

The Rise of Blade Computing

By Anil Vasudeva, Principal Analyst & Founder, IMEX Research.com

This article appears in BladeLetter by permission of IMEX Research.

Excerpted from “Blade Servers Industry Report 2003”

©2003 IMEX Research. All rights reserved.

The rise of Blade Computing	2
Deficiencies of Current Architectures	2
CISC vs. RISC.....	3
HW vs. SW Provisioning	3
The evolution of Blade Architectures	3
Fabric Architecture for Blades.....	4
Shared I/O architectures	4
Growth Drivers for Blade Servers	5
Towards de facto standardization	5
Market Segmentation	6
Market Segmentation by Tiers	6
Tier-1 Edge Computing	6
Tier 2 Application Servers	6
Tier-3 Back End Data Base Servers	6
Market Requirements by Tiers-1, -2 & -3 Computing	7
Market Requirements in Tiered Computing	7
Sizing of I/O Traffic in Tiered Computing	7
Market Segmentation by Applications	8
Market Segments by Technology	9
References	10
Acknowledgements	10
Key Blade Server Technologies	11
Key technologies in Computing & Communications Blades	11
State of Progression of Key Technologies	11
IP SAN.....	11
Status of iSCSI.....	11
Industry Initiatives	12
Cisco, Network Appliance, qLogic	12
Microsoft Initiative - Chimney	13
TCP Chimney	13
RDMA Chimney	13
Technology Challenges	14
RDMA	14
Market Acceptance of IP SANs	15
Players in Blade Servers by Market Segment	15
Vendors by Technology Segment	15
The Future of Data Centers	17
Emergent Technologies driving Data Centers of the Future	17
Server provisioning software for Blades	17
Investment opportunities in Servers	17
Definitions	18
References	18

The rise of Blade Computing

Deficiencies of Current Architectures

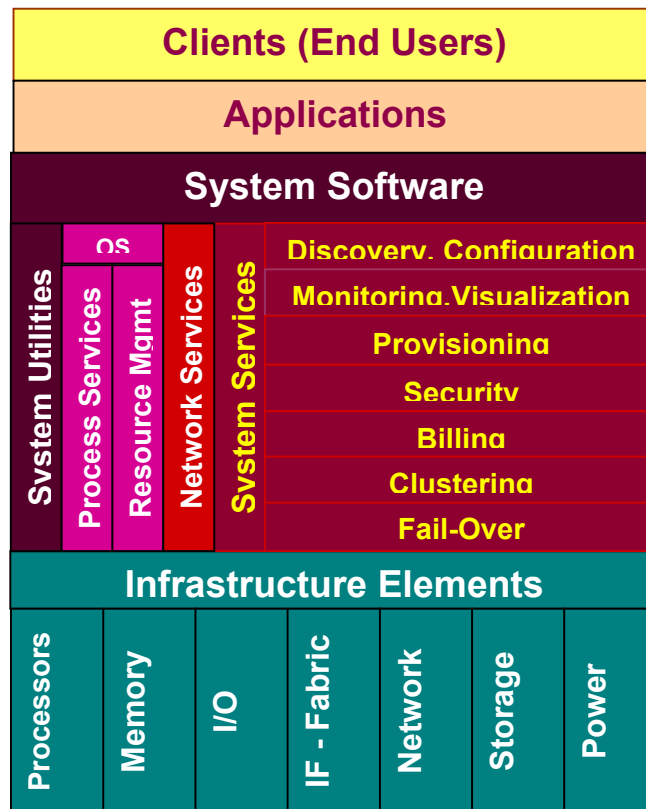
Pressured by CEOs to “do more with less” in this challenging economic environment, interest in server consolidation has soared unabated. And with average CPU utilization in many large companies running around 25% while stranded direct-attached storage assets being utilized at 50%, improving the productivity of its assets has become a top mandate for CIOs in many companies. This low resource utilization has been, in large part, due to the performance and flexibility deficiencies of current server architectures wherein most servers are I/O rather than compute bound.

IT Managers have always drooled at the cost effectiveness of Windows and Linux based servers but the lack of support for multiple workloads concurrently within the operating system have kept them away from deploying them in mission critical environments. In contrast, the Mainframe and Unix based systems that they already own have sophisticated software tools providing hardware level provisioning. This enables them to deploy several independent workloads in each operating system instance, with each workload executing in a separate virtual server environment.

CIOs are willing to invest in new technologies so long as deploying them provides them the benefits inherent in cost-effective industry standard, interoperable IT gear. But with ad hoc proliferation of servers within the organizations and each server requiring a multitude of connections such as several power cords, multiplicity of adapters for primary and backup storage, networking, keyboard, video, and mouse outlets, the process of provisioning server resources has become complex, human resource intensive, error prone and very disruptive.

The ultimate cost reduction opportunity for servers, however, lies in the elimination of monolithic, proprietary server environments altogether and instead augment the industry-standard server operating systems from Microsoft and Linux community with advanced access, availability, and administration functions emulated or replicated from proprietary mainframe and Unix environments and implement them on scalable and modular computers.

The ability of new generation blade servers to be scalable (incrementally added), to manage shifting peak workloads through software provisioning versus the inflexibility of large UNIX based SMP servers dedicated to certain types of workloads puts them at a distinct advantage for adoption in various environments from small businesses to large enterprises as well as system integrators and service providers.



CISC vs. RISC

Mainframe and proprietary, RISC-based Unix environments typically allow provisioning of processors, memory, and I/O resources across a large number of concurrently executing processes. As a result, IT staff can consolidate multiple applications and services onto a single server, thereby dramatically lower operating complexity.

The price of these high-end Unix systems, is typically influenced by several additional value-added silicon, system, and software features that offer dramatically improved server access, availability, and administration features relative to commodity Intel servers running industry-standard operating systems from Microsoft or the Linux community.

However the economic significance of Intel's push into 64-bit computing should not be underestimated. HP and IBM have committed to migrating their RISC-based server environments to IA-64. Sun Microsystems ultimately will follow the other vendors in adopting the IA-64 architecture. Sun Microsystems may be the last holdout but the ongoing price/performance improvements of Intel's designs relative to the SPARC platform (despite their rumored outsourcing of high end Servers to Fujitsu Japan) will leave Sun with no visible alternative but to embrace it.

Intel's new 64-bit architecture with advanced features rivaling RISC platforms promises to relegate the proprietary RISC platforms to the annals of enterprise computing history, over time. The price/performance position of Intel's 64-bit computing should find a receptive ear from the current "do more with less" mentality of IT personnel.

HW vs. SW Provisioning

One of the key contributions of value-added core logic solutions relates to hardware enforced provisioning. Hardware enforced provisioning involves the segmentation of server resources at the processor, multiprocessors, or board level. Each of these segmented hardware resources then may be allocated to one of multiple operating system instances. This form of provisioning provides a high degree of fault tolerance, simplifies the hardware upgrade process, and may allow hardware at different performance levels running diverse applications to co-exist in a single system or chassis. This is the basis of blade architecture.

Intel server designs will end up replicating the provisioning capabilities of mainframes and proprietary Unix platforms to drive greater adoption in large application and database systems. In the future expect Fujitsu, IBM, Hitachi, HP, NEC and Unisys to apply their system design expertise to create hardware level provisioning to create differentiated platforms on industry-standard processors.

The evolution of Blade Architectures

The lack of sophisticated provisioning capabilities has so far been a key inhibitor of adoption for Intel servers in business critical environments. Now emerging "blade server" platforms offer some further relief for the Intel ecosystem.

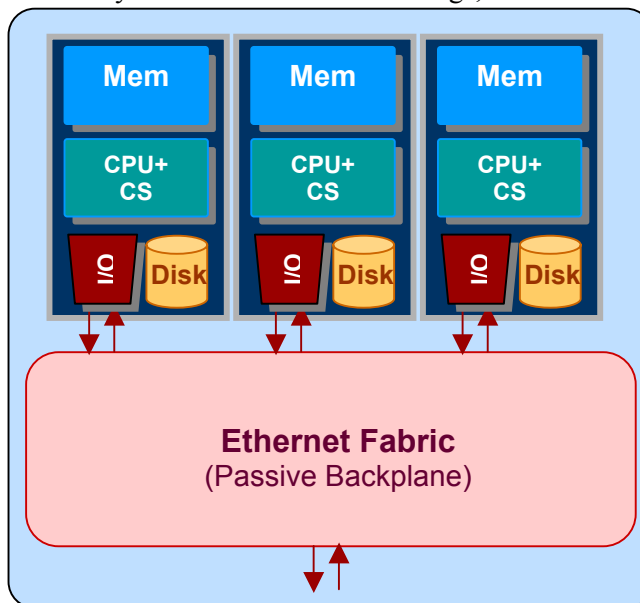
Blades facilitate convergence of computing (channels) and Networks (nodal interconnectivity) for delivery of transactions and content between modular system nodes sharing resources to achieve low cost computing. This is accomplished by use of standard high volume modular servers while mitigating scalability and low reliability bottlenecks of existing monolithic von-Neumann computer architectures of today.

Servers are packaged as one-, two-, or four-way Intel CPU “blades”, and these blades may be added or removed from the system through a single connection to a highly integrated back plane, (called midplane in blades parlance) dramatically simplifying cable management.

Major vendors offering such configurations include Dell, Egenera, Fujitsu-Siemens, Hewlett-Packard, IBM, and RLX Technologies and Sun Microsystems. Earlier blade offerings, such as those from HP, Dell, Fujitsu and Sun, targeted the “stateless” workloads in the front-end web access and presentation tier. On the other hand IBM, Intel, Egenera and second-generation blades from HP directly targeted the heavy-duty enterprise workloads for application and associated database servers in backend tier as well.

Fabric Architecture for Blades

The first generation blade servers rely on captive, bus-oriented connections for storage and network I/O wherein the operating system assumes complete control of the connected storage and network I/O resources at the expense of balancing server, network, and storage resource utilization or their scalability.

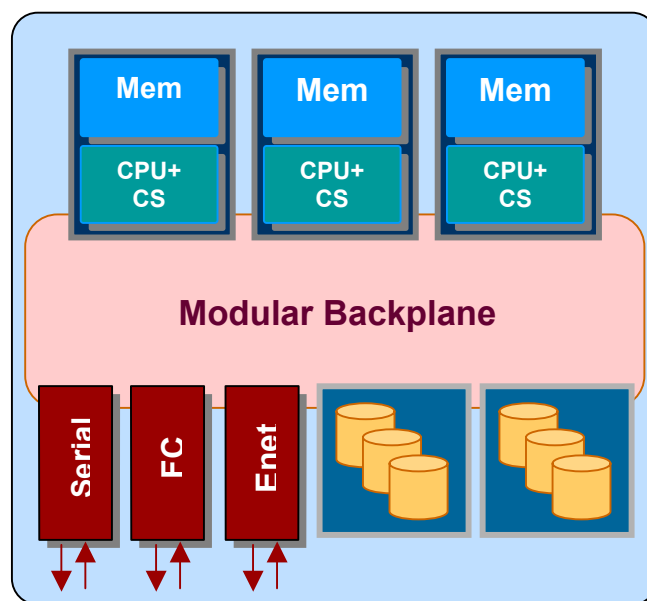


Data Center managers, unable to predict the performance from a given I/O capacity in place end up throwing new servers to scale aggregate performance when the network or storage I/O connection becomes saturated, despite the fact that server processors remain underutilized. Shared I/O architectures (a concept borrowed from the mainframe and Unix server markets) such as those promised by blade architectures eliminate these concerns by decoupling the I/O complex from individual servers. Multiple server blades share storage and network I/O resources and are provisioned with minimal effort, to individual or multiple blades as needed.

Shared I/O architectures

Future blade server options are likely to employ fabric back plane designs that support shared I/O configurations.

Most Intel server platforms, including first generation blade servers, rely on captive, bus-oriented connections for storage and network I/O. In these configurations, the operating system assumes complete control of the connected storage and network I/O resources. This presents challenges when attempting to balance server, network, and storage resource utilization. Shared I/O architectures eliminate these concerns by decoupling the I/O complex from individual servers.



Some of IBM's latest Intel platforms offer a Remote I/O feature that allows multiple Intel servers to share I/O resources located in a separate I/O adapter rack. Egenera uses an internal SAN fabric as the back plane for its BladeFrame architecture for concurrent access to shared I/O resources.

Growth Drivers for Blade Servers

The Data Centers migration is underway to leverage the volume driven economics of industry standard servers using Linux/Windows operating environment. The integrated modular architecture of blade systems providing Virtualization, Provisioning and Self-Driven Automation capabilities already are starting to strike a pleasant resonant chord with CIOs.

Some of the major growth drivers accelerating this phenomenon include:

- UNIX to Linux Migrations to leverage cost effectiveness of open source Linux
- HPC Linux Clusters have struck an explosive growth in academia, national laboratories for scientific computing and migrating to commercial world for Bioinformatics, decision support financials, visualization etc.
- Server Consolidation onto fewer servers with virtualization of resources, provisioning and lights out automation capabilities
- Blades Density to provide economic benefits in real estate and operational environmental.
- Standardization on fewer OS types for servers, storage and network protocols deployed.

Towards de facto standardization

Improved hardware provisioning provides a compelling value proposition for blade servers, but lack of standardization is undermining this message. Each server vendor initially has attempted to create proprietary designs to protect margins and secure long-term footprint in data center configurations. This lack of standardization is impeding adoption, as customers want standard form factors supported by multiple vendors.

IBM has made Blade servers as a cornerstone of its x-series (Intel) product line and is heavily pushing its BladeCenter along with larger SMP computers. To mitigate concerns of non-standardization and to leapfrog other competitors IBM chose to parlay its relationship with Intel to create a partnership in which the two companies share technology and development costs for blade server designs and drive a de facto standard. IBM is hoping that the partnership will drive the commoditization of blade server hardware, allowing the company to win the day on its superior software and services. In exchange, Intel got a faster ramp in blade server designs and the blade market and, in turn, allows Intel to drive volume of its Xeon (and later Itanium) processors into higher performance, higher margin enterprise markets.

IBM has taken considerable share in the blade server market in recent months. It has now become the new leader in blade server revenues in 2Q2003, according to IMEX Research, a San Jose CA Technologic Research & Consultancy. (See Blade Servers Industry Report 2003 – IMEX Research Doc.BS-203 Sep2003)

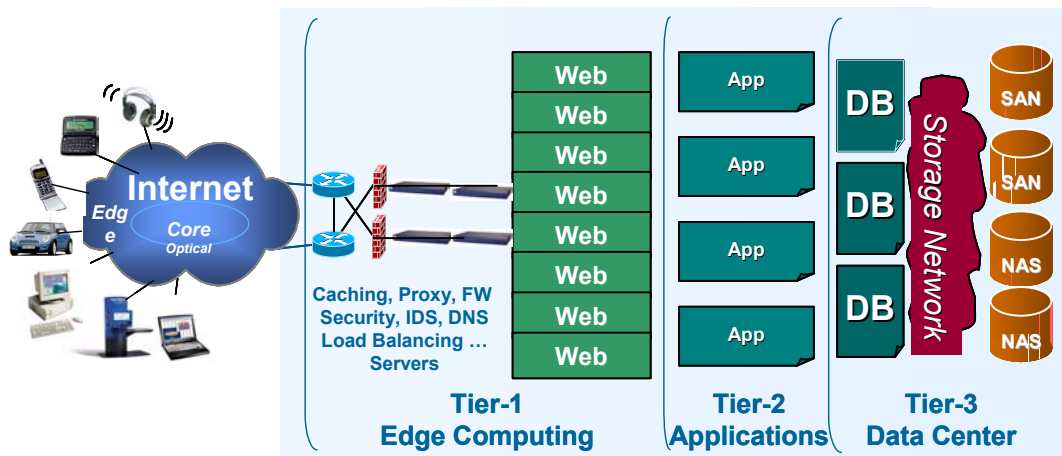
Intel is now selling its BladeCenter “clones” to other OEMs. It has signed 12 OEMs and System Integrators worldwide including Bull Computer etc. Indications are that Dell and HP are giving strong consideration to a license agreement with Intel for the BladeCenter design in order to accelerate their blade server roadmaps.

Market Segmentation

Market Segmentation by Tiers

Blade Servers are architected to be flexible for software provisioning to match the requirements of applications. AS such it is easy to address the various specialized requirements of Tier1, 2 and 3 applications in Blades.

3 Tier Computing Architecture



©2003 IMEX Research. All rights reserved

Tier-1 Edge Computing

Segment Front End & Web Servers
Markets ISPs, Low End Bus
Addressed Applications Web Servers, Infrastructure ...
Work Loads Caching, Directory Services, DNS, Proxy, Firewall, Security, Intrusion Detection Servers, Low End File/Print, Low End Messaging ...

Tier 2 Application Servers

Segment Back End - Mid Range Applications
Markets Small/Medium Businesses
Addressed Applications Application Servers, Infrastructure, High Performance Computing, Streaming ...
Work Loads Midrange OLTP, ERP, SCM, CRM, File/Print, Messaging, HPC Linux Clusters ...

Tier-3 Back End Data Base Servers

Segment Back End – Data Base Applications
Markets Enterprise Businesses
Addressed Applications OLTP/DB Servers, Numeric Intensive Computing, Decision Support ...
Work Loads High End Data Sets, OLTP, ERP, CRM, SCM, High End SMP/HPC Clusters

Market Requirements by Tiers-1, -2 & -3 Computing

Market Requirements in Tiered Computing

	Tier-1	Tier-2	Tier-3
Applications	Front End – Web Apps	Back End - Mid Range Apps	Back End – Data Base Apps
Markets	ISPs, Low End Bus	Small/Medium Bus	Enterprise Businesses
Typical Applications	Web Servers, Infrastructure ...	Application Servers Infrastructure High Perf. Computing Streaming ...	OLTP/DB Servers Numeric Intensive Computing, Decision Support ...
Work Loads	-Caching, -Directory Services, -DNS, - Proxy, - Firewall, -Security, - LE File/Print, - LE Messaging ...	MidRange – OLTP, – ERP, - SCM, -CRM, - File/Print, - Messaging - HPC Linux Clusters ...	High End Data Sets, - ERP, CRM, SCM HE SMP/HPC Clusters
Performance	Connections Intensive	Transactions	Data Access & Movement
Chassis*	1-2U	2-4 U	4-8U
CPUs/Blade*	1 - 2 way	2 - 4 way	4 - 8+ way
CPUs*	LV Pentium/ Embedded	Pentium/Athlon/ U-Sparc2/IBM	Xeon/Opteron/ U-Sparc3/IBM
Storage*			HE RAID & Backup
I/O Requirements* - Network Traffic - IPC - Storage	Very High** Very Low High	Low High Very Low	Very Low High High
Manageability	OS based	Add-On Tools	Proprietary Tools
Typical Examples	Fujitsu-Siemens BX 300 Dell 1600 MC	HP BL 40p Egenera	IBM Blade Center HS 20 Intel HP BL40e Egenera

*Typical
Registered names and trademarks belong to respective companies
**Legend

Sizing of I/O Traffic in Tiered Computing

I/O Traffic	Very Low	Low	Medium	High	Very High
Gbps	< 0.1 Gbps	0.1-0.5 Gbps	0.5-1Gbps	1-5 Gbps	> 5 Gbps

Market Segmentation by Applications

Web Infrastructure – Caching, Firewall, DNS, Security (Secure Sockets, Data Compression, Encryption..), VPN, Load Balancing, Web Content Acceleration, Directory Services, Proxy Servers, ...

Data Center Infrastructure – File/Print, Data Backup & Restore, Resources Allocation & Administration, Data & Network Management, System Utilities ...

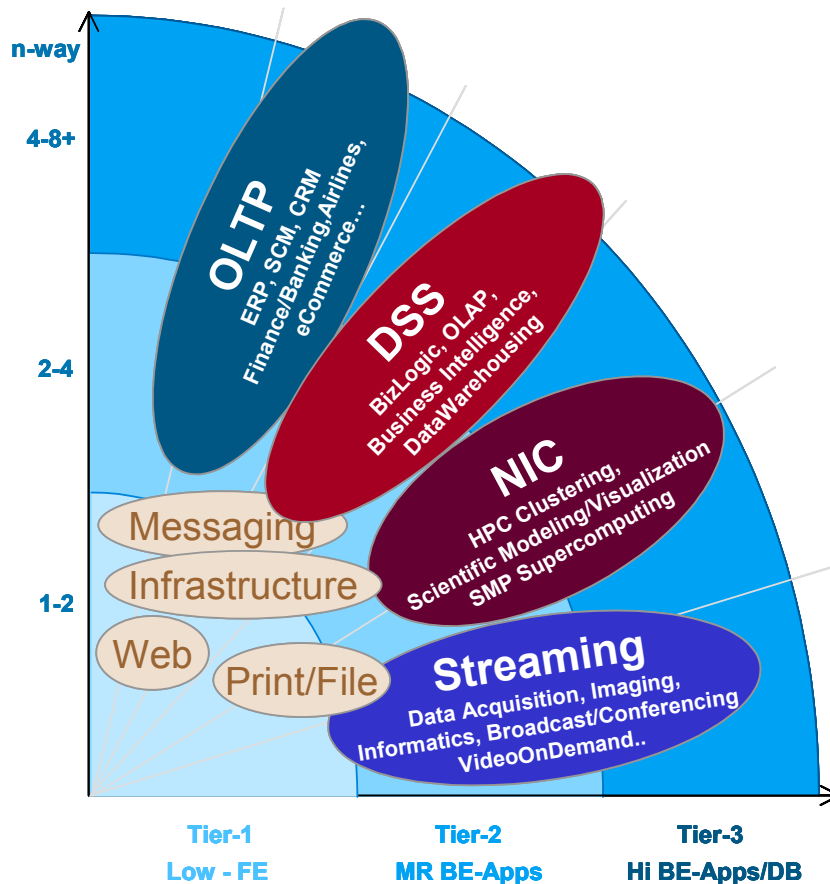
High Performance Computing – Clustering, Parallel Processing, Numeric Intensive Computing, Supercomputing ...

Online Transaction Processing – OLTP, SCP, CRM, ERP ...

Decision Support – Application Development, Business Logic, OLAP, Business Intelligence, Data Warehousing ...

Streaming – Data Acquisition, Medical/Bioinformatics, Video On Demand ...

Applications in 3-Tier Computing



©2003 IMEX Research. All rights reserved

IMEX

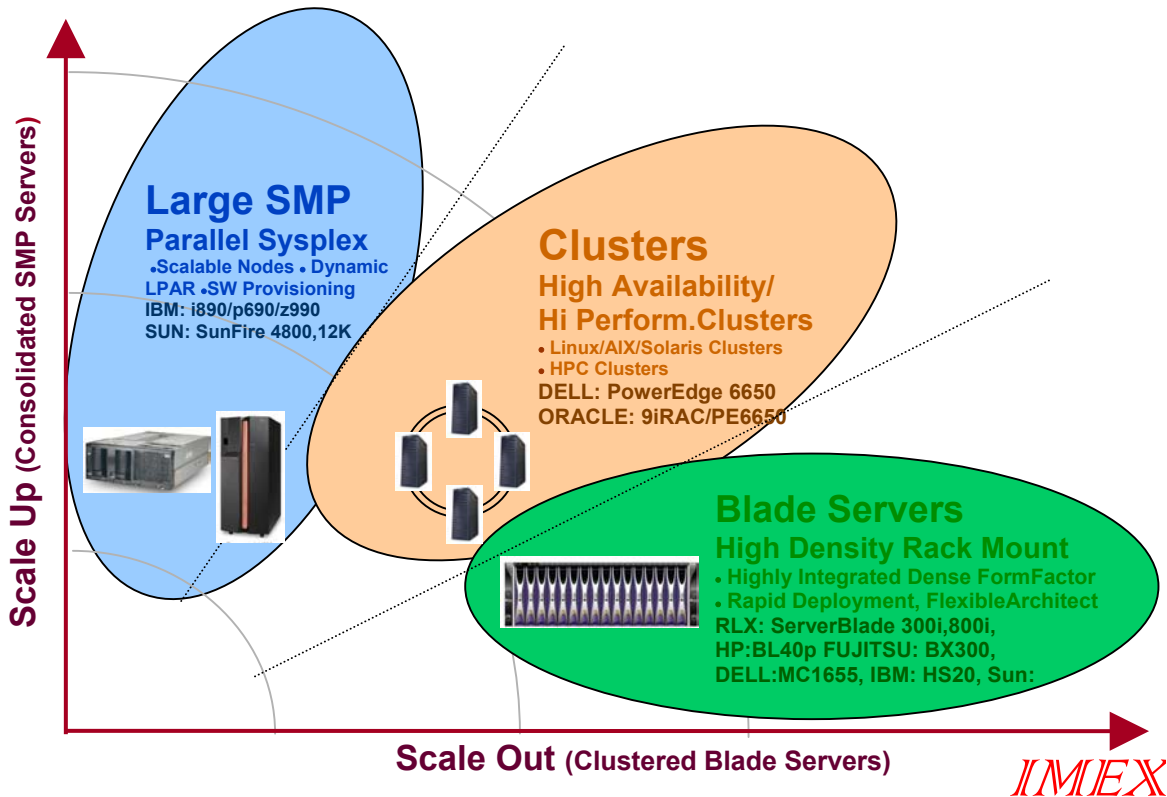
Market Segments by Technology

Fundamentally volume driven economics of industry standard servers is driving the scale out architectures for servers. Software Provisioning technologies used in single server are giving way to hardware provisioning of cluster of low cost nodes. Technologies such as Clustering, Back Plane Fabrics (e.g. using Infiniband or TCP Offloaded IP to achieve low latency, high bandwidth), Virtualization/Provisioning Software, High Density/ Power Management, High Availability architectures, etc.

The platform market can be segmented in three distinct areas based on Scale Up and Scale Out technologies deployed:

- Large SMP Servers using Software Provisioning for Parallel Processing
- Cluster Architectures for High Availability or High Performance
- High Density, Highly Integrated Blades with non-blocking Interconnect Fabrics

Server Market Segmentation - Platforms



References

- Report: **Blade Servers Industry Report 2003** *Document RPT-107 (Aug 2003)*
Brief: **UNIX versus Windows based Server Architectures.** *Document BRF-112 (Feb 2003)*
Report: **Competitive Analysis of Startups and Incumbents in Blade Computing & Communications** *Document RPT-115 (Oct 2003)*
Brief: **Market Opportunities in Blade Computing – Scalable architecture from PCs to Supercomputers** *Document BRF-128 (Dec 2003)*

Acknowledgements

Royal Bank of Canada (RBC) Papers (Oct 2003) for incisive views on Fabric Computing
Steven Beedle, ZNA Communications for continuous helpful support
Dr. Lance Leventhal, Emulative Systems, for being the mentor