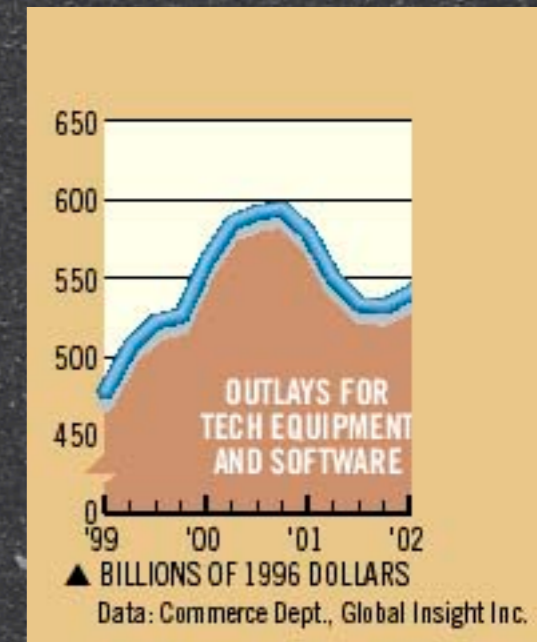


HPC Computing Applied to Business Applications

Moshe Bar
Qlusters, Inc
openMosix

Why are clusters happening now?

- Budget cuts all over
 - 71% of all Fortune500 see savings (TCO) as biggest IT spending driver
- Proprietary SMP systems and supercomputers ridiculously expensive
 - Commodity computer is here to stay
 - (Sun (ie McNealy) is doomed)



Clustering has been studied for the last 25 years

however, several pre-requisites were missing

Clustering Pre-requisites

- Cost per node
- single CPU processing power
- Single node form factor (size)
- Single node heat emission
- Inter-connect speeds

The Moshe Bar Coefficient

$$\text{Moshe's Coefficient} = \frac{(\text{heat emission} * \text{US\$} * \text{cubic inch})}{\text{Gflops}}$$

Heat Emission


System	Performance (sysmark)	Heat Dissipation
Dell 2650 (2 CPUs)	532	32W
Sun FireBlade v100 (1 cpu)	189	36W

Cost per node

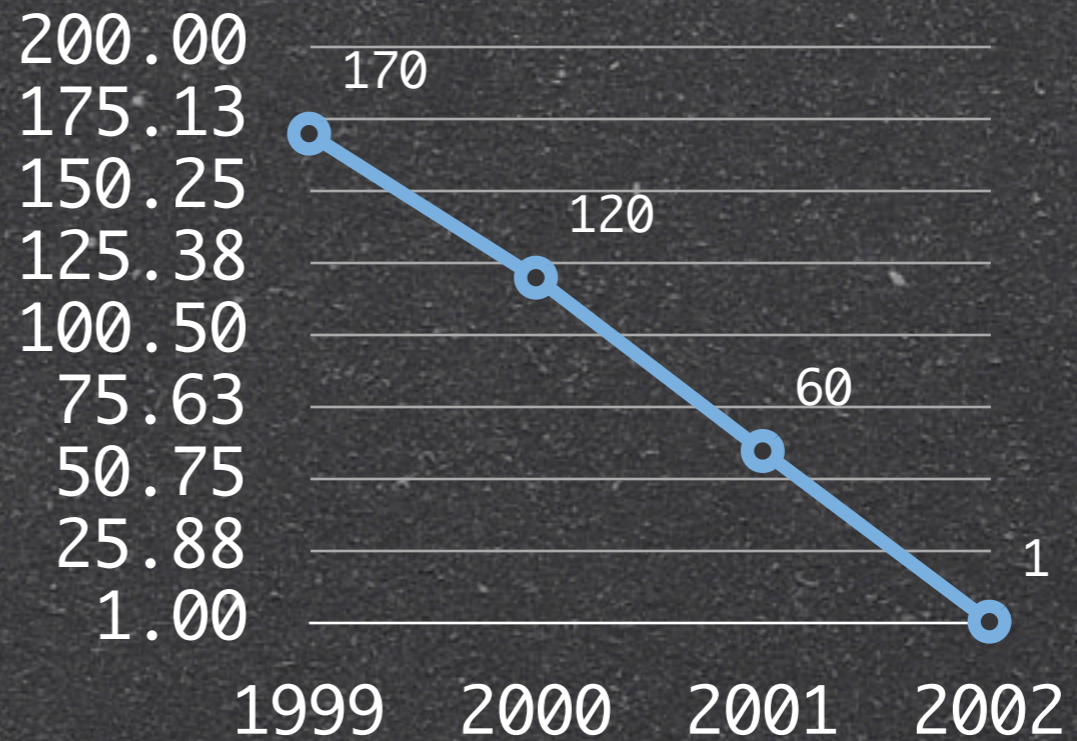
System	Config	US\$
Dell 1750	2 CPU 2.4 Ghz 2GB RAM	2527 \$
Fireblade v210	2 CPUs 1 Ghz 2 GB RAM	7599 \$

Inter-connect Speeds

Bandwidth



Latency



2000

- Only in 2000 did all the components fall in place
- HPC is early adaptor, tech savvy
- Thousands of HPC clusters installed in short time
- Most of top 500 supercomputers by now are clusters, displacing Crays etc.


HPC Clustering

- Two main paradigms:
 - Parallel Computer Emulation (Beowulf)
 - Single System Image (SSI)

Beowulf

- Emulates parallel computer in software
- Must use external libraries
- Must modify application program
- Relies on central controller
- Gene Amdahl's Law has big impact

Beowulf



Gene Amdahl's Law

Ideal World

$$\frac{t}{\text{number of CPUs}}$$

Real World

$$\frac{P}{(S*P)+(1-S)}$$

5% Serial

Nodes	Overall Speedup
1	1.00
10	6.90
50	14.49
100	16.81
1000	19.63
10000	19.96

SSI Clusters

- Same scalability and overhead for 2 nodes or for 2000 nodes
- Users do not see the individual nodes
- Programs do NOT need to be changed to take advantage (unlike pvm, mpi etc.)
- Heterogeneous node configurations are automatically dealt with always automatically load-balanced
- Ease of management

No controlling master in SSI

Single Image Computer



SSI Clusters Today

- 📌 Qlusters ClusterFrame
(www.Qlusters.com)
- 📌 openMosix
- 📌 openSSI

From HPC to Business Applications

- Business applications are a significant investment for companies
- Modifications not possible (shrink-wrapped SW), difficult or too expensive
- No many number crunching problems
- Many scalability problems, however

Other Urgent Problems for Companies

- Management of hundreds of servers (or thousands) is difficult and expensive
- Lower TCO
- Increase availability of applications

SSI Solution

- SSI scales up any kind of applications, unmodified
- SSI enables instant (sub-second) failover of stateful applications
- Provisioning system is easy to add to SSI

Solution Requirements

- Any solution addressing enterprise's needs of today must:
 - Provide scalability
 - Provide extreme availability
 - Lower TCO by automating mgmt of grid of servers

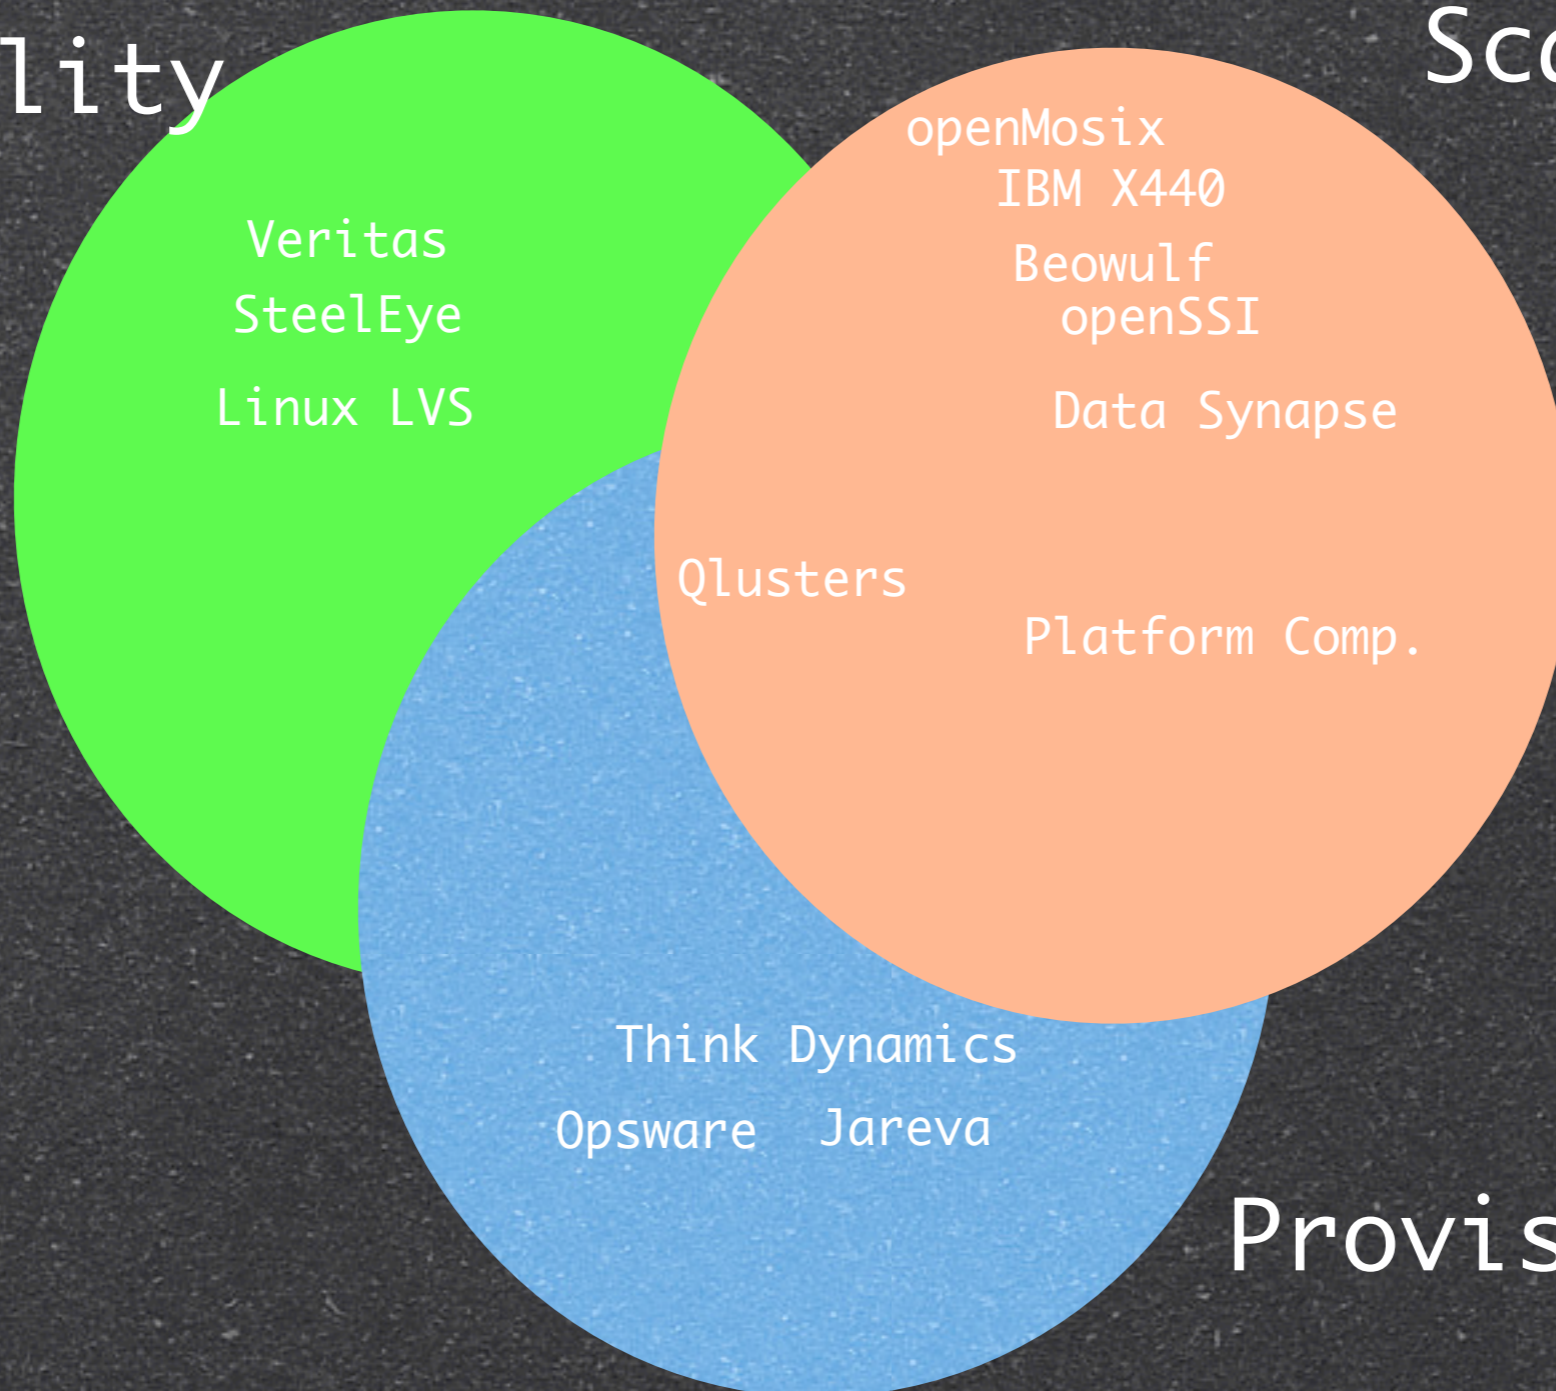
The next big computer
company will be a
software company

Kevin Harvey, Partner at Benchmark Capital

Solutions Today

High
Availability

Scalability



Veritas
SteelEye
Linux LVS

openMosix
IBM X440
Beowulf
openSSI
Data Synapse

Qlusters

Platform Comp.

Think Dynamics
Opsware Jareva

Provisioning/Mgmt