


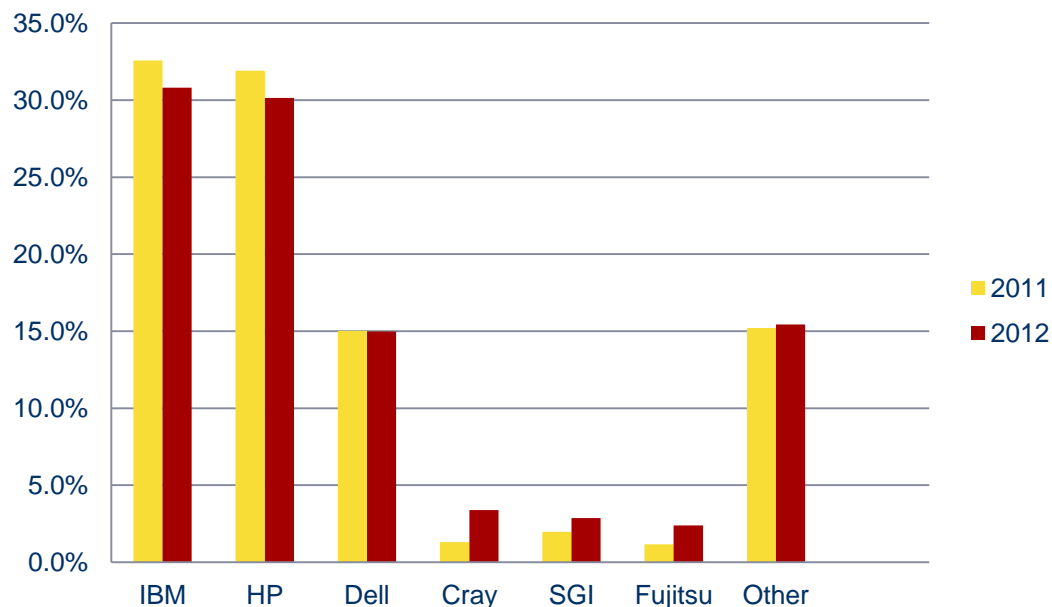
HPC Market Update: 2012 Second Quarter

WW HPC Server Market Revenues

Q2 2012 WW Broader Market



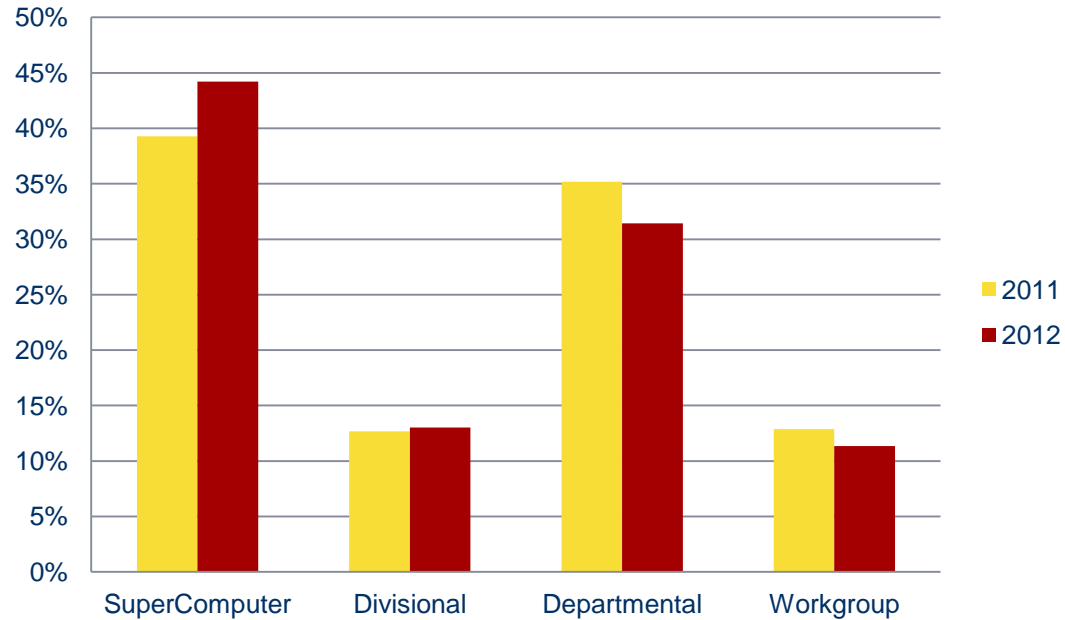
2011 vs 2012 Half Year Comparison



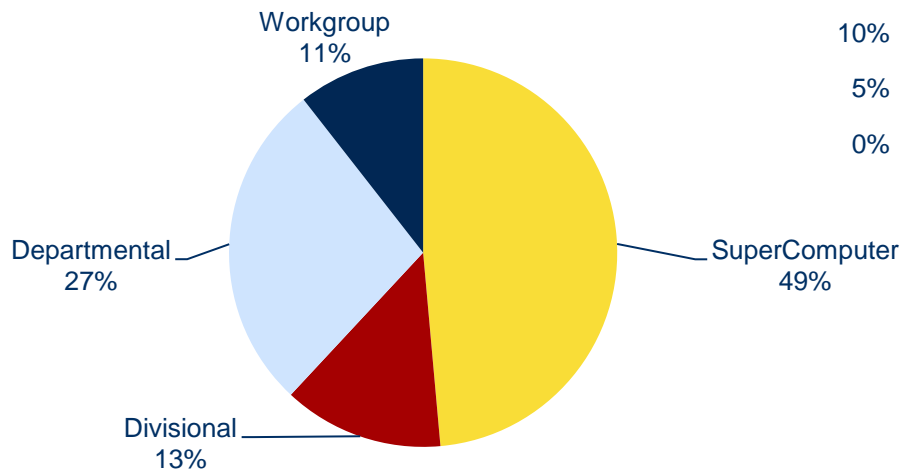
In Thousands	2011	2012
1/2 yr Rev.	4,912,964	4,875,688

Competitive Segments: Half-Year Comparison

2011 vs 2012 Competitive Segments



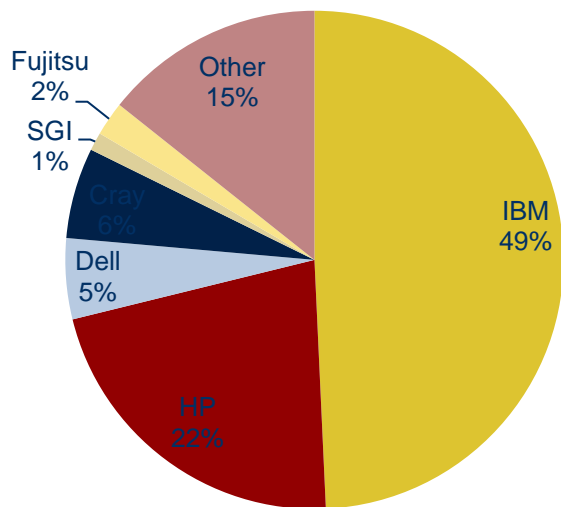
Q2 2012 Competitive Segments




In Thousands	2011	2012
1/2 yr Rev.	4,912,964	4,875,688

Supercomputer Segment: Half-Year Comparison

Q2 2012




2011 vs 2012 Supercomputer Segment




	2011	2012
1/2 yr Rev.	1.9 Billion	2.1 Billion

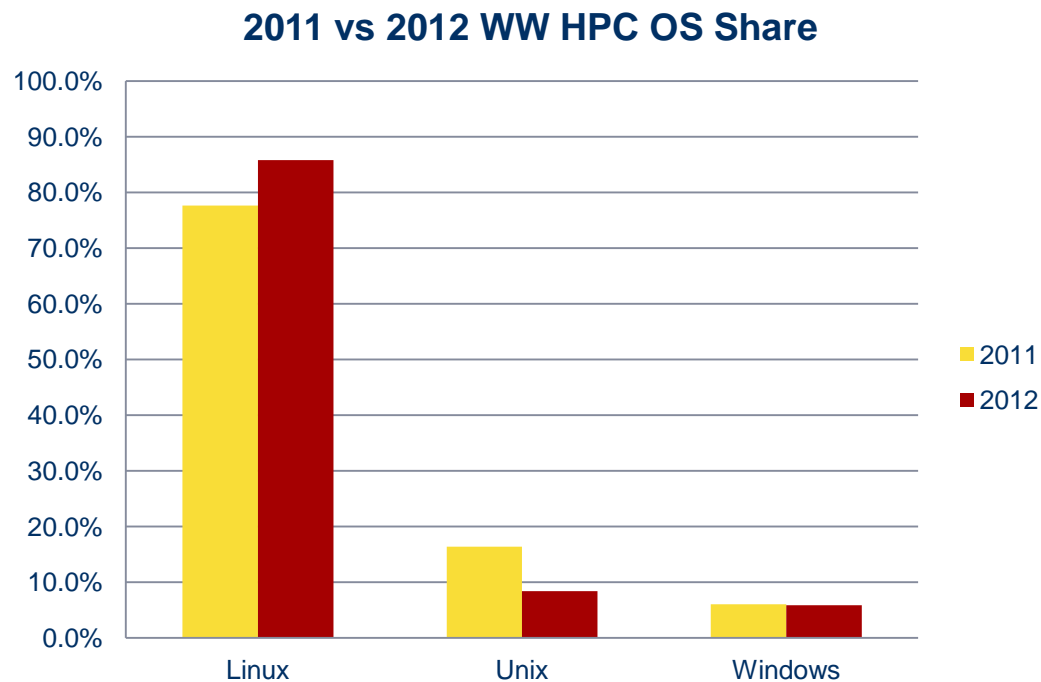

Base Processors: x86 Dominates




2011 vs 2012 Processor




WW HPC Server Revenue By OS: Linux Grows!



Regional HPC Server Revenue Shares: Half-Year 2012



2011 vs 2012 Regional Split by Segments



HPC and Big Data

- IDC forecasts that revenue for HPC servers acquired primarily for Big Data use will grow robustly (10.4% CAGR) to approach \$1 billion in 2015.
- Despite rapid growth, Big Data systems will have only a single-digit share (~7%) of all HPC server revenue in 2015, because of the small starting point.
- The larger Big Data ecosystem (servers, storage, networking, software) will approach \$2 billion in 2015.


IDC Worldwide Data Intensive (Big Data) Focused HPC Server Revenues

(\$ Millions)

	2009	2010	2011	2012	2013	2014	2015	CAGR '10-'15
WW HPC Server Sales	8,637	9,504	10,034	10,564	11,397	12,371	13,485	7.2%
Big Data Workloads	535	603	655	708	786	881	989	10.4%
Big Data in HPC Portion	6.2%	6.3%	6.5%	6.7%	6.9%	7.1%	7.3%	3.0%

Source: IDC 2012

The Boundaries Between HPC Big Data and High-End Commercial Analytics Are Dissolving



- HPC vendors are targeting commercial markets (IBM, HP, Dell, Cray, SGI, Intel, AMD, EMC, et al.).
- Commercial vendors are seeing HPC requirements (SAS, SAP, Oracle, LexisNexis, et al.)
- Early crossover use cases:
 - Fraud/error detection
 - Genomics/proteomics
 - Personalized medicine
 - FSI

Capacity “Big Data” Work

- Most jobs are here
- Big job is partitionable into independent small jobs.
- No global memory space
- Great SW can't overcome not-great HW
- Standard clusters:
 - Physically distributed memory
 - Logically distributed memory
 - Standard memory size, speed
 - Standard bandwidth, latency



Capability “Big Data” Work

- Toughest jobs are here
- Big job is not easily partitionable
- Need global memory space
- Great SW can exploit great HW
- Turbo-charged clusters:
 - Physically distributed memory
 - Logically shared memory
 - Better memory size, speed
 - Better bandwidth, latency



Mayo Clinic: Outcomes-Based Medical Diagnosis and Treatment Planning

- **Enter the patient's history and symptomology.**
- **While patient is still in the office, sift through 10 million archived patient records for relevant outcomes.**
- **Provider considers the efficacies of various treatments for "similar" patients (but is not bound by the findings).**
- **Ergo, this functions as a powerful decision-support tool.**
- **Benefit: learn from much larger patient populations than in typical medical studies.**
- **Could become very economically important**



Apollo Group/University of Phoenix: Student Recruitment and Retention

- **University of Phoenix: 280,000 students and growing**
- **Must target millions of students to produce this yield.**
- **Also tracks student performance for early identification of potential dropouts – “churn” is very expensive**
- **Solution: sophisticated, cluster-based Big Data models**



Global Courier Service: Fraud/Error Detection

- **Check 1 billion-plus packages per hour in central sorting facility**
- **Benchmark won by unannounced vendor with a turbo-charged interconnect and memory system**



- **Most HPC and Big Data utilize same / similar hardware substrates**
 - Interconnects: Infiniband, 10 GbE, 1 GbE,
 - Core Processors; Xeon, Opteron, SPARC64, Co-Processors / GPU, FPGAs, Massive Multithreading Processors
 - Storage Infrastructures: SSD, Flash augmented SANs to accelerate metadata
- **Big Data Workflow execution pattern common to several HPC scenarios**
 - Software frameworks in HPC are highly tuned to executing data-intensive workloads
- **Cross pollination of technologies and ideas between Big Data and HPC**
 - File Systems : GPFS and Lustre from HPC are being inherited into several Big Data System Architectures



- **HADOOP – ubiquitous Map Reduce processing framework**
- **Big Data processing becoming increasingly complex requiring capabilities beyond passive batch processing**
- **Stream processing and live processing of data at scale are driving innovation of new Big Data systems technologies**
- **Data Movement: A Persistent Challenge**
 - Moving Work to Data (minimizing data movement)
 - Moving Data to Work (speeding up data movement)
- **New Software technologies that can adaptively handle diverse workloads needed.**
 - Simplifying Programmability by exposing a common API to developers
 - Using Adapters to invoke different execution modalities

2012 WW HPC End-User Multi-Client Study

IDC 2012 HPC research reports combine surveys of end-users / buyers with our supply-side data for a complete view of HPC market

For the first time we will cover the fast evolving HPC Big Data Market

Reports planned from the study:

- HPC industries / applications workloads report
- HPC systems software and middleware report
- HPC storage and interconnect report
- HPC processors / co-processors / accelerators
- The emerging synergy of HPC and Big Data report

- How fast users are growing their systems
- How installed systems evolve over time with added/upgraded OEM equipment , as well as after-market purchases of other hardware/software
- To track how different industries are growing their use of HPC servers
- How buyers are changing the way they make their purchase decisions
- How Big Data trends are impacting HPC centers, including the ways in which HPC sites are addressing Big Data needs, and the convergence of HPC Big Data and high-end commercial analytics
- Identify new market trends and momentums, e.g., accelerators, clouds
- Show how vendors can better deliver to user requirements

\$15K per report. All 5 reports discounted at \$60K.

In addition, the 2012 IDC HPC end-user MCS reports will include the IDC five year market forecasts for each area and for HPC servers overall.

Opportunity to add custom questions to the study

- Interviews start in first week of October
- Order by October 1: add 1-2 custom questions for all interviews
- After October 1: add 1-2 questions to remaining interviews

IDC plans to complete reports in year-end 2012 time frame.