



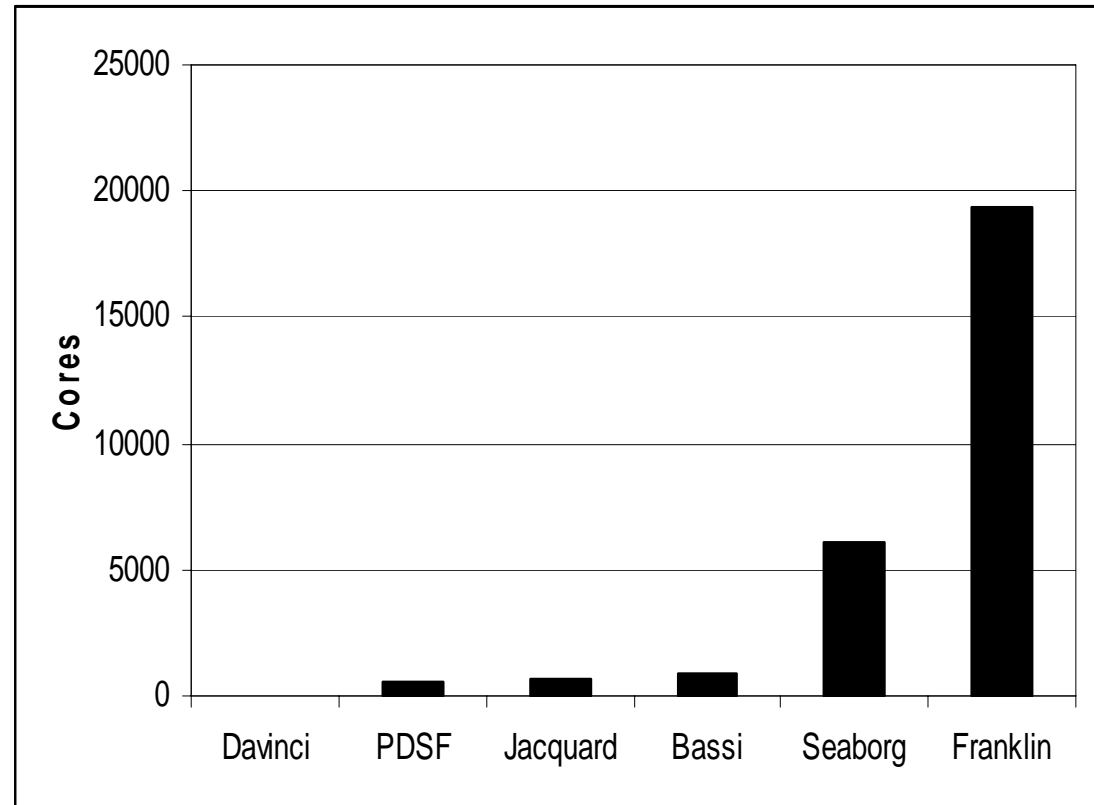
National Energy Research Scientific Computing Center (NERSC)

Scaling Into Tomorrow

Nicholas P. Cardo
NERSC Center Division, LBNL
May 8, 2007

NERSC System's Growth

- Davinci (32)
- PDSF (550)
- Jacquard (712)
- Bassi (888)
- Seaborg (6080)
- Franklin (19376)



Silence

- April 2, 1722, Benjamin Franklin introduces “Silence Dogood”
- 2 Chassis single cabinet XT4
- 10Gb, 1Gb networking
- SeaStar2
- DDN 9550 and DDN 3000
- FC-4 to CISCO 9513 switch
- Physically separated from production system

Franklin




Benjamin Franklin
America's First Scientist

- Cray XT4 (*102 cabinets*)
- SeaStar 2
- XT 1.5.43+
- DDN S2A9550s (*11 cabinets*)
 - 5 couplets, 160 tiers, 300GB FC (*scratch*)
 - 1 couplet, 16 tiers, 300GB FC (*home*)
- 2 x Cisco 9513 FC Switches
- 17 x 24 x 24 Class 3 Torus

Fun Facts

- 9740 Nodes
- 38960 DIMMs
- 39168 GBs Memory
- 1628 Disk Drives
- 3366 SeaStar Cables
- 9133.8 Meters of SS Cables
29968 Feet (5.7 Miles)
- 96 Wires/Cable
- 2876927.8 feet of wire
544.9 Miles

Facility Preparation



Floor Space



Power


15 to 22.5KW per Cabinet

* 102 Cabinets

= 1530 to 2295 KW



=



Partitions

- Service Nodes (52)
 - 8 GBs memory
 - 2.6 GHz dual-core Opteron
- Compute Nodes (9688)
 - 4 GBs memory
 - 2.6 GHz dual-core Opteron

Node Allocations

- Login (16): dual-port GigE
- Network (4): 10 GigE
- I/O(28): Lustre
- System (4): boot, SDB + backups
- 2 SMWs

Lustre

- Home:
 - 5 OSSs/10 OSTs
 - 1 MDS/1 MDT
 - ~2 TB LUN
 - 1048576 Stripewidth
 - 2 Stripecount
 - 4096 blocksize
- Scratch
 - 20 OSSs/158 OSTs
 - 1 MDS/1 MDT
 - 1 MDS Backup
 - ~2TB LUN
 - 1048576 Stripewidth
 - 4 Stripecount
 - 4096 blocksize

Batch System


→ interactive

→ debug

→ premium

→ regular

TORQUE/Moab



→ low

Networking

