23-06-2014

ITIC ANALYSIS OF SERT® WORKLET RESULTS

Jay Dietrich, IBM; jdietric@us.ibm.com

ITI PRELIMINARY ANALYSIS OF SERT DATA: AGENDA

- General Relationship of Performance, Power, and Server Configuration Under SERT
- Impact of Components on SERT worklet scores by component type:
  - Overall Correlations
  - CPU
  - Memory
  - Storage
  - Impact of CPU and Java Capabilities on the Crypto Score
- How should Configuration Types be Defined?
- Data Needs
- Future Work

NUMBER OF SERVER TYPES IN THE SERT DATABASE

Server Type	Form Factor	# of Sockets	# of Machine Types/Models	# of Configurations
Unmanaged				
	Rack	1	2	10
	Tower	1	3	15
Managed				
	Blade	2	9	47
	Blade	4	3	13
	Rack	1	7	35
	Rack	2	18	79
	Rack	4	2	10
	Tower	1	3	15
	Tower	2	3	15
Resilient	Rack	2	3	24
	Rack	2	3	21

- **Dataset numbers as of June 20, 2014**
- **Majority of the Machine Type/Models are ENERGYSTAR® Certified**
- **ITI is working with it's members and EPA to keep the data set up to date; it has been shared with China CNIS, EU Consultant for Lot 9, JEITA, and Korea TTA.**

THE RELATIONSHIP OF PERFORMANCE AND POWER IN SERT

- SERT Worklet scores are a ratio of measured performance to measured power over a 4 or 8 power use intervals.
- Adding component(s) to a system introduces:
  - Additional performance
  - Additional Power Use
- The basic equation for assessing the SERT worklet score is:

$$\frac{\text{(Baseline Perf + Perf-change with config-difference)}}{\text{(Baseline Active Power + Power-change with config-difference)}}$$

In turn, this can be converted to:

$$\frac{\text{(1 + (Perf-change/Baseline Perf))}}{\text{(1 + (Power-change/Baseline Power))}} * \frac{\text{Baseline Perf}}{\text{Baseline Power}}$$

- Given the limits of the SERT worklets themselves, we find that:
  - In the low end configurations; additional components typically add performance benefits.
  - In the typical and maximum configurations, some components will add only power debt on a given worklet type

IMPACT OF COMPONENT TYPES ON THE SERT WORKLET SCORES

CORRELATION OF WORKLET SCORES

	CPU							Memory		Storage		Hybrid	Power	
	Compress	CryptoAES	LU	SOR	XML Validate	SORT	SHA256	Flood	Capacity	Sequential	Random	SSJ	Maximum Power	Idle Power
Compress	1.00	0.31	0.90	0.89	0.95	0.86	0.91	-0.20	-0.16	0.33	0.35	0.94	-0.10	-0.35
CryptoAES	0.31	1.00	0.50	0.18	0.29	0.21	0.30	-0.20	-0.17	0.28	0.12	0.36	-0.18	-0.30
LU	0.90	0.50	1.00	0.89	0.94	0.86	0.91	-0.25	-0.18	0.31	0.28	0.88	-0.19	-0.39
SOR	0.89	0.18	0.89	1.00	0.92	0.97	0.94	-0.11	-0.05	0.26	0.27	0.83	-0.05	-0.25
XML Validate	0.95	0.29	0.94	0.92	1.00	0.89	0.95	-0.24	-0.19	0.27	0.29	0.92	-0.17	-0.36
SORT	0.86	0.21	0.86	0.97	0.89	1.00	0.94	-0.07	-0.02	0.23	0.24	0.84	-0.03	-0.19
SHA256	0.91	0.30	0.91	0.94	0.95	0.94	1.00	-0.17	-0.11	0.27	0.27	0.86	-0.17	-0.36
Flood	-0.20	-0.20	-0.25	-0.11	-0.24	-0.07	-0.17	1.00	0.94	0.02	0.09	-0.23	0.53	0.65
Capacity	-0.16	-0.17	-0.18	-0.05	-0.19	-0.02	-0.11	0.94	1.00	0.06	0.10	-0.18	0.46	0.58
Sequential	0.33	0.28	0.31	0.26	0.27	0.23	0.27	0.02	0.06	1.00	0.84	0.31	-0.09	-0.18
Random	0.35	0.12	0.28	0.27	0.29	0.24	0.27	0.09	0.10	0.84	1.00	0.33	0.04	-0.08
SSJ	0.94	0.36	0.88	0.83	0.92	0.84	0.86	-0.23	-0.18	0.31	0.33	1.00	-0.12	-0.33

High positive correlation
Medium positive correlation
Low positive correlation
No correlation
Low negative correlation

- **CPU Worklets:**

CORRELATION BETWEEN WORKLETS FOR 2 SOCKET SERVERS

	Compress	Crypto	LU	SOR	XML	SORT	SHA256	Flood	Capacity	Hybrid ssj	Maximum Power	Idle Power
Compress	1.00	0.38	0.86	0.88	0.90	0.87	0.86	-0.03	0.03	0.87	-0.14	-0.35
Crytpto		1.00	0.58	0.23	0.32	0.28	0.30	-0.13	-0.02	0.21	-0.18	-0.29
LU			1.00	0.90	0.92	0.89	0.87	0.05	-0.02	0.81	-0.23	-0.37
SOR				1.00	0.94	0.98	0.93	0.15	0.11	0.76	-0.06	-0.23
XML					1.00	0.94	0.96	-0.01	-0.04	0.83	-0.25	-0.39
SORT						1.00	0.95	0.10	0.09	0.81	-0.05	-0.23
SHA256							1.00	0.07	0.04	0.76	-0.08	-0.25
Flood								1.00	0.87	-0.01	0.72	0.56
Capacity									1.00	-0.02	0.77	0.34
Sequential										1.00	-0.21	-0.39
Maximum Power											1.00	0.88

- (green) High Positive; >=.80
- (blue) Medium Positive; .4<correl<.8
- (yellow) Low Positive; .1<correl<.4
- (red) No Correlation; -.1<correl<.1
- (orange) Low Negative Correlation; <-.1

Correlations for the two socket servers roughly mimic the correlation of the overall database.

It suggests that the impact of "dead power" in the maximum configurations manifest itself similarly in all server types.

CPU WORKLETS

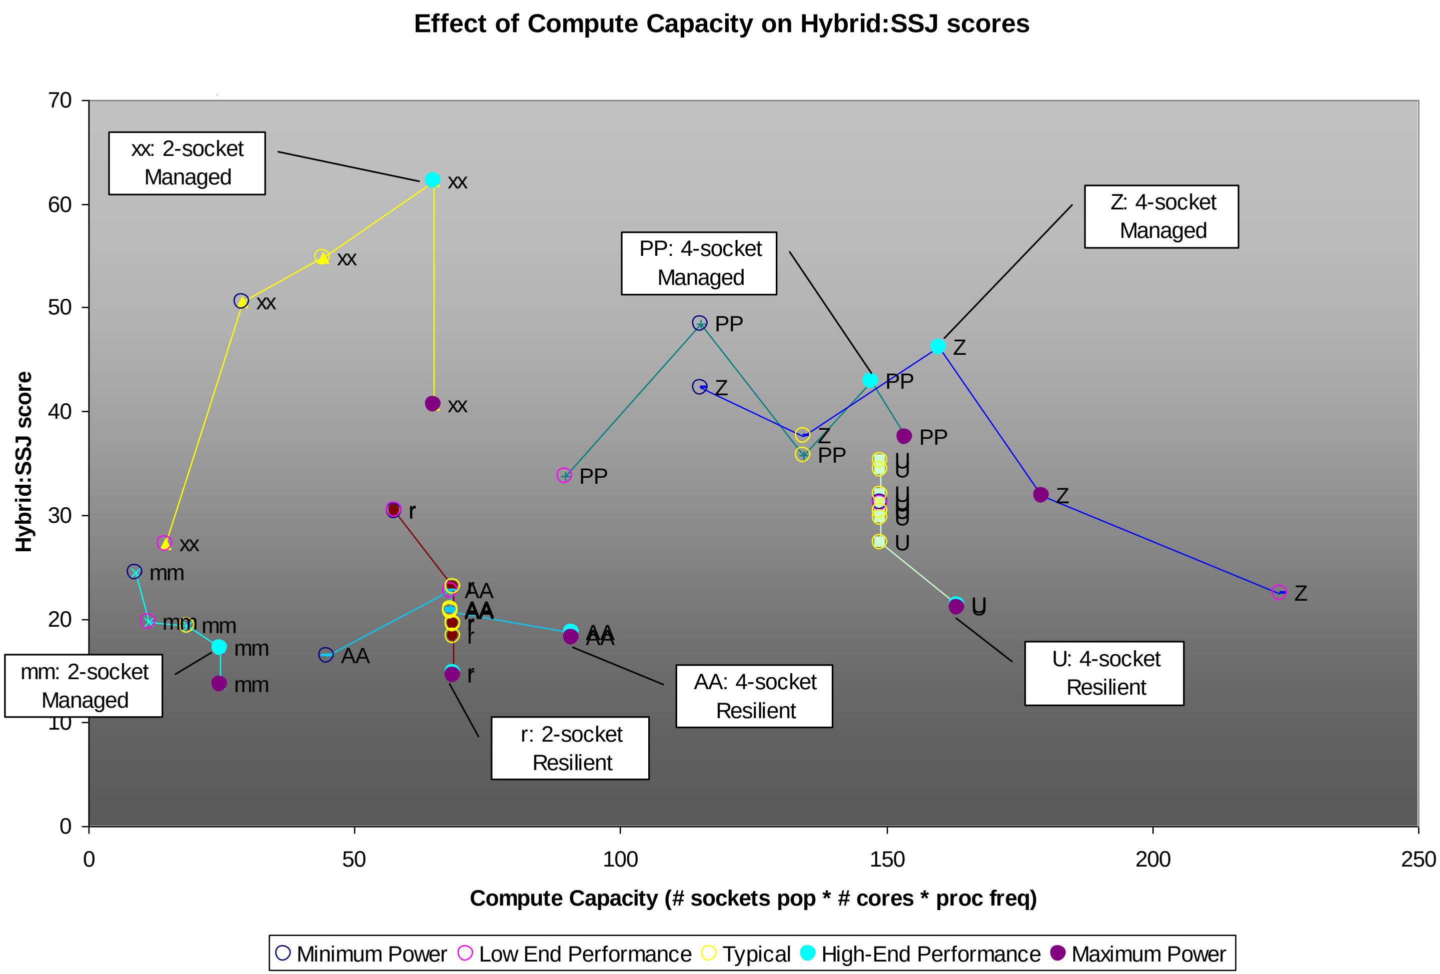

CPU WORKLETS: ACROSS CONFIGURATIONS

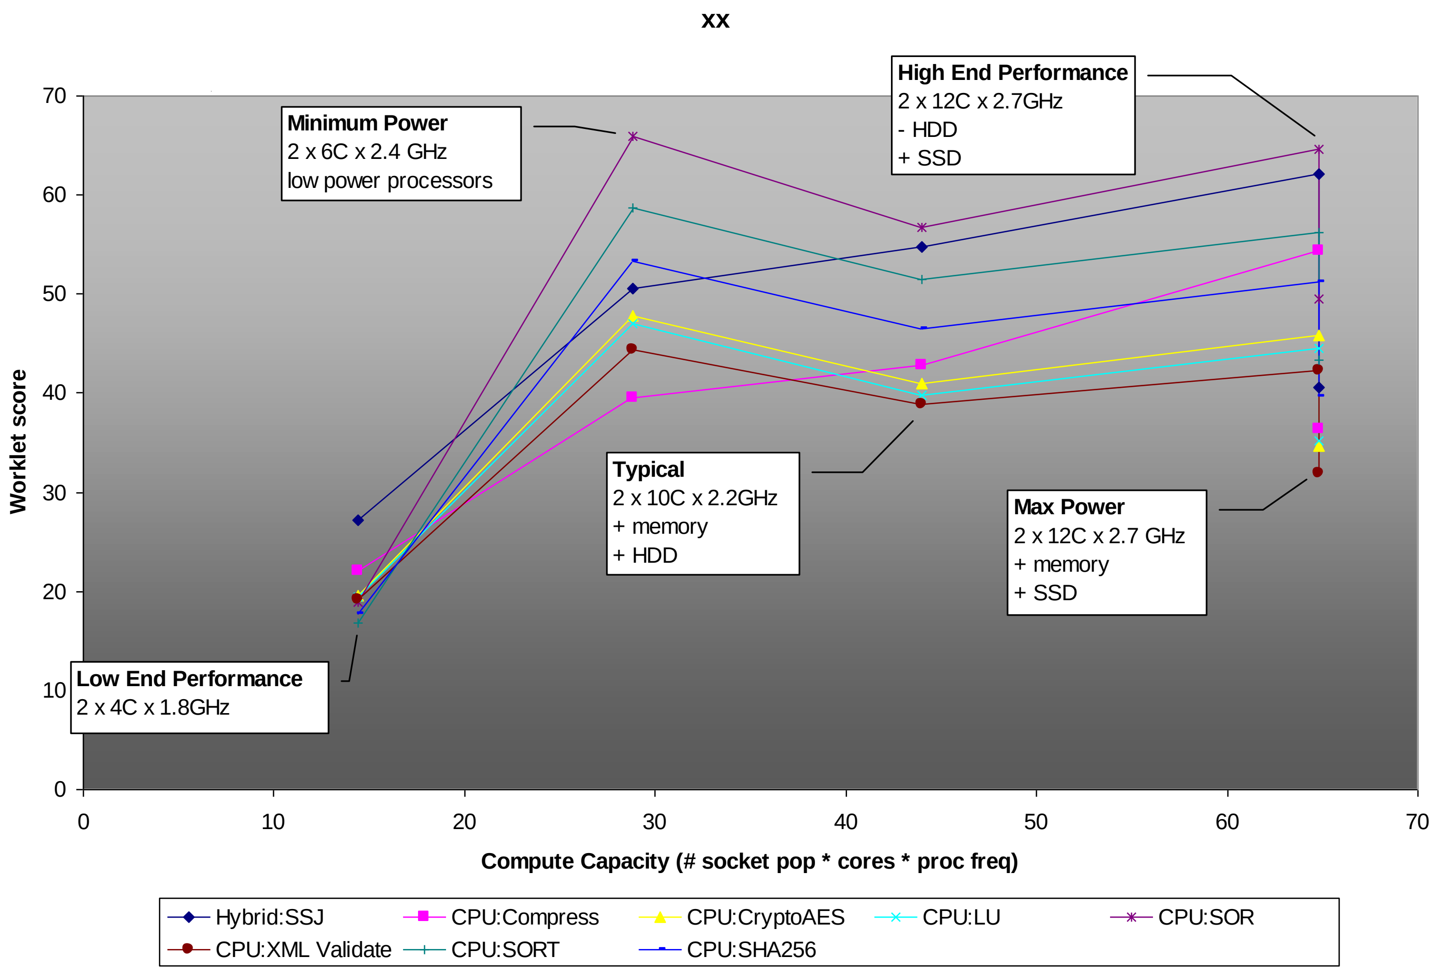

MEMORY SYSTEM ANALYSIS

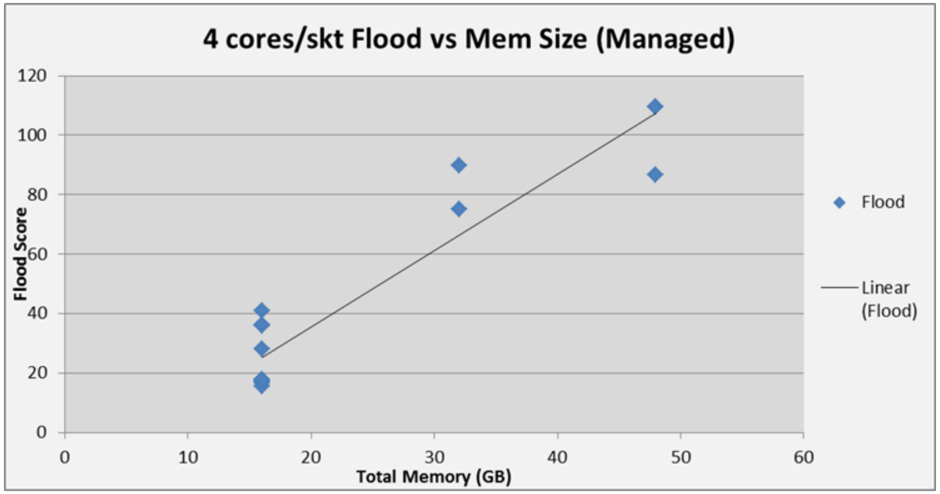

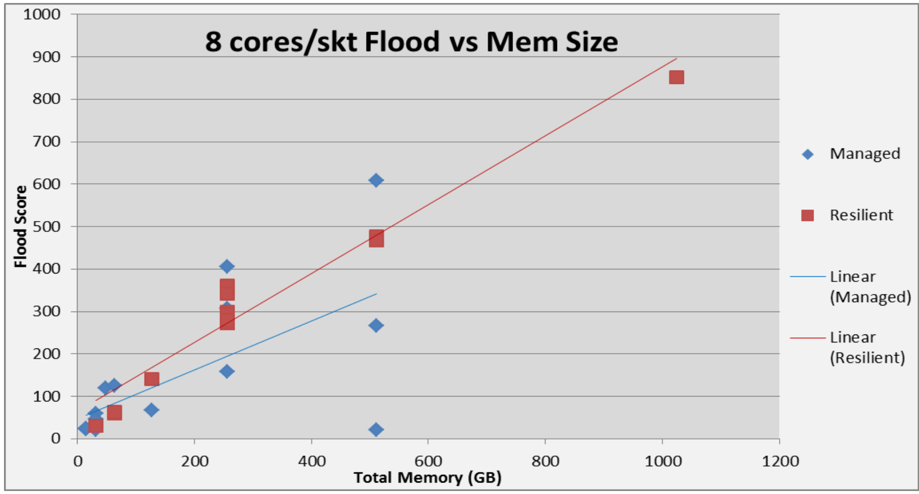

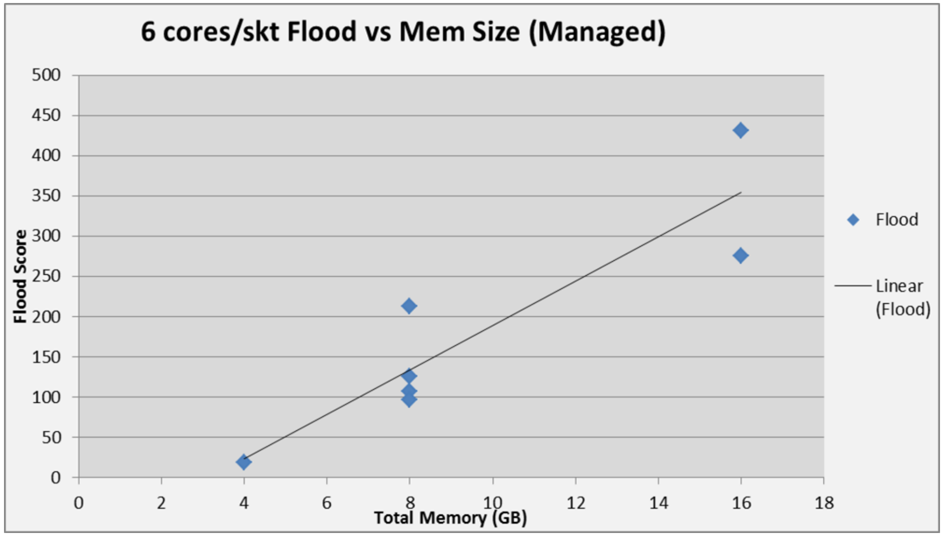

Observations:

- **Flood and Capacity worklets have 94% correlation.**
- **Flood scores increase linearly with increasing GB.**
- **Significant score distribution at GB points. Need to consider impact of memory type, and chip and DIMM size on score.**

Information Technology Industry Council

Innovation. Insight. Influence.

IMPACT OF GB OF MEMORY ON CPU SCORE

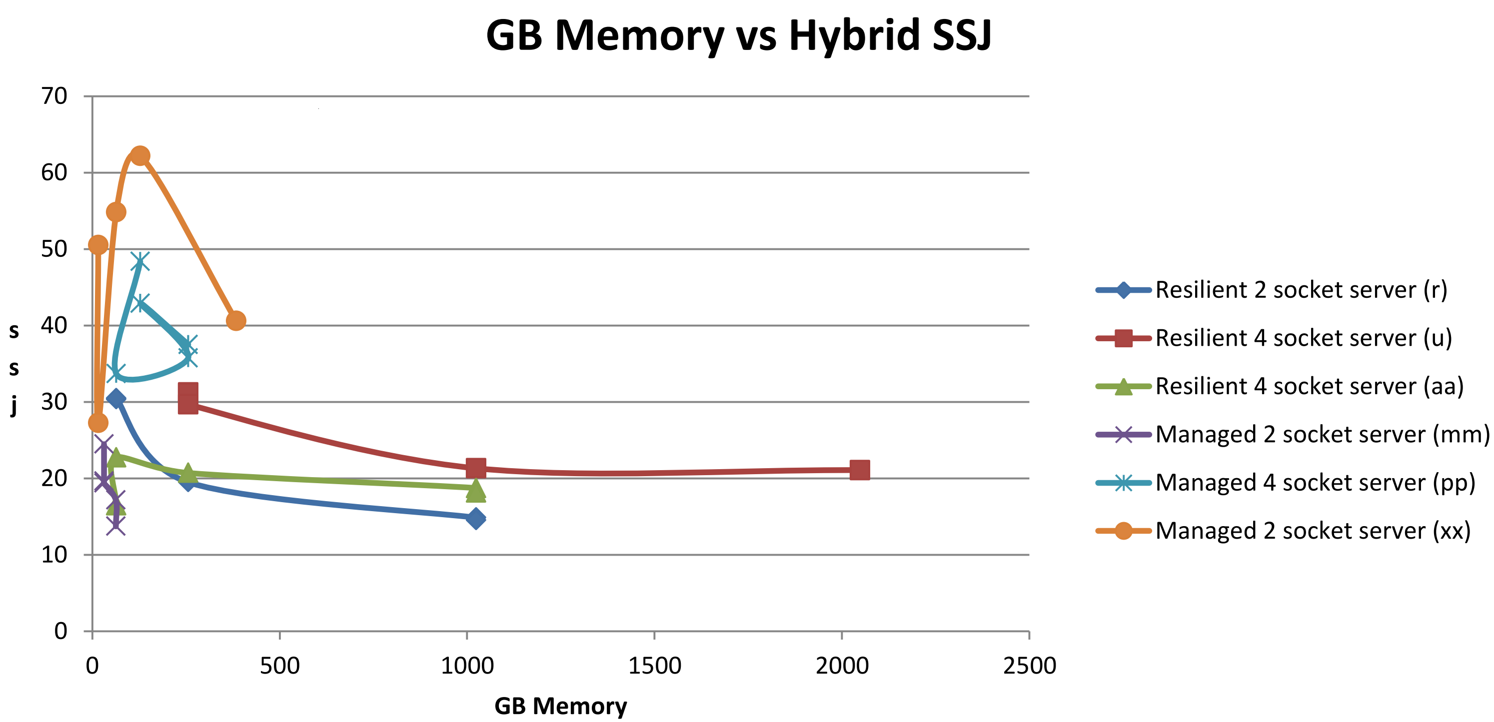

Observations:

Individual Configuration lines flow from minimum power to maximum power configurations

Added Components in maximum configurations reduce CPU worklet scores

Choice of components can create different outcomes at each configuration point.

STORAGE WORKLETS: GENERAL OBSERVATIONS

- SSDs have higher Sequential and Random Worklet Scores than HDDs.
  - SSD systems have higher random than sequential scores.
  - HDD systems have higher sequential than random scores.

- Storage scores have low or no correlation to CPU and memory worklet scores.

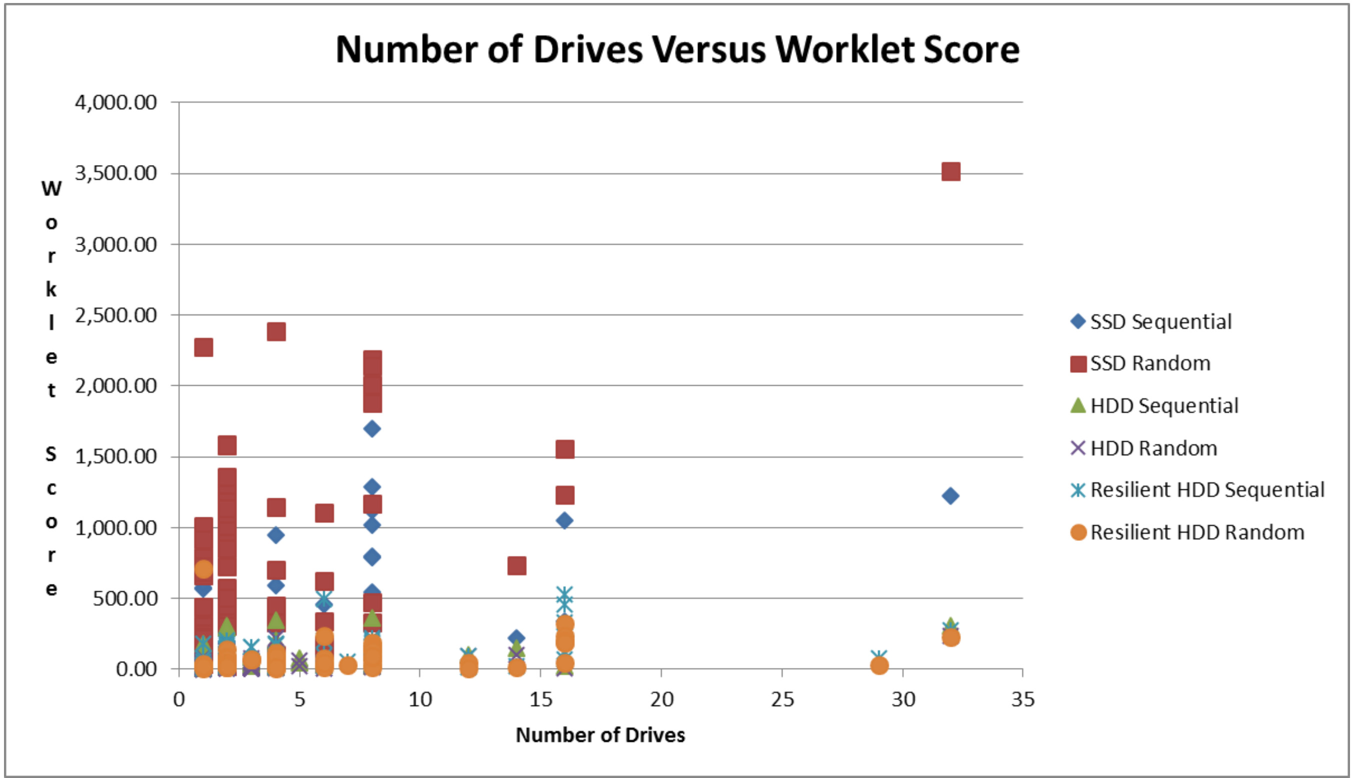

CORRELATION OF # AND SPEED OF DRIVES TO STORAGE WORKLET SCORES

	2 Socket Managed Servers				Total Server Dataset			
	HDDs		SSDs		HDDs		SSDs	
	# of Drives	Drive Speed	# of Drives	GBs	# of Drives	Drive Speed	# of Drives	GBs
Sequential	0.3437187	0.52270517	0.7001124	0.294116	0.4020192	0.1314018	0.5303168	0.307528
Random	0.3476777	0.16055263	0.8121809	0.204996	0.3674241	0.0389353	0.5711544	0.175033

OBSERVATIONS:

- **# of SSDs has a strong correlation to storage scores for 2 socket servers.**
- **# of HDDs has a slight correlation to storage scores for 2 socket servers and the total DB.**
- **Drive speed has medium correlation to the sequential scores on two socket systems.**
- **Overall database has significantly lower correlations, indicating a breakdown based on server type or number of sockets.**

RELATIONSHIP OF STORAGE AND CPU WORKLET SCORES

Correlation Scores for 2 socket managed servers & total server dataset

	2 socket managed servers				Total Server Dataset			
	SSDs		HDDs		SSDs		HDDs	
	CPU Geo	Hybrid SSJ	CPU Geo	Hybrid SSJ	CPU Geo	Hybrid SSJ	CPU Geo	Hybrid SSJ
Sequential	0.192747	0.192747	0.002395	0.139552	0.357595	0.381178	0.16678	0.125642
Random	0.032139	0.41115	-0.11367	-0.07491	0.304261	0.436578	0.090442	0.086535
Hybrid ssj	0.771691		0.891555		0.853552		0.925653	

- **No identifiable correlations between the CPU and Storage worklet scores.**
    - **Overall correlation chart showed low positive correlation between CPU and storage worklet scores.**
    - **Correlations for the 2 socket, managed servers are lower than the overall data set.**
    - **SSDs CPU GeoMeans have a lower correlation to hybrid ssj than HDDs.**

- **Indicates a differentiation between storage and CPU tests.**

IMPACT OF DRIVE COUNT ON CPU SCORES

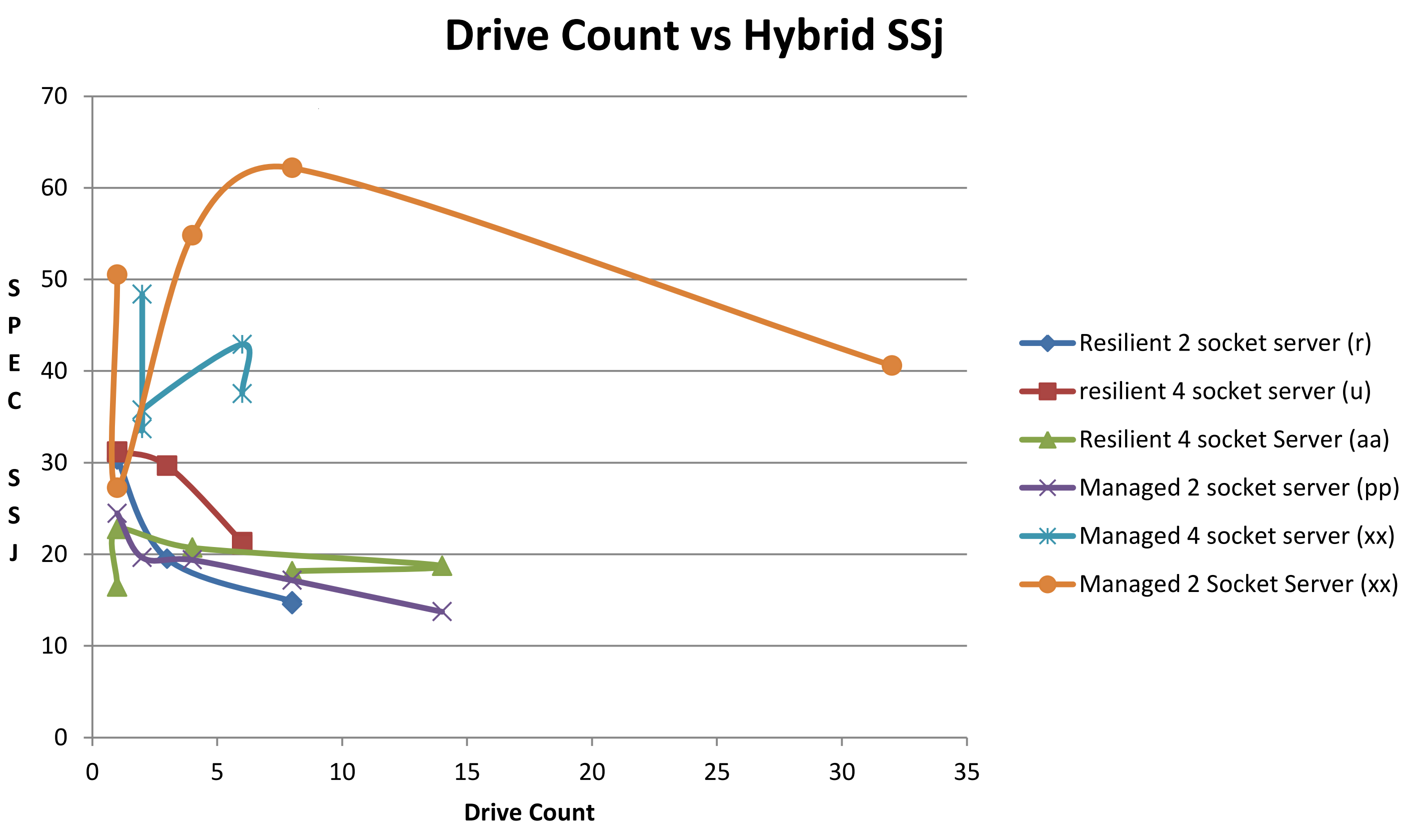

Observations

- The drives within a machine type are a mix of SSD and HDD
- Higher drive counts reduce Hybrid ssj scores, but is also a function of added memory.
- Configuration lines move from Minimum power to Maximum Power

ANALYSIS OF CRYPTO SERT WORKLET

Family	# Configurations with CryptoAES < 69 (range of score)	# Configurations with CryptoAES>69 (range of score)	Comments
a	1 (21)	4 (98-123)	It is possible that the low-end processor that got the low result does not support hyperthreading/options that makes CryptoAES better
b	1 (19)	4 (102-124)	Same as above
d	4 (19-43)	1 (214)	Same processor reached score of 43 and 214
oo	2 (21-23)	3 (92-138)	Seems to use Opteron processors
qq	4 (19-44)	1 (110)	Seems to use Intel processors
rr	3 (25-53)	2 (99-150)	Seems to use Intel processors

Crypto Score provides a look at potential score impacts from:

- **Hardware based: hyper-threading**
- **Software based: jvm accelerator**

Need to consider how you want to represent system efficiency versus hardware efficiency.

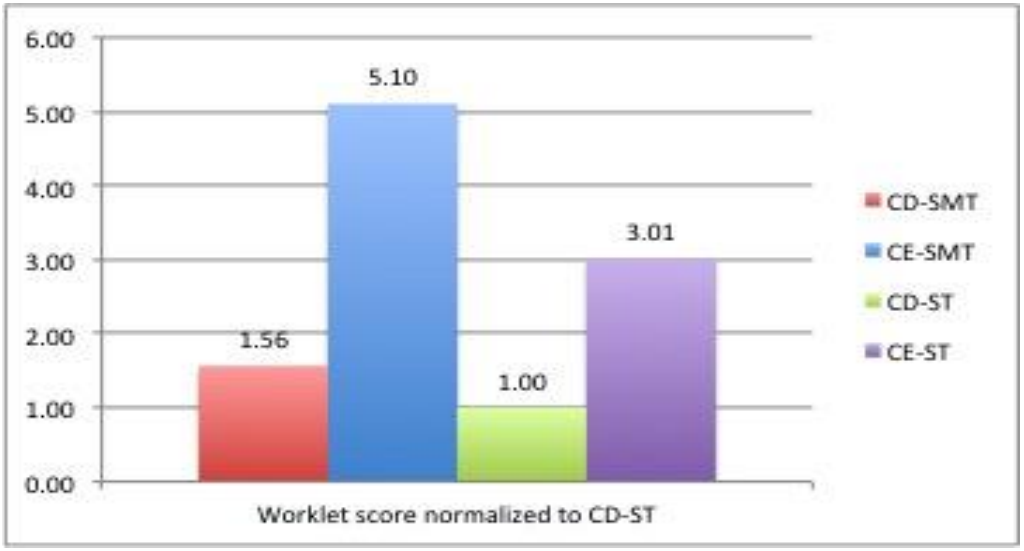

Minimum Memory Requirements for Larger Systems

- Theoretical Minimum System Cannot be Tested Using SERT
  - SERT optimal performance requires between 512 MB and 1 GB of memory per logical core (thread)
  - A 2-socket, 12-core/processor (24 cores total) server with 4 threads/core requires a minimum of 96 GB of memory
- Resilient Servers and 4 socket servers are rich in memory
  - Approximately 90% of servers shipped have ≥ 128 GB of memory
  - Customer workloads drive memory requirements

COMPLICATIONS IN SELECTING THE MAXIMUM CONFIGURATIONS

- Usage of V2.0 by IT Equipment Manufacturers has Demonstrated the Need for Clarity of the Product Family Definitions
  - Both high-end performance and maximum power definitions suggest heavily populated machines
  - Performance and power consumption are proportional for components such as processors and memory
- Very Difficult to Achieve High Performance on All SERT Worklets compared to Maximum Power
  - Different manufacturers chose to maximize performance on different worklets.
  - Choice of configurations was also influenced by need to create the “envelope” of certified configurations for a given machine type/model.
  - The number of options for a given machine type and model are a limiting factor
- An Analysis of the Data Set Suggests that Different Manufacturers Used Different Criteria in Their Selection of Configurations.

HOW ARE THE LOW AND HIGH PERFORMANCE CONFIGURATIONS DEFINED?

- Performance/Power relationship is optimized at different points for different worklet types:
  - CPU worklets deliver their best performance with an optimal amount of memory and storage.
  - The best Memory worklet scores are at the high end of GB capacity for a given machine type.
  - The best storage scores depend on an optimal number of drives matched to the capabilities of the processors.
  - The performance and power use of a server are highly dependent on the specific mix of components: Technologies, # of threads, interface speeds, etc.

- Should configurations be matched to worklet types?
  - Maximum performance to memory score(s)?
  - Typical configuration to maximize storage score(s)?
  - Typical configuration to optimize CPU score(s)?
  - Minimum performance to assess all worklet types:
    - Optimum memory GB: 64-256 GB depending on server type.
    - Consider HDDs and SSDs separately?

- Is Additional Segregation of the Dataset Warranted?

SELECTING TEST CONFIGURATIONS

- Eliminate Maximum Power Configuration
- A machine type/model should be able to be defined by 2-4 configurations:
  - Specified Configurations
  - Specified Minimum Configuration with 2-3 manufacturer selected configs.
  - A normalization scheme: needs to be defined.
- SERT Data complexity in the face of server configuration complexity illustrates that any performance/power metric for servers will have limitations in how it is used.

DATA NEEDS

- Memory Architecture (# of channels) and idle data:
  - SERT Scores likely affected by DIMM size, number of chips and technology.
  - Idle power adder will be dependent on same: need to validate.
- Presence and size of SSD/DRAM cache on HDDs
- jvm type and version

FUTURE WORK: Page 1

- Evaluate how memory idle power changes with technology and DRAM size:
  - Get IDDq6 data from manufacturers
  - Evaluate change in power use from change in DIMM and DRAM size
- Does drive interface type affect SERT score?
- How should I/O be integrated into the analysis scheme:
  - Absence of >10 GB ports.
  - Minimum # of 1<x<=10 GB Ports
- Analysis of Blade data:
  - Should a chassis level score be created?
  - How do you assess number of blades per chassis, differences in shared overhead.

FUTURE WORK: page 2

- Impact of Software on score:
  - Specification of software stack can create problems:
    - Change in version and unsupported software
    - Functionality is increasingly delivered through software
    - Different OS require different supporting software
  - SERT Requirements for Software
    - Declare Software Used
    - Tuner per SERT Requirements
- Idle Analysis:
  - Correlations to Worklet Scores: None evident
  - Scale Idle using Specific SERT Scores (Undefined)
  - Use of Adders
  - Assessing Idle Power Relationships on larger systems.

Information Technology Industry Council

Innovation. Insight. Influence.