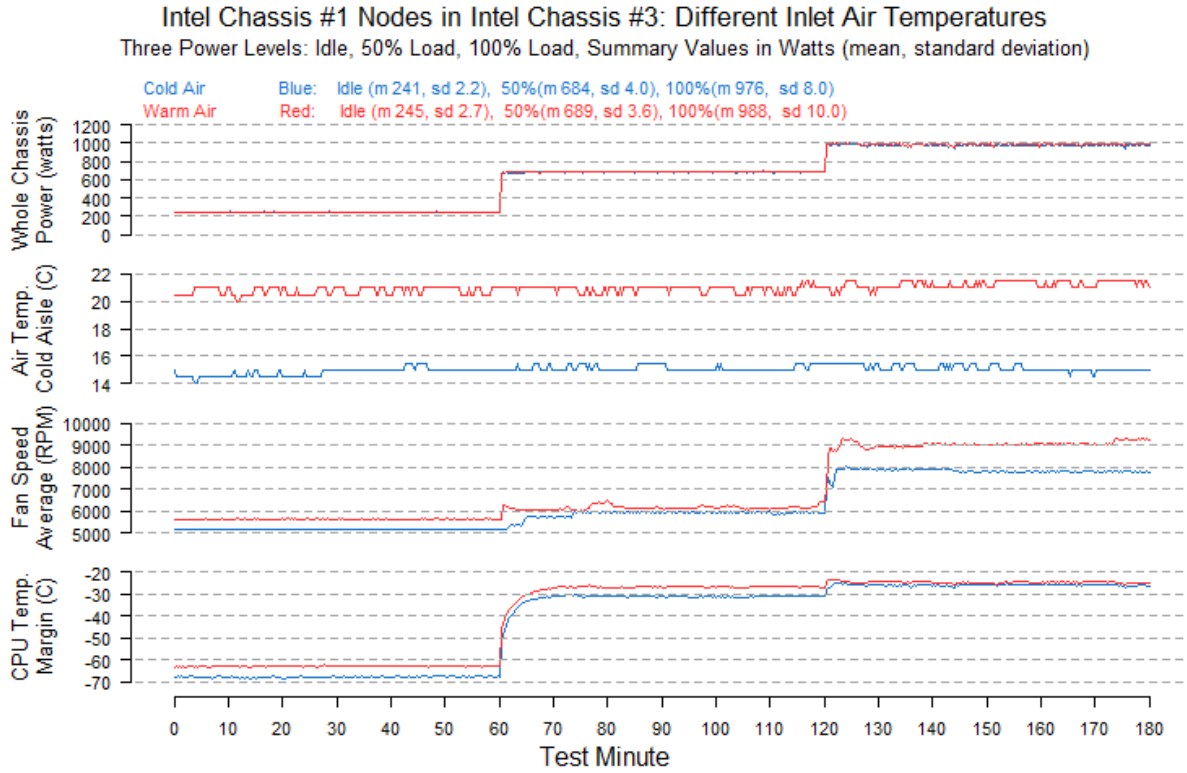
tests, the Intel server appeared to have an efficiency advantage of 5 percent, with an average power measurement of 232 W compared to Dell at 245 W.

Figure 3-8: Intel Server Individual Power Supply Measurements



Reports, including Moss and Bean (2011) and anecdotal accounts often cite that as the front bezel air inlet temperature rises, the server power also increases due to an increase in leakage current and/or an increase in server fan power. During our testing two identical (except for the inlet air temperature) test runs were completed with same Intel server. The raw data for these two runs are plotted (Figure 3-9). The power increased 1.2 percent as the inlet temperature increased from 15°C to 21°C, as the average server fan speed increased from 8,000 to 9,000 rpm. We conclude that the 1.2 percent increase in power consumption was due to the increase in server fan power and not leakage current. The leakage current was assumed to be constant because there was almost no change in processor temperatures. This data point can add to the body of knowledge on the subject of how higher inlet temperatures affect server power consumption. The 1.2 percent power increase for a 6°C inlet temperature increase does not likely apply at higher inlet temperatures because the difference in fan power is not constant as a function of fan speed, especially at the highest fan speeds, which were not encountered in our tests.

Figure 3-9: Intel Server Test with Different Inlet Air Temperatures



CHAPTER 4:

Conclusions and Recommendations

4.1 Conclusions

Server power consumption and computational efficiency depend on many complex interacting phenomena for a given computing load. The variables include the quantity and model of processor(s) and memory, power supply sizing and efficiency, motherboard design (such as voltage regulator circuit efficiencies), BIOS settings, and cooling design including fan quantity, type, redundancy, and controls.

The three brands, equipped with a single set of processors and memory, had different compute efficiencies at high software loads. The 95 percent confidence intervals for both power consumption and more important, efficiency, had some overlap. Therefore, considering the measured results and small sample size, it cannot be concluded that one brand is more computationally efficient than either of the other two at a high computational load.

This study's data showed that the Supermicro server had an idle power level 27 percent higher than the average of the other brands when configured with the single key component set. The higher idle power for the Supermicro server may not be typical of the population of this model of Supermicro server. There could be multiple causes for the Supermicro high idle power, including a low-probability combination of unusually inefficient components, higher-than-needed fan power, poor power supply efficiency, or an undetected fault. Our conclusion is that the tested Supermicro model population, using the BIOS indicated in Appendix A, has higher idle power than the other brands.

Idle power is more easily compared by measuring power at a given load because specialized load-producing software is not required. If testing resources are limited, idle power comparison testing alone should be considered, and may yield useful results.

The as-received power differences across the three Intel servers, combined with the Swap A and Swap B results indicated that a significant power use variance was related to processors and memory. The measurements for the three as-received Intel servers were 1,166 W, 1,101 W, and 1,127 W. The range was 66 W, or 6 percent of the average power. The results when the single set of processors and memory were installed in all brands (Swap C) were 986 W, 899 W, and 933 W. This Swap C range was 87 W, or 9 percent of the average power. For Swap C, the power, not the efficiency, was used for convenience, as the performance (GFLOPS) varied by less than 1 percent across the three brands at the 100 percent software load power level. In our data, the processor and memory as a group, and all other components as another group, contributed the same magnitude to the total server power variation. The variations for both groups were on the order of +/- 3 percent to +/-4.5 percent each of the mean server power value at 100 percent load. We conclude there are significant opportunities to improve server efficiency outside of the constraints provided by using commodity processors and memory components.

This study found that the best set of 8 processors and 16 DIMMS consumed about 45 W less power than the worst set, out of the three sets available to us. Out of the 24 processors in our servers, it is very unlikely that all 8 of the best (highest-efficiency) ones are in our best set and that all 8 of the worst ones are in the worst set. Similarly, the best set probably does not contain all of the best DIMMS. If we were to go board-by-board or component-by-component to select the best boards or individual components, it should be possible to construct a set of components that significantly outperforms any of our three sets.

The three tests (Swap A, Swap B, and Swap C) have a common theme. Each of these tests measures the power used, and in some cases, efficiency, when one component subset is tested in two or three servers while all other components remain constant. This type of testing can provide insight into the power use of certain components or component subsets without measuring power at various locations inside the server.

As discussed in the Methods section and in Appendix A, currently available servers have a variety of power use and performance features that are enabled or disabled using the BIOS. Individual BIOS settings or combinations of settings can affect the power use and efficiency of a server significantly. The BIOS setting options are not consistent across server brands, even when configured with the same optional components such as processors, memory, and storage. Understanding the BIOS setting options across the servers being compared is an important part of server energy use and performance comparison testing.

Each server has two power supplies. The power was measured at each power supply inlet. For the Dell and Supermicro servers, the total load was shared equally by both power supplies. In contrast, at the idle power level, the Intel server power supplies were not equally loaded. As mentioned, the power supplies in the Intel servers were managed using cold redundancy, which is part of Intel's Efficient Power Technology. This technology puts one power supply in a standby mode and transfers virtually the entire load to one power supply. This should provide better overall efficiency, since power supply efficiency typically is better at higher load factors. Analysis of the Intel server power data showed a 5 percent reduction of power use at idle compared to the Dell server. This suggests the cold redundancy scheme may be effective at reducing idle power. To confirm this, future experiments could be performed to test the Intel servers with the cold redundancy disabled.

4.2 Recommendations

When selecting server-type IT equipment, the purchaser should test models from different brands or within the same brand to find those that meet the computational needs at high overall energy efficiency. Simple testing as described in this paper can identify models that have low energy use at idle and good efficiency at high computational loads.

The processor model selected has a significant effect on computational efficiency. It is suggested that purchasers become aware of the processor and memory technology differences as technology changes to specify components with better efficiency. For example, there are significant efficiency improvements in the newer Ivy Bridge processors compared to the Sandy Bridge processors.

Energy use at low computational loads or at idle is significant. Selecting servers that have efficient power supplies or power supply controls, such as cold redundancy, is a key factor for saving energy.

A more in-depth study of this subject is suggested for future work, to better understand and improve server energy efficiency and test methods and to provide additional guidance to help purchasers select IT equipment with optimum energy efficiency.

GLOSSARY

Term	Definition
BIOS	basic input/output system
CPU	central processing unit
CRAC	computer room air conditioning
CRAH	computer room air handler
DIMM	dual inline memory module
GB	gigabyte
GFLOPS	giga-floating-point operations per second
GigE	GiGabit Ethernet
HCA	host channel adapter
HPC	high-performance computing
HPL	High Performance LINPACK
IPMI	intelligent platform management interface
IT	information technology
kW	kilowatt (1000 W)
LBNL	Lawrence Berkeley National Laboratory
LINPACK	LINear equations software PACKAge
MFLOPS	million floating point operations per second
OEM	original equipment manufacturer
PDU	power distribution unit
PSU	power supply unit
SM	Supermicro
SNMP	Simple Network Management Protocol
TDP	thermal design power
W	watt

REFERENCES

- Coles, Henry, and Steve Greenberg (Lawrence Berkeley National Laboratory), and Corinne Vita (Vigilent). 2012. *Demonstration of Intelligent Control and Fan Improvements in Computer Room Air Handlers*. Publication number LBNL-6007E.
- Coles, Henry (Lawrence Berkeley National Laboratory). 2014. *Demonstration of Rack-Mounted Computer Equipment Cooling Solutions*. Publication number LBNL-6659E.
- Greenberg, Steve (Lawrence Berkeley National Laboratory). 2013. *Variable-Speed Fan Retrofits for Computer-Room Air Conditioners*. Retrieved September 30, 2014 from http://energy.gov/sites/prod/files/2013/10/f3/dc_fancasestudy.pdf.
- Koomey, J. 2011. *My new study of data center electricity use in 2010*. July 31. Retrieved February 7, 2014 from <http://www.koomey.com/post/8323374335>.
- Moss, David (Dell Inc.), and John Bean, Jr. (Schneider Electric). 2011. Schneider Electric Data Center Science Center. White paper 138. Retrieved October 8, 2014 from <http://www.apcdistributors.com/white-papers/Cooling/WP-138%20Energy%20Impact%20of%20Increased%20Server%20Inlet%20Temperature.pdf>
- Uptime Institute. 2013 Uptime Institute Annual Data Center Industry Survey Report and Full Results. Retrieved September 30, 2014 from http://uptimeinstitute.com/images/stories/DataCenterSurvey_assets/uptime-institute-2013-data-center-survey.pdf.

Appendix A: BIOS Adjustment Procedure

Intel (BIOS version: SE5C600.86B.02.02.0002 [12/23/2013])

1. Load default setting profile
2. Turbo Boost: Enabled -> Disabled
3. Hyper-Threading: Enabled -> Disabled
4. C3: Disabled -> Enabled

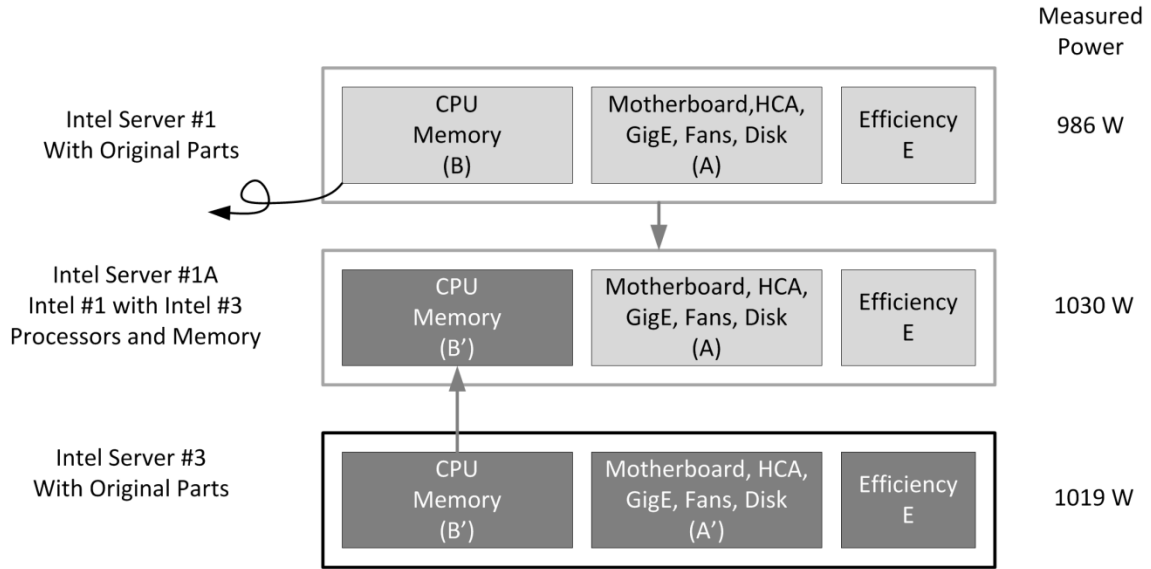
Supermicro (BIOS version: 3.0a [02/19/2014])

1. Load default setting profile
2. Disable the following
 - Hyper-Threading
 - VT
 - VT-d
3. Power Technology: Energy Efficient -> Custom
 - Turbo Mode: Enabled -> Disabled
 - C3 Report: Disabled -> Enabled

Dell (BIOS version: 2.3.1 [01/02/2014])

1. Load default setting profile
2. Hyper-Threading: Enabled -> Disabled
3. Turbo Mode: Enabled -> Disabled

Appendix B: Swap B Analysis



$$(A + B) / E = 986$$

$$(A + B') / E = 1030$$

$$(B' - B) / E = 44$$

$$(B' + A') / E = 1019$$

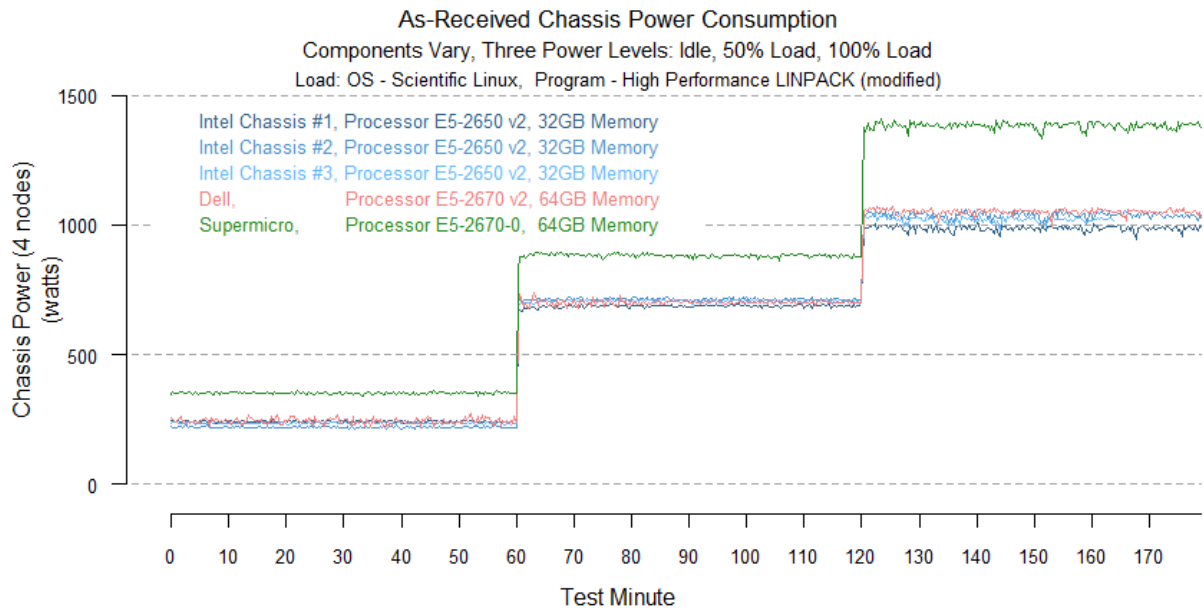
$$(B' + A) / E = 1030$$

$$(A' - A) / E = -11$$

$$(B' - B) = 42 \text{ [@ 95\% E]}$$

$$(A' - A) = -10 \text{ [@ 95\% E]}$$

Appendix C: As-Received Measured Power Data



Appendix D: Test Results Data

Note: Performance listed per node. To obtain per server performance multiply by four.

Efficiency is listed in this Appendix D as GFLOPs/joule. To obtain server MFLOPS/watt multiply GFLOPS/node by 4 and divide by number of mean watts at 100% load. Then multiply by 1000. Example: First row of table D-1, MFLOPS/watt = $((287.1 \times 4) / 986) \times 1000 = 1166$ shown on Figure 3-2.

Table D-1: As-Received Test Description and Results

Server	Node #	Processors Memory	Idle Mean SD	50% Mean SD	100% Mean SD	Performance (GFLOPS)	Efficiency (GFLOPs/ joule)
Intel #1	0-3	Intel #1 (Gold) 2650 v2	242 3.2	687 4.1	986 10.2	287.1	0.291
Intel #2	4-7	Intel #2 2650 v2	220 2.9	713 4.9	1037 13	285.4	0.275
Intel #3	8-11	Intel #3 2650 v2	235 2.4	706 3.6	1019 10.9	287.1	0.282
Dell	12-15	Dell 2670 v2	245 11.2	701 7.2	1050 10.7	358.2	0.341
SM	16-19	SM 2670 v0	351 3.6	882 5.9	1382 13.4	301	0.218

SD = standard deviation

Table D-2: Intel #1 Nodes Installed in Three Intel Server

Server	Node #	Processors Memory	Idle Mean, SD	50% Mean, SD	100% Mean, SD
Intel #1	0-3	Intel #1 (Gold) 2650 v2	242, 3.2	687, 4.1	986, 10.2
Intel #2	0-3	Intel #1 (Gold) 2650 v2	244, 2.3	684, 3.9	982, 9.6
Intel #3	0-3	Intel #1 (Gold) 2650 v2	245, 2.7	689, 3.6	989, 9.0

Table D-3: Intel #3 Components Installed in Intel #1 Server and Nodes

Server	Node #	Processors Memory	Idle Mean, SD	50% Mean, SD	100% Mean, SD
Intel #1	0-3	Intel #3	245, 3.1	717, 4.2	1030, 10.6
Intel #3	8-11	Intel #3	235, 2.4	706, 3.6	1019, 10.9
Intel #1	0-3	Intel #1 (gold) 2650 v2	242, 3.2	687, 4.1	986, 10.2

Table D-4: Intel #1 "Gold" Components Installed in Dell and Supermicro

Server	Node #	Processors Memory	Idle Mean, SD	50% Mean, SD	100% Mean, SD
Intel #1	0-3	Intel #1 (Gold) 2650 v2	242, 3.2	687, 4.1	986, 10.2
Dell	12-15	Intel #1 (Gold) 2650 v2	236, 10.3	637, 5.9	899, 8.5
Supermicro	16-19	Intel #1 (Gold) 2650 v2	304, 2.1	663, 3.6	933, 8.8

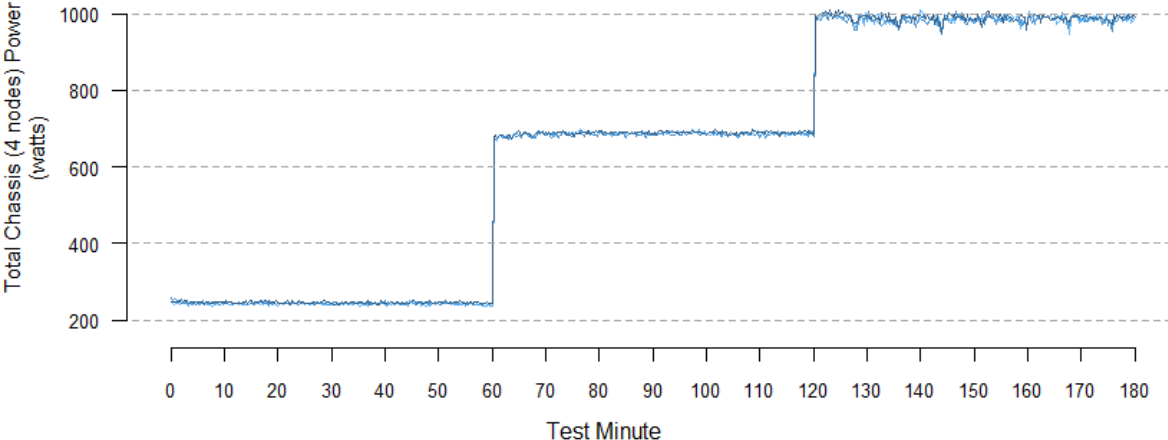
Table D-5: Intel #1 "Gold" Components Installed in Dell and Supermicro - Full Power Efficiency

Server	Node #	Processors Memory	100% Mean (watts)	Performance (GFLOPS)	Efficiency (GFLOPs/ joule)
Intel #1	0-3	Intel #1 (Gold) 2650 v2	986	287.10	0.291
Dell	12-15	Intel #1 (Gold) 2650 v2	899	287.45	0.320
Supermicro	16-19	Intel #1 (Gold) 2650 v2	933	285.26	0.306

Appendix E: Swap A Measured Power Data

Three Intel Servers Using One Set of Intel Nodes

Nodes from Intel Chassis #1 Installed In Intel Chassis #1 - #3
Three Load Levels: Idle, 50%, 100%, Summary Mean Values in Watts
Chassis #1: Idle: 242, 50%: 687, 100%: 986
Chassis #2: Idle: 243, 50%: 684, 100%: 981
Chassis #3: Idle: 245, 50%: 689, 100%: 989



Appendix F: Swap C Measured Power Data

All Brands - Use Key Processors and Memory Component Set

