Asia-Pacific Economic Cooperation

2012/SCSC/WKSP4/014

Considerations in Categorizing or Grading Server Systems

Submitted by: IBM

APEC RUSSIA 2012

Aligning Energy Efficiency Regulations for ICT Products: Developing a Strategic Approach
Seoul, Korea
18 July 2012

APEC Conference on Aligning Energy Efficiency Regulations for ICT Products II

July 18, 2012 – Seoul Korea

CONSIDERATIONS IN CATEGORIZING OR GRADING SERVER SYSTEMS

LIU HAI
IBM

AGENDA

- Server Product Power Signatures Vary By:
  - Processor Socket Count
  - Resilient and Scalable Characteristics
  - Range of Configurations Created with Components
- Complexity of Server Products Require Different Approach for Classifying and Testing Products for Energy Efficiency:
  - Limit the Number of Configurations Tested
  - Group by High Level Product Identifier: Machine Type or Model.

SERVERS HAVE HIGHLY VARIED CONFIGURATIONS AND CAPABILITIES:

- Processor Families:
  - Different Number of Cores: 2 to 8 (Likely Increases in future)
  - Different Socket Power, Dictated by:
    - Number of Cores; and
    - Resiliency and Capabilities Designed into the Processor
- Varied Number of Components:
  - Memory
  - Hard Disk Drives and Solid State Drives
  - Input/Output Cards
  - Graphics CPUs
- Resiliency, Availability and Scalability Characteristics Adds Power Use Through (see tables 3-5 in back-up for details on criteria):
  - Additional Circuitry and Capability in Processor
  - Additional Capability in the Memory
  - Redundant Power Supplies and Cooling Systems
  - Scalability Features which add circuitry
  - Green Grid has proposed a Resilient and Scalable Server Definition to ENERGY STAR to account for differences between Managed and Resilient and Scalable Servers.

System Configurations	**Minimum**	**Typical**	**Maximum**
Configuration ID	7377=xxx	7377-xxx	7377-xxx
Processor Information	2x Xeon X5675 3.06GHz 6C 95W	2x Xeon X5675 3.06GHz 6C 95W	2x Xeon X5675 3.06GHz 6C 95W
Memory Information	2 DIMMs, 2GB each	8 DIMMs, 4GB each	12 DIMMs, 16GB each
Internal Storage	1x 160GB 7200 SATA 2.5" SS HDD	2x 146GB 15K SAS 2.5" HS HDD	28x 146GB 15K SAS 2.5" HS HDD
I/O Devices	On-board dual 1Gb NIC	On-board dual 1Gb NIC	On-board dual 1Gb NIC, 2x pciE dual 1Gb NIC option
Power Supply Number and Redundancy Configuration	2, Redundant	2, Redundant	2, Redundant
Management Controller or Service Processor Installed?	YES	YES	YES
Other Hardware Features / Accessories	PCIe ServeRAID M1015 Controller	PCIe ServeRAID M1015 Controller	PCIe ServeRAID M1015 Controller

Power Data	**Minimum**	**Typical**	**Maximum**
Idle Category (1S and 2S only)	Category D: Managed Dual Installed Processor (2P) Servers		
ENERGY STAR Idle Power Allowance (1S and 2S only)	170.0	234.0	770.0
Measured Idle Power (watts)	144.4	175.7	350.7
Power at Full Load* (watts)	294.4	366.1	544.2
Benchmark / Method Used for Full Load Test	mPrime v25.9 [2]		
Test Voltage and Frequency for Idle and Full Load Test	230V/60Hz		
Range of Total Estimated Energy Usage ** (kWh/year)	2,530 to 5,158	3,078 to 6,414	6,144 to 9,534

Table 1: Range of Power for a 2 Socket Server System
Data from EPA ENERGY STAR Data Sheet

Server Product Family is Best Differentiated by:

- Maximum Number of Processor Sockets
  - Processor Power Characteristics of a System
  - More Processor Sockets Increases the Circuitry and the Range of Component Configurations.
  - The ENERGY STAR program has shown that the Number of Processor Sockets is a Good Means to Differentiate Servers.
- RAS and Scalability Characteristics of the Server:
  - Resilient and Scalable Servers Have a Higher Power Signature than Managed Servers
  - Resilient and Scalable Characteristics are Valued by Segments of Users in the Market.
    - Provide Increased Assurance of Operational Continuity
    - Resilient and Scalable Servers with a Given Socket Count will not be Competitive Against Managed Servers Due to Higher Base Power Load.

USE PRODUCT FAMILY DESIGNATION:

- One Size Fits All Approach is Not Appropriate for Enterprise ICT Equipment
- Use ENERGY STAR Product Family Definition for Testing Servers:
  - Testing Each Configuration is Impossible
  - Bracketing High and Low Power Configurations Provides Reasonable Assessment of Overall Product Capability
    - Recommend Use 3 Configurations for Testing for A Grading System.
  - Test/Evaluate a Subset of Configurations within a High Level Product Designator:
    - Machine Type; or
    - Model Number
- Server System Complexity and Wide Range of Power Signatures Necessitates a Unique Approach to Manage Product Energy Efficiency Evaluation
  - Enterprise ICT products are much more complex than appliances

CONSIDERATIONS FOR CREATING SERVER CATEGORIES

BACK-UP SLIDES

System Configurations	Minimum	Typical	Maximum
Configuration ID: See Qualified Configuration Section	7145-xxx	7145-xxx	7145-xxx
Processor Information	4, Intel Xeon X7560	4, Intel Xeon X7560	4, Intel Xeon X7560
Memory Information	8 Dimms, 4GB each	32 Dimms, 2GB each	64 Dimms, 16GB each
Internal Storage	1x 2.5 inch, 10kRPM SAS	3x 2.5 inch, 10kRPM SAS	8x 2.5 inch, 10kRPM SAS
I/O Devices	onboard dual 1 GB NIC	onboard dual 1 GB NIC + 2x PCIe dual 1GB NIC	onboard dual 1 GB NIC + 4x PCIe dual 1GB NIC + 1x 1-port Fibre HBA + 2x ServeRaid BR10i
Power Supply Number and Redundancy Configuration	2 redundant	2 redundant	2 redundant
Management Controller or Service Processor Installed?	YES	YES	YES
Other Hardware Features / Accessories	Internal SAS controller	Internal SAS controller	Internal SAS controller

Power Data	Minimum	Typical	Maximum
Idle Category (1S and 2S only)	N/A (3S or 4S)		
ENERGY STAR Idle Power Allowance (1S and 2S only)	N/A	N/A	N/A
Measured Idle Power (watts)	386.0	498.0	992.0
Power at Full Load* (watts)	670.0	888.0	1620.0
Benchmark / Method Used for Full Load Test	STREAM rev 5.8 / TRIAD (Linux - OpenMP compilation)		
Test Voltage and Frequency for Idle and Full Load Test	230Vac / 60 Hz		
Range of Total Estimated Energy Usage ** (kWh/year)	6,763 to 11,738	8,725 to 15,558	17,380 to 28,382
Link to Detailed Power Calculator (if available)	www.ibm.com/systems/bladecenter/resources/powerconfig.html		

Table 2: Range of Power for a 4 Socket Server System

Requirements	Characteristic	Resilient and Scalable Server	Managed Server
Must Have All 3	Processor RAS and Scalability:		
	1. Processor Error Correction	x	
	2. Interconnect with additional chipsets to add function.	x	x
	3. Provide High Bandwith I/O interfaces	x	
Must Have All 6	Memory RAS and Scalability		
	1. Extended Memory Correction ECC	x	x
	2. Failure Recovery of 2 fails in adjacent chips	x	
	3. Data Migration on Fail	x	x
	4. Memory Buffers to connect DIMMs to DDR channels.	x	
	5. Resilient links between processor and memory buffer.	x	
	6. Lane sparing in processor memory links.	x	
Must Have	Redundant Power Supplies	x	x
Must Have	Redundant Thermal Cooling Systems	x	x

Table 3:
RESILIENT AND SCALABLE SERVER CHARACTERISTICS:

1. Processor and RAS Scalability
2. Memory RAS and Scalability
3. Redundant Power Supply
4. Redundant Thermal Cooling System.

Requirements	Characteristic	Resilient and Scalable Server	Managed Server
Must Have 6 of 10	System Resiliency		
	1. Redundant Storage Controller or Paths to External Storage		
	2. Redundant Service Processors		
	3. Redundant DC-DC Regulator Stages		
	4. Supports Run-time processor Deallocation	x	
	5. Hot swappable I/O and HDD/SSD	x	x
	6. End to End Bus Retry on processor to memory or processor to processor links	x	
	7. Run-time expansion/retraction of hardware resources without reboot	x	
	8.Processor Socket Migration	x	
	9. Memory Patrol or Scrubbing for Error Detection and Correction	x	x
	10. Internal Storage resiliency with at least RAID V Controller		

Table 4: RESILIENT AND SCALABLE SERVER CHARACTERISTIC: SYSTEM RESILIENCY

Requirements	Characteristic	Resilient and Scalable Server	Managed Server
Must Have all 3	System Scalability		
	1. Higher Memory Capacity >=8 DDR3 or DDR4 DIMM Ports per socket	X	
	2. Logical Partitioning Support	X	X
	3. Greater I/O Expandability	X	X

Table 5:
RESILIENT AND SCALABLE SERVER CHARACTERISTIC: SYSTEM SCALABILITY

VARIED POWER USE FOR 4 PROCESSOR SOCKETS SERVERS :

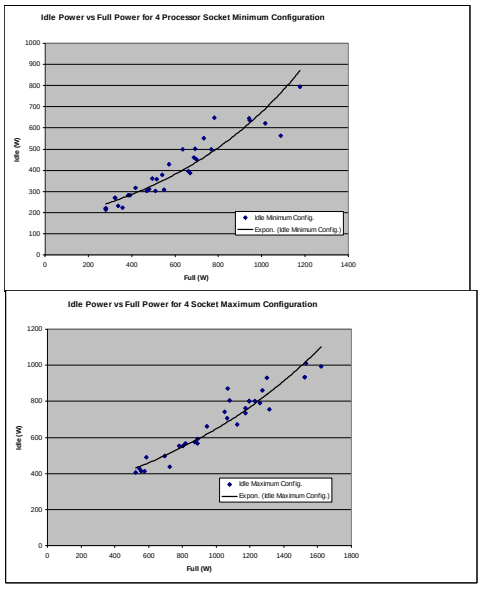

- Depending on Configuration and Architecture, Different Products Have Highly Varied Idle Power.
- A Low Idle Power System May Not be the Best Fit for Specific Applications:
  - High Performance Computing
  - Virtualization
  - Specific Application Types

Data from ENERGY STAR® qualified 4 processor socket systems

Lower Power Does Not Mean Higher Efficiency (4 Socket Comparison)

Individual Server A	
Max Power	663 W
Idle Power	336 W
Capacity	**334.3 Kbops/system**

Individual Server B	
Max Power	715 W
Idle Power	358 W
Capacity	**411.7 Kbops/system**

Low Power & High Power Server Options to Meet IT Requirement of 5.0 Mbops

Rack A		Rack B	
Capacity	5.02 Mbops	Capacity	5.35 Mbops
Systems	15	Systems	13
Total Idle	5.04 kW	Total Idle	4.65kW

Choosing Server B is MORE Efficient
Saves 4.54 MWhrs to 3.73 MWhrs per yr.**

Performance is Key to Server Energy Efficiency

** savings based on 50% utilization to 15% utilization. Does NOT include HVAC or environmental controls

Source: Intel Test Labs

IMPACTS ON INNOVATION AND CONSUMER CHOICE:

- Favors Low Power Systems:
  - High Power Systems Provide Better Overall Energy Efficiency and Total Workload Delivered per Unit of Energy Consumed.
  - High Power Systems Necessary for Some Applications such as:
    - Graphics
    - High Performance Computing
- Limits System Productivity Growth
  - Energy Efficiency is a Function of Attributes Beyond Simple Power Use
  - Innovation and Productivity Depends on Optimized Improvement of All Attributes
- Virtualized Systems Often Require Higher Energy Footprint
  - Increased System Utilization Increases Efficiency for Power, Space, and Material Use
- Data Center Systems Availability and Reliability Require Higher Energy Footprint and System Infrastructure
- Overly Prescriptive Requirements Regulate Yesterday's Technologies and Delay or Prevent Introduction of Tomorrow's Innovation.