

SIMPLIFYING IT

Marcel van Drunen

**Enterprise Technologist
High Performance Computing**



HIGH PERFORMANCE CLUSTER COMPUTING

REDEFINING HPC PRODUCTIVITY

“We’re focused on scalable and flexible solutions that simplify high-performance computing by reducing cost and complexity.


What we’re learning about HPC technology will redefine productivity throughout the research, discovery and business computing ecosystem.”

Michael Dell - 2008



MARKET SEGMENTATION BY IMPLEMENTATION

**High Performance Computing Servers
\$10 Billion**



**Biggest market
Fastest growth**



MARKET SEGMENTATION BY SECTOR

**High Performance Computing Servers
\$10 Billion**

```
graph TD; A[High Performance Computing Servers $10 Billion] --> B[Government]; A --> C[Industry/ Private Sector]; A --> D[Academic/ University];
```

Government

**Industry/
Private Sector**

**Academic/
University**

New HPC Markets

- Mid-size Engineering companies
 - Used to work with heavy workstations
 - CFD software got 'heavier'
 - Models increase
 - Time to solution doesn't decrease
 - > Conversion to Cluster computing
- Genetics department of smaller university
 - Used to do only theoretical work
 - The tools got cheaper and more accepted
 - >Now have a ton of data, need a cluster to work with it, but don't have 'human resources'



ISSUES FACING RESEARCH DEPARTMENTS

COST CONTROL

ACCESS TO
MORE COMPUTE

REDUCED TIME
MANAGING

BOTTOMLINE:
HPC CLUSTERS ARE TOO COMPLEX

ENSURING
SCALABLE
SOLUTION

REDUCED
FUNDING

OUR MISSION

Dell™ is committed to
Simplifying HPC
by removing cost and
complexity



IDC QUOTE

Much remains to be learned about the low end, especially the white-hot sub-\$50,000 Workgroup technical server segment. **But this much is already certain these users typically require ease-of-everything; purchasing, installation, operation (including software compatibility and performance), and upgrading.**


Source: <http://www.hpcwire.com/topic/systems/17905264.html>



HOW WE DO IT

FINANCE

Through DFS leasing rather than purchasing you can achieve cost-efficient acquisition, management, replacement and disposal of technology of your HPC technology.

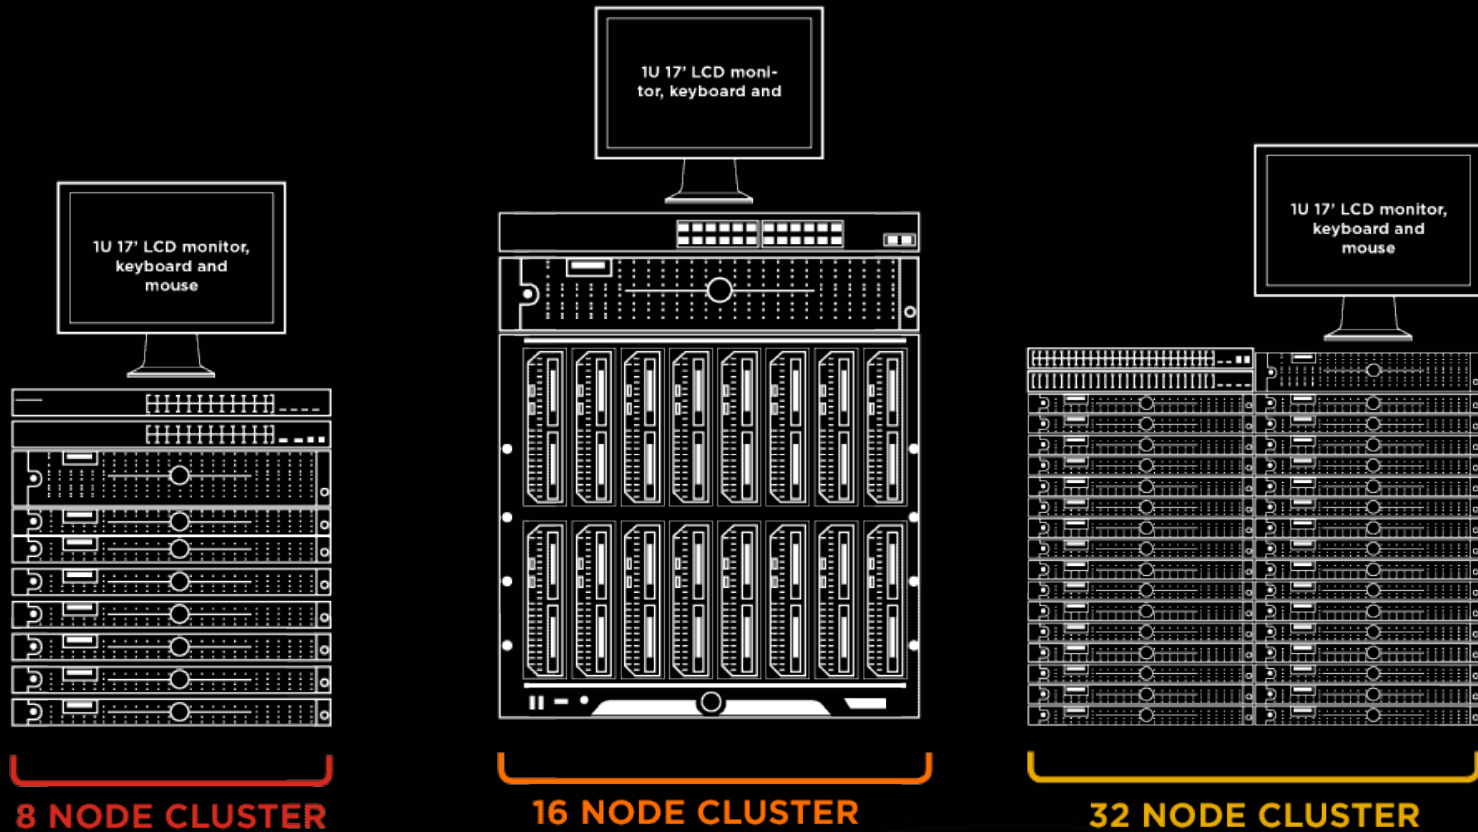


Standard offerings

- We recommend a 'supportable' OS
 - Could be Red Hat, Windows, SuSe
- Dell engineers can build a standard cluster
 - Platform Cluster Manager
- For Application integration and other advanced installations we recommend valued partners like:
 - ParTec
 - ClusterVision
- 'Non standard' storage:
 - DDN
 - Panasas



RANGE OF VALIDATED CONFIGURATIONS

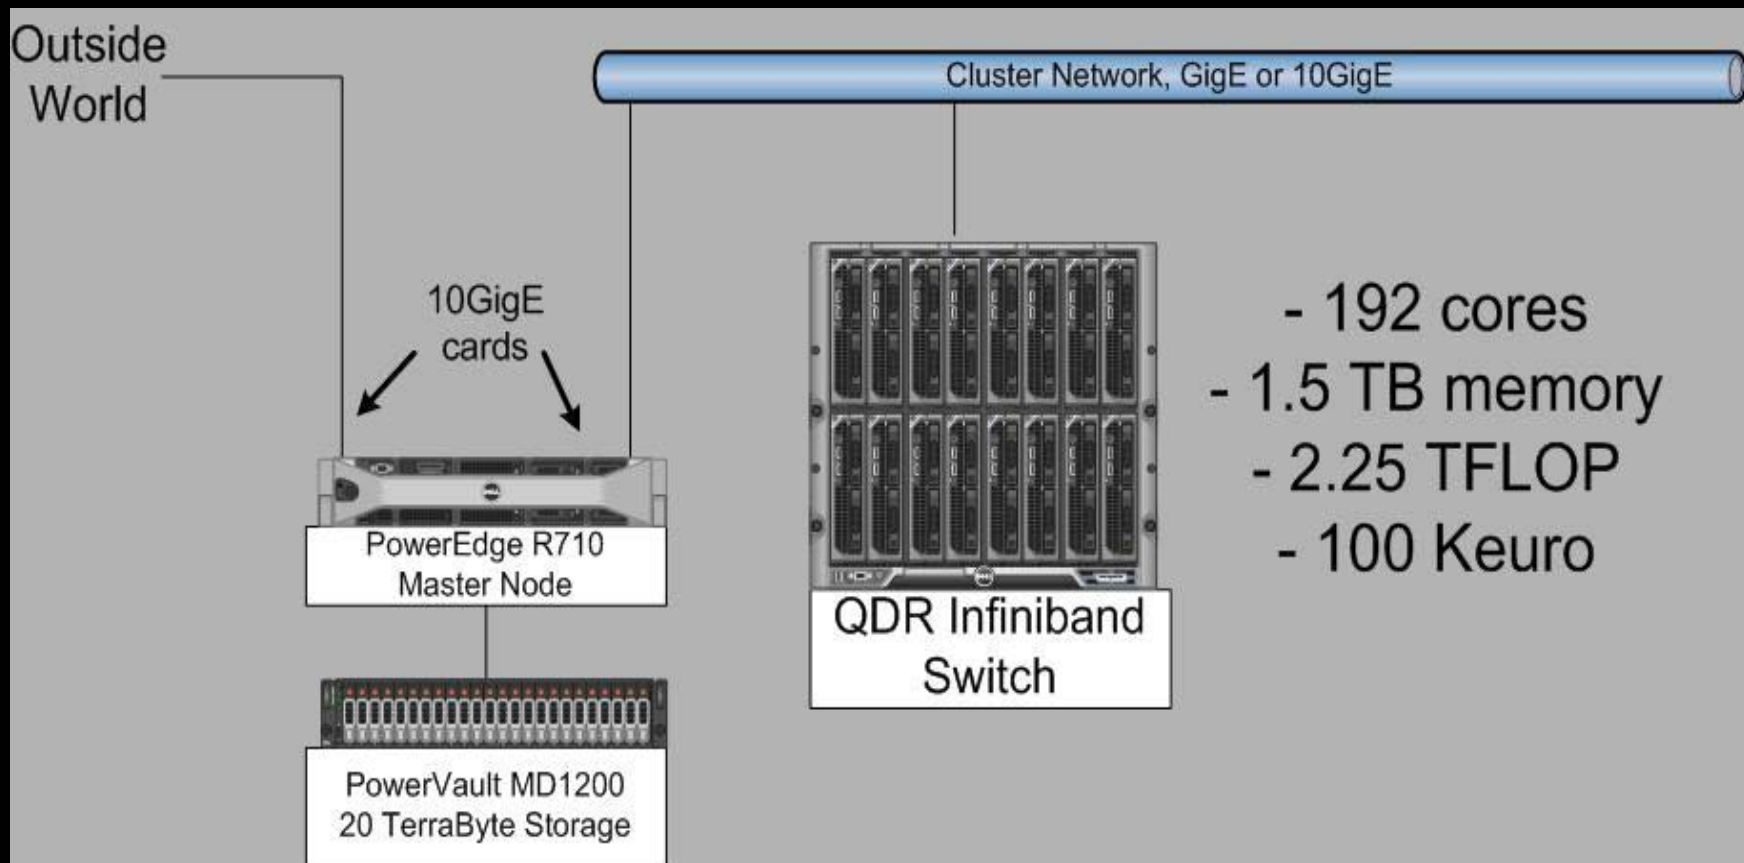


Most popular server format

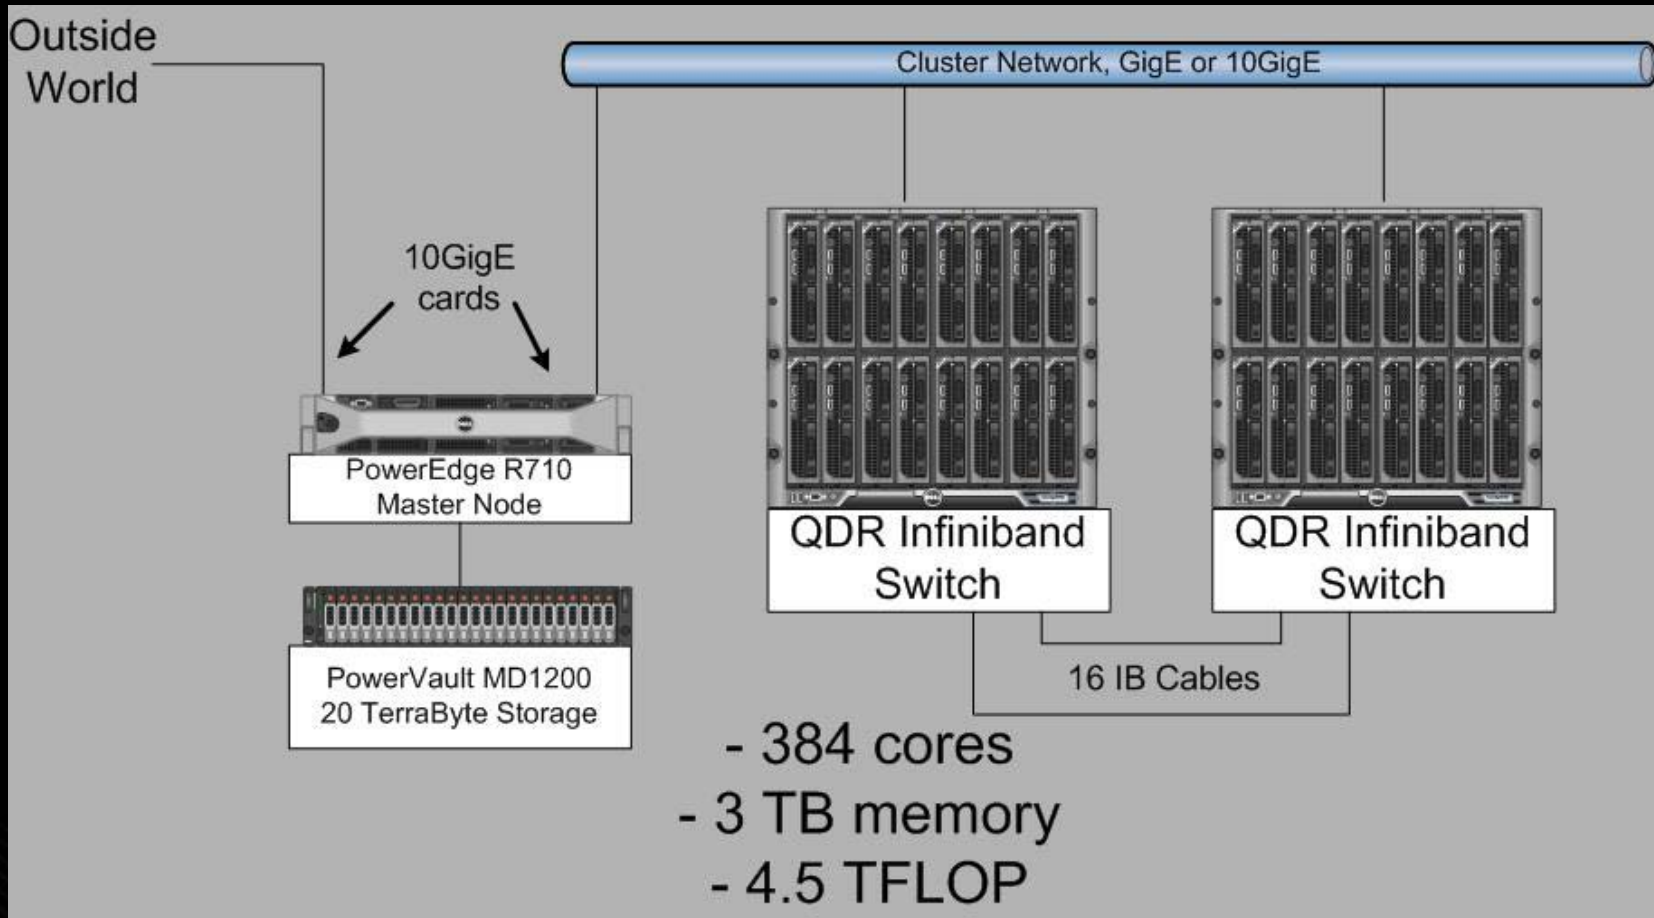
- 'Good enough' density
- Superior management features
- Almost no cables, specially with sub-16 node clusters



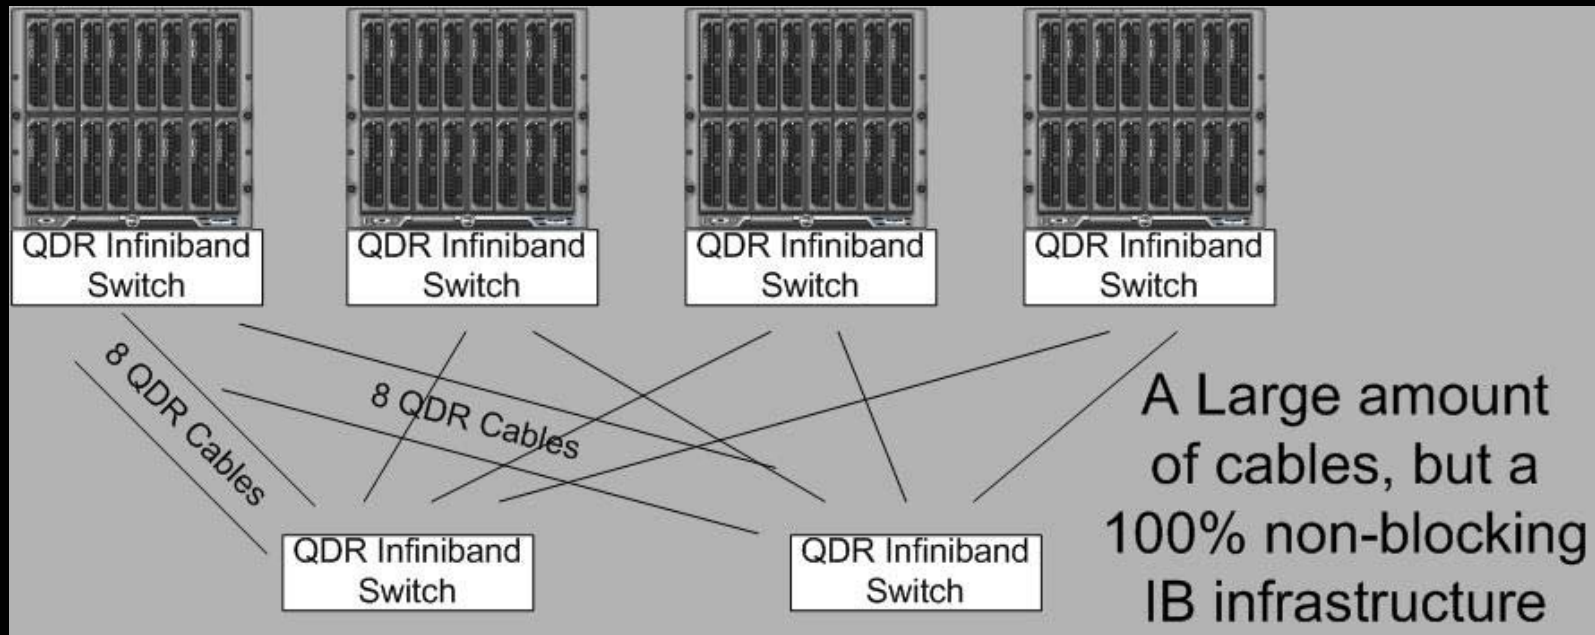
Compute power in a box



Compute power in two boxes



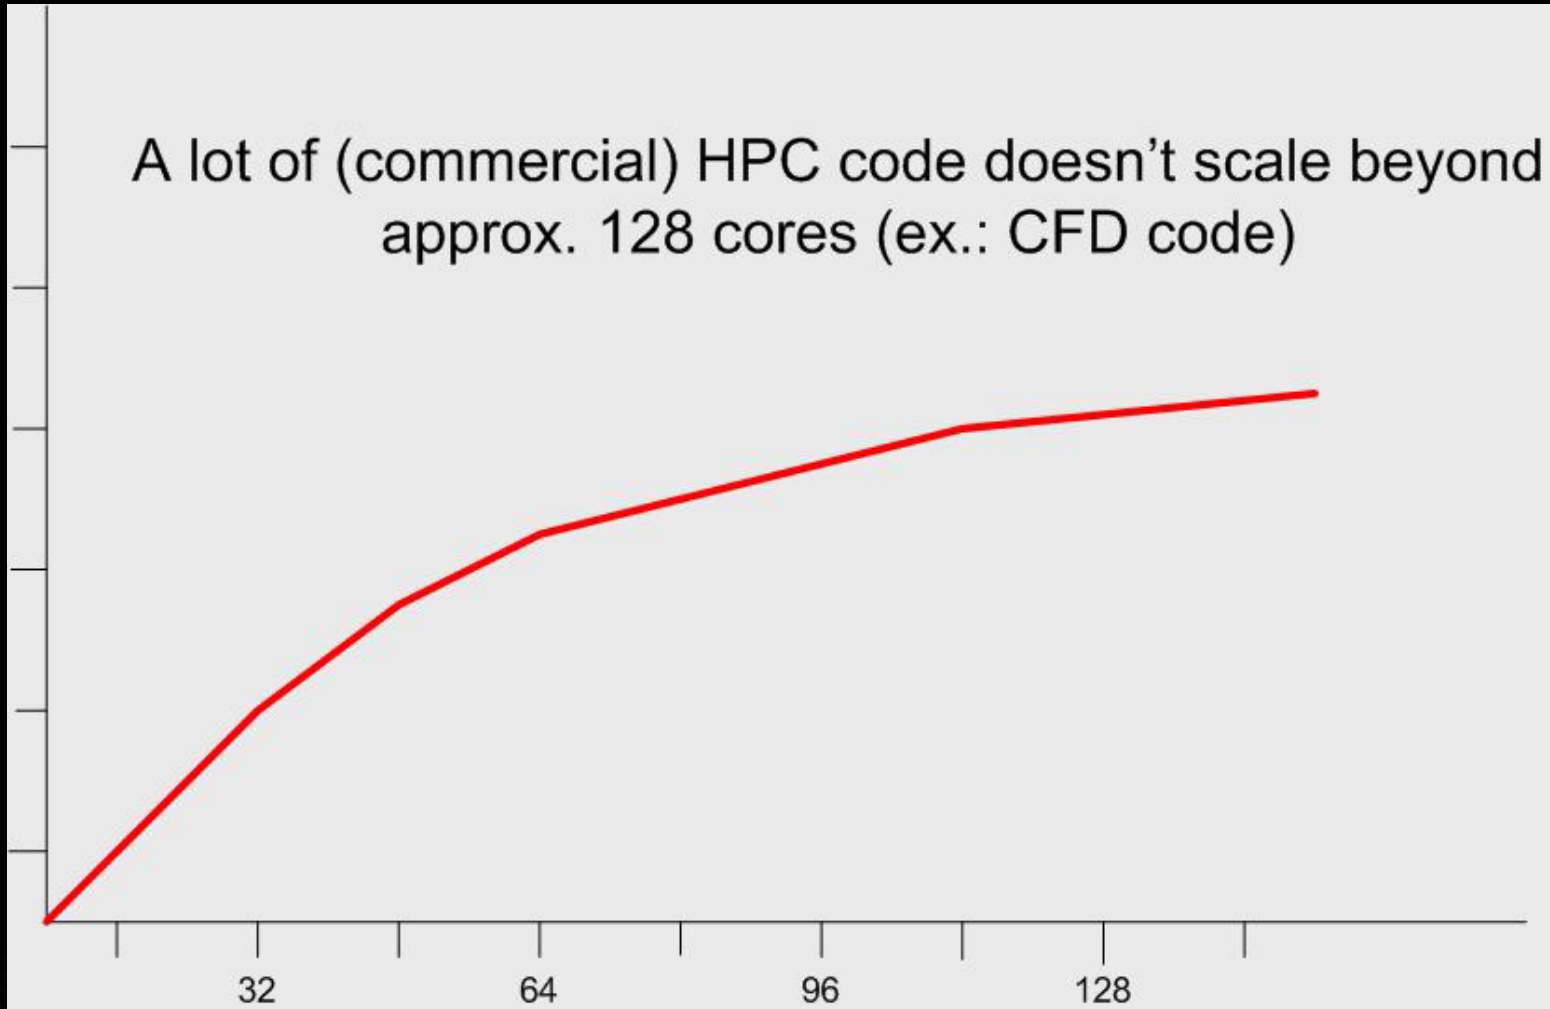
Cable sprawl...



Note to presenter: You're close to Italy, so you cannot say you don't like spaghetti

Do you need a non-blocking Infiniband infrastructure?

A lot of (commercial) HPC code doesn't scale beyond approx. 128 cores (ex.: CFD code)



Density painpoints

Large customer in Middle-East

- Has 500+ node clusters
- Would like more density
- Can only bring 7 KW to a rack and cool it
- > Ends up with half-filled racks

Large academic customer in BeNeLux

- Has 300+ node cluster
- Likes more density
- Has a weight limit of 680 Kg/m²
- > ends up with half-filled racks



Standards

- QDR Infiniband is the standard cluster Interconnect that Dell supports
- 10GigE is becoming standard for cluster communication, for job submission and storage access
- 10GigE latency is coming down, but still more expensive, and less bandwidth, than QDR IB
- Converged networking almost everywhere, but not much in HPC, depending on definition



HPC Challenge: feeding the beast

Dell PowerEdge R815 Server:

- 4 x 12-core AMD CPU's -> 48 real cores per system
- Up to 512GB of memory per system
- Network and storage are now bottleneck
- How many Infiniband HCA's per system?
- Future brings even more fun challenges with growing GPU integration



Developments

- Price of memory is coming down
 - 3 or 4 GB per core is affordable
 - Price of SSD drives is coming down
 - No need for tricks to hide disk latency
 - SSD's make a good match with IB latency and bandwidth (NFS over IB?)
 - Will need one more generation for adequate storage size
 - Price of GPU is already low, but little support in commercial software
- The software people are at bat




Partner Solution: vSMP

ENTERPRISE APPLICATIONS

Applications requiring **fraction** of the physical server resources

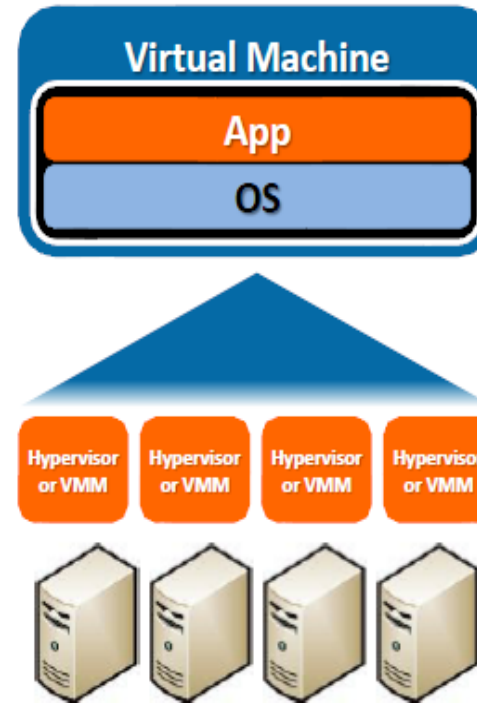
Existing: Partitioning




HIGH PERFORMANCE COMPUTING

Applications requiring **superset** of the physical server resources

New: Aggregation



Quickest way to build a large SMP system

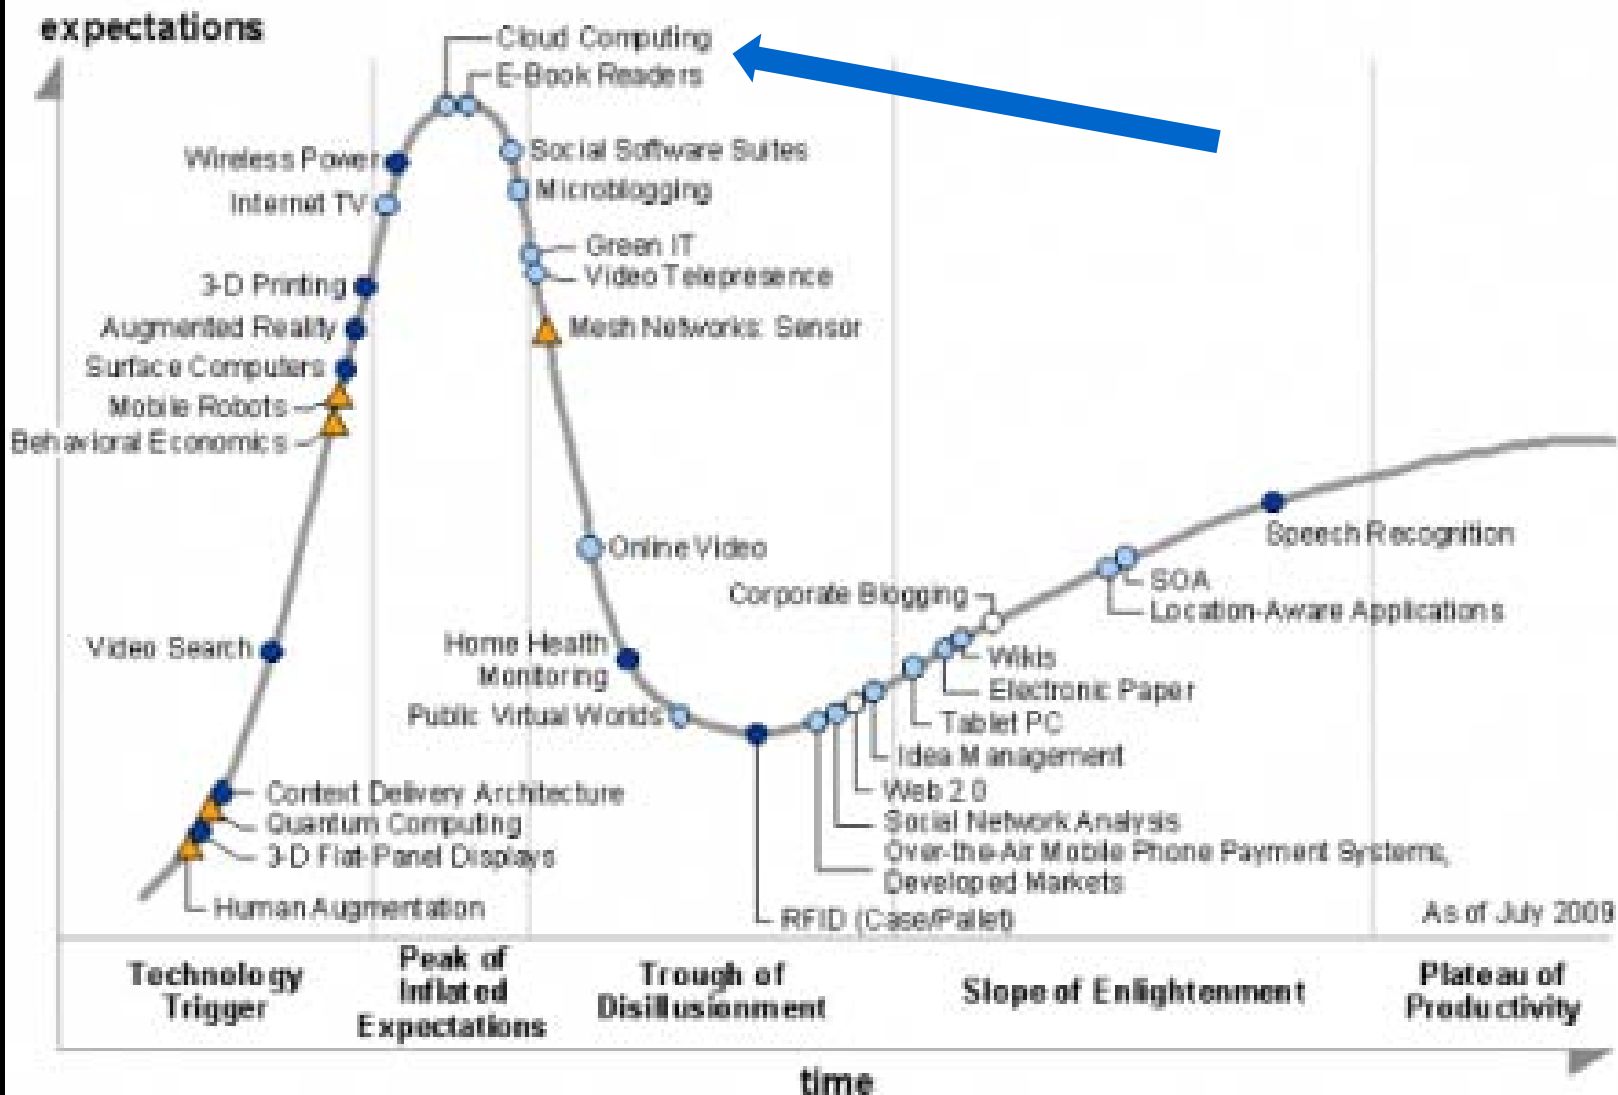


ScaleMP features

- Virtual SMP machine:
 - Up to 32 CPU's, 192 cores
 - Up to 4 TB of memory
 - These figures will soon increase (Nehalem EX etc.)
- Dell PowerEdge blade chassis plus a hypervisor
 - Support at fraction of cost for 'real' SMP machine
 - Hundreds of engineers live around the corner
 - The spare parts are everywhere
 - No proprietary hardware
 - works with Mellanox Infiniband
- SMP performance for the price of commodity hardware



Hype Cycle of Emerging Technologies, 2009



Source: Gartner (August 2009)



Cloud

- Dell is one of the major Cloud vendors in the world
 - Large search engines and online shops
 - Large public cloud providers
 - Will soon launch private cloud offerings
 - Launching Cloud-class hardware THIS WEEK
- HPC customers may turn to the cloud
 - Final step in simplification
 - Pay-per-use HPC
 - Common cloud objections also apply to HPC



SUMMARY: A SMARTER PATH TO HPC CLUSTERS

**GET
FASTER**

- Flexible, industry-standard HPC Clusters deployed in days not weeks
- Fast start services streamline transition to productive research

**RUN
BETTER**

- Validated end to end solutions for all organizations
- Leverage Dell's knowledge building 24 of the world's fastest supercomputers*
- Balance energy efficiency and high levels of performance

**GROW
SMARTER**

- Broad choice in inter-connect, form-factor and management technologies
- Scalable HPC storage capacity & performance

* The TOP500 project ranks the sites that operate the 500 most powerful computer systems based on the best performance on the Linpack benchmark. The Top500 List is published twice a year and as of November 2007 Dell ranked 24 of the top 500 computing systems representing a total of 77,828 processors. Source: www.top500.org



DELL SIMPLIFIES IT

THANK YOU

THANK YOU

© 2008 Dell Corporation Limited. Dell, the Dell logo and PowerEdge are registered trademarks or trademarks of Dell Inc. Dell disclaims proprietary interest in the trademarks or tradenames of other entities used to refer to them or their products. Intel and the Intel Logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Dell Corporation Limited, Reg. No. 02081369, Dell House, The Boulevard, Cain Road, Bracknell, Berkshire RG12 1LF.

