


OS Issues for Ultrascale Architectures


Arthur B. (Barney) Maccabe
Computer Science & HPC@UNM
University of New Mexico




Cougar versus Linux on ASCI/Red



NAS B: mg



NAS B: cg



Apples and Pears



- Unfair to Linux
 - ◆ Portals 3 versus Portals 2
 - ◆ Random node allocation
- Unfair to Cougar
 - ◆ Old compilers and libraries (elf versus coff)
- Is it really unfair?
- Linux issues
 - ◆ non contiguous allocation (TLB misses)
 - ◆ file centric
- LANL (Tru64) and LLNL (AIX) observations



Linux Advantages



■ Community support

◆ Lots of device drivers

- how many do we really need?

◆ Lots of bug fixes

- great when you have to use the same junk hardware that everyone else is using

◆ There are people who know or want to know Linux

■ Application support

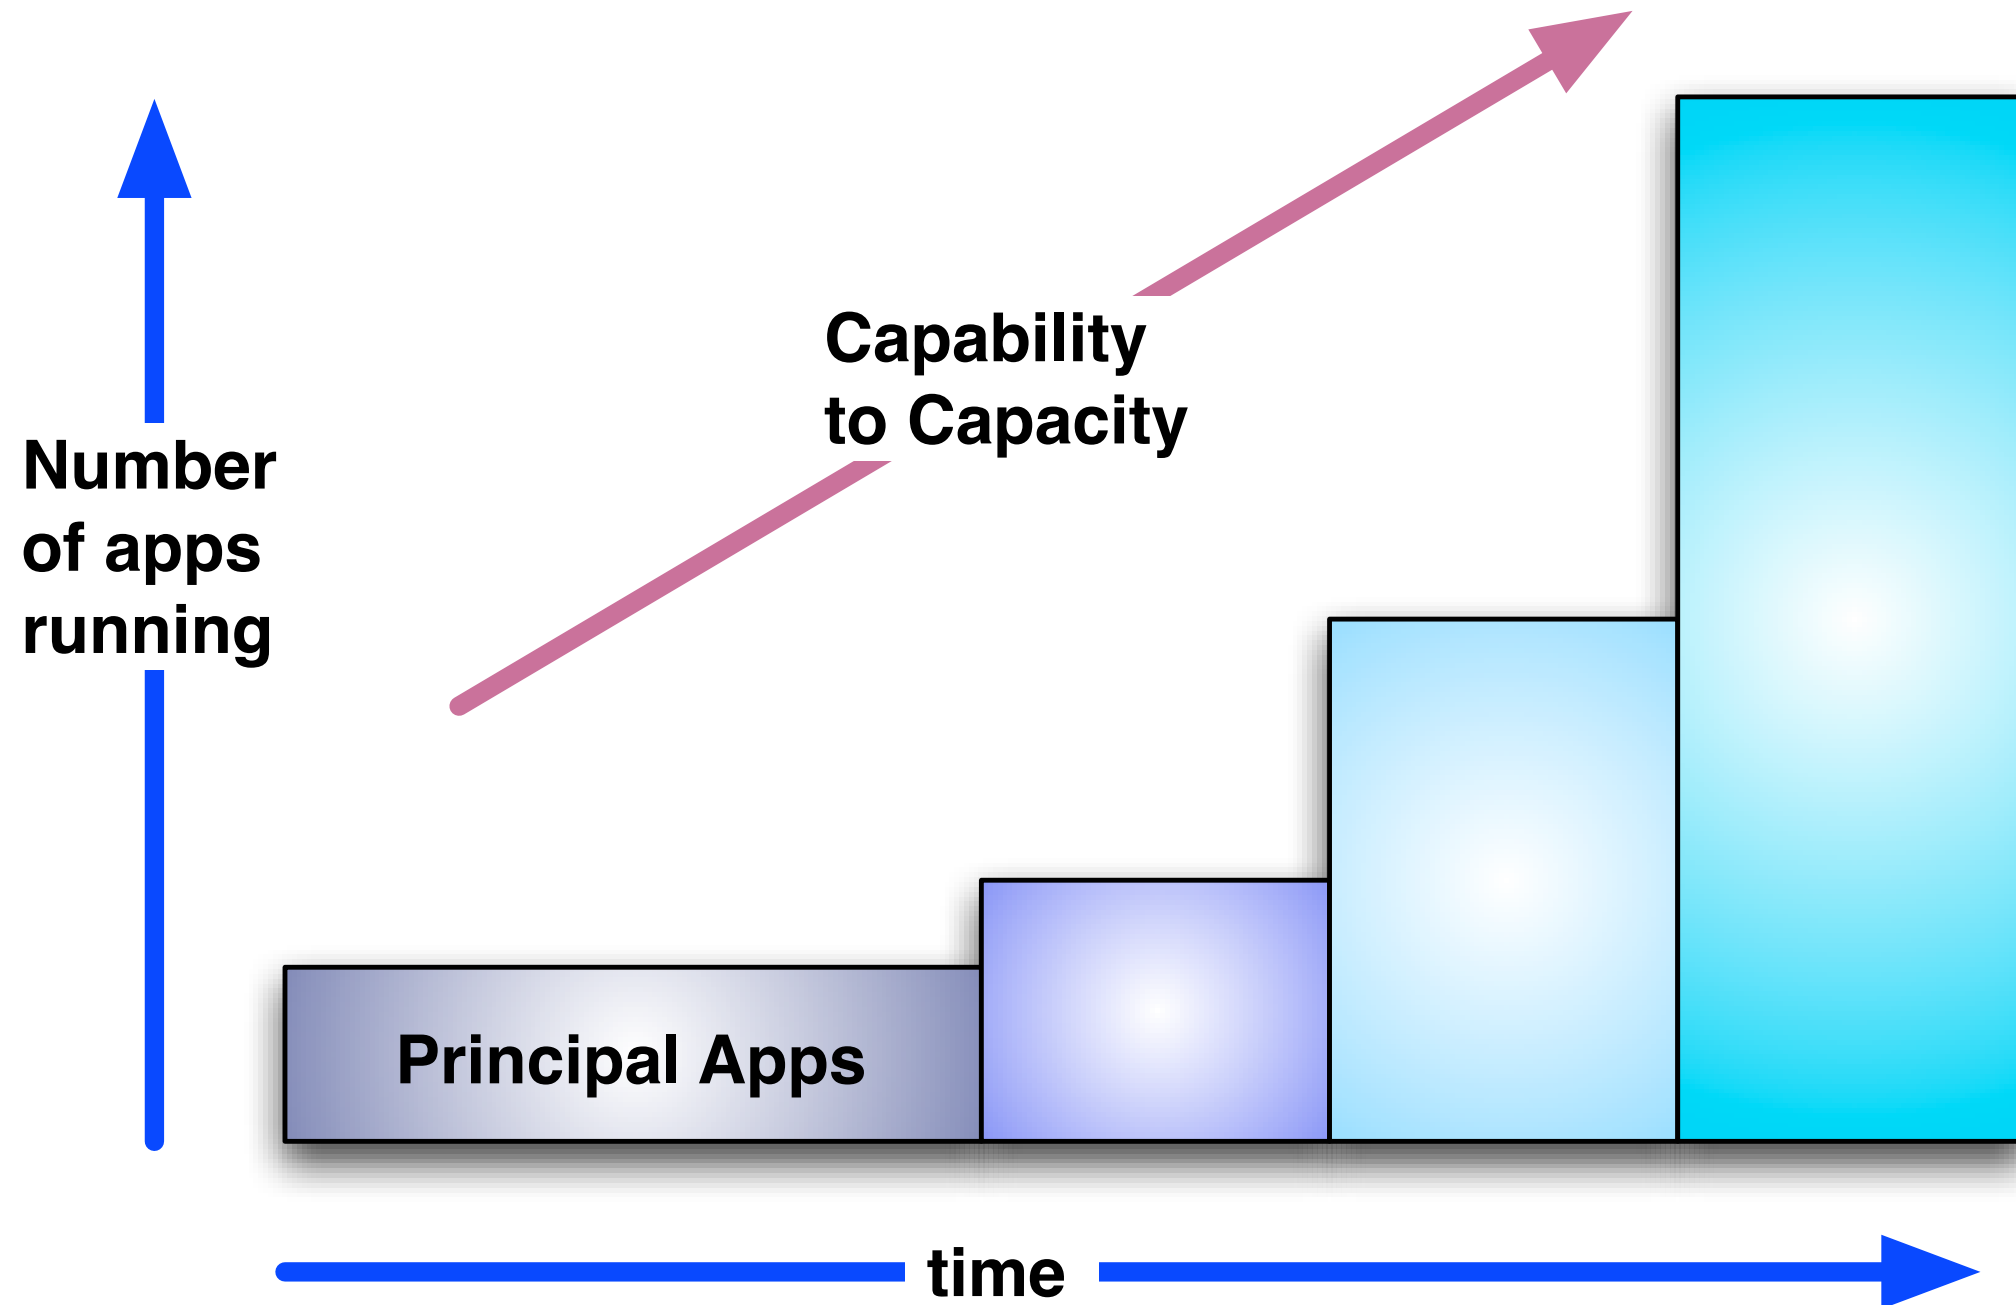
◆ Well understood programming environment

◆ Is it the right model?


- process model



System Lifetime



Factors



■ Lightweight system

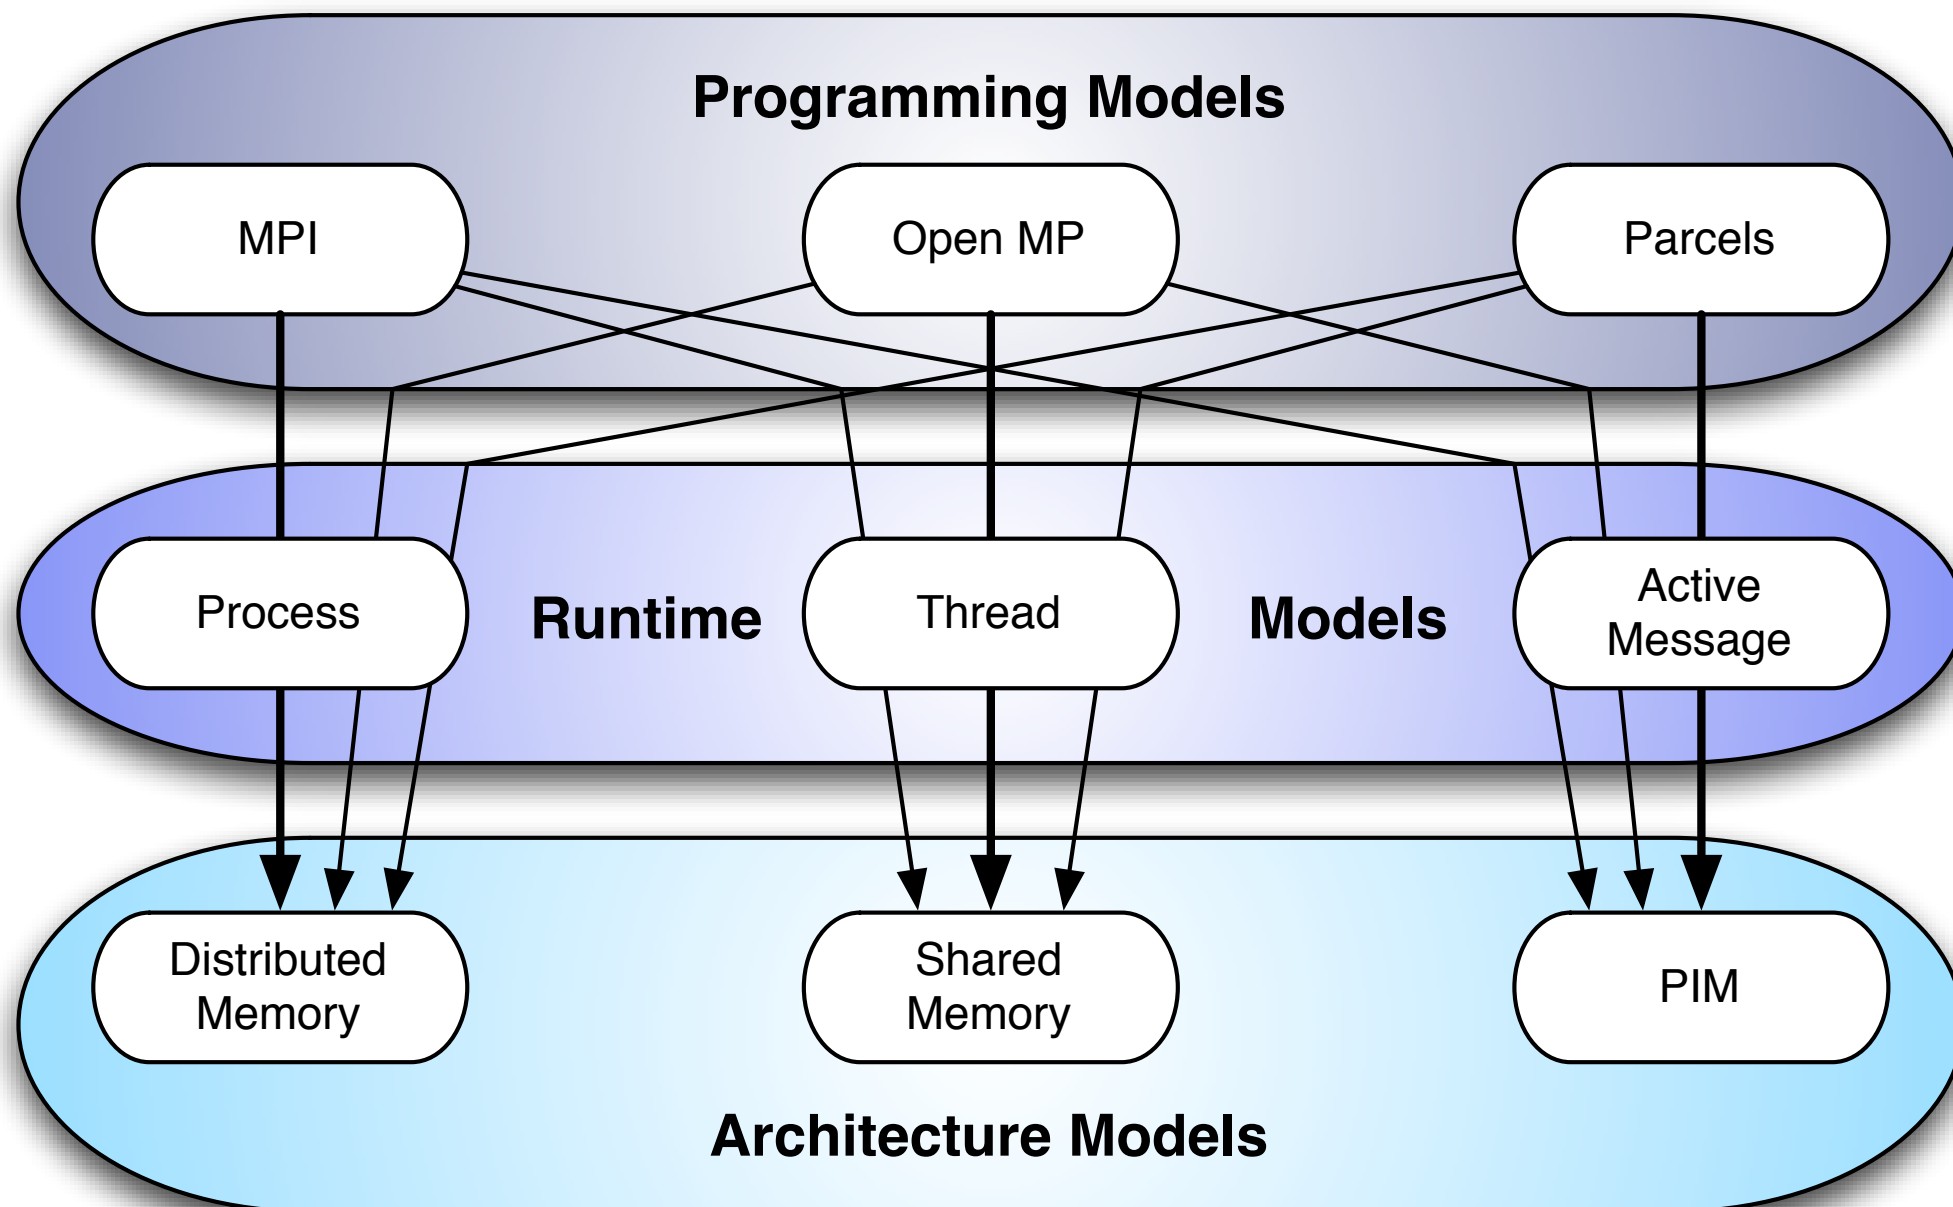
- ◆ small collection of apps
 - single programming model
- ◆ single architecture
- ◆ single usage model
- ◆ small set of shared services
- ◆ no history :)

■ Puma/Cougar

- ◆ MPI
- ◆ Distributed memory
- ◆ space shared
- ◆ parallel file system



Programming Models



Observations



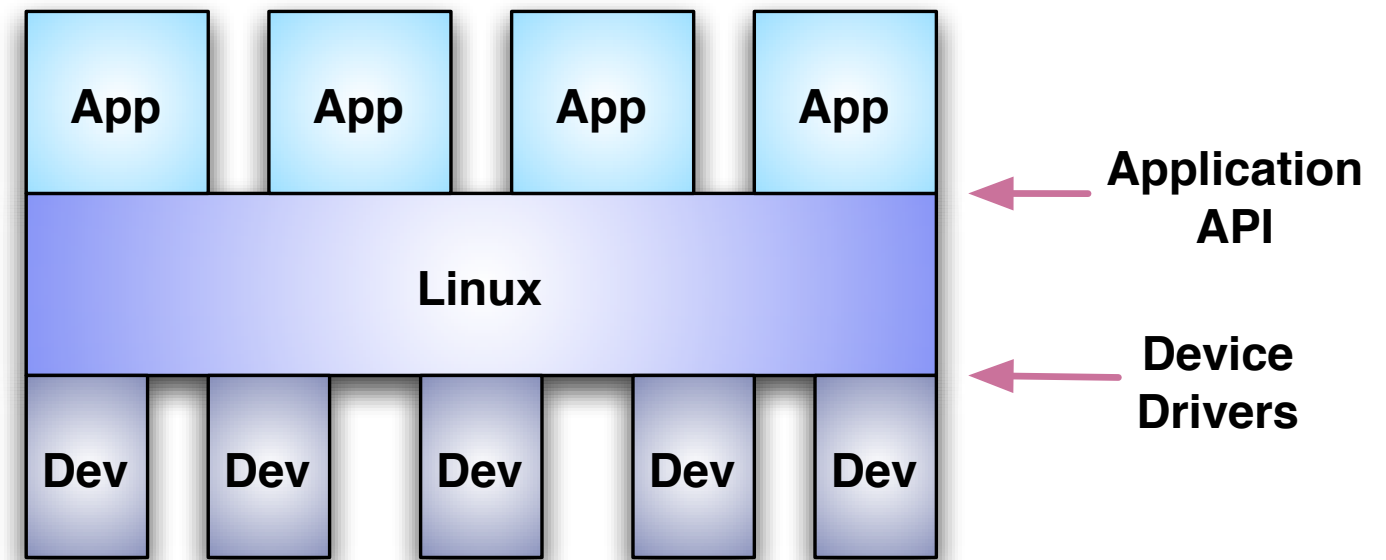
- If Application variety dominates, Linux wins
 - ◆ as long as apps are process based
- If new programming models dominate, Lightweight approaches win
- If you need to use junk hardware, Linux wins
 - ◆ I know I don't want to mess with a PCI bus...
- If Architecture changes dominate, Lightweight approaches win



What is Linux?



- An API for applications
 - ◆ build a Linux personality
 - performance tradeoffs will be different



- A collection of device drivers
 - ◆ runs on a variety of systems
 - as long as you can live with an x86 memory model
 - lots of support for broken hardware



Concluding Thoughts



- Lightweight approaches are the best way to support the development of novel architectures and programming models

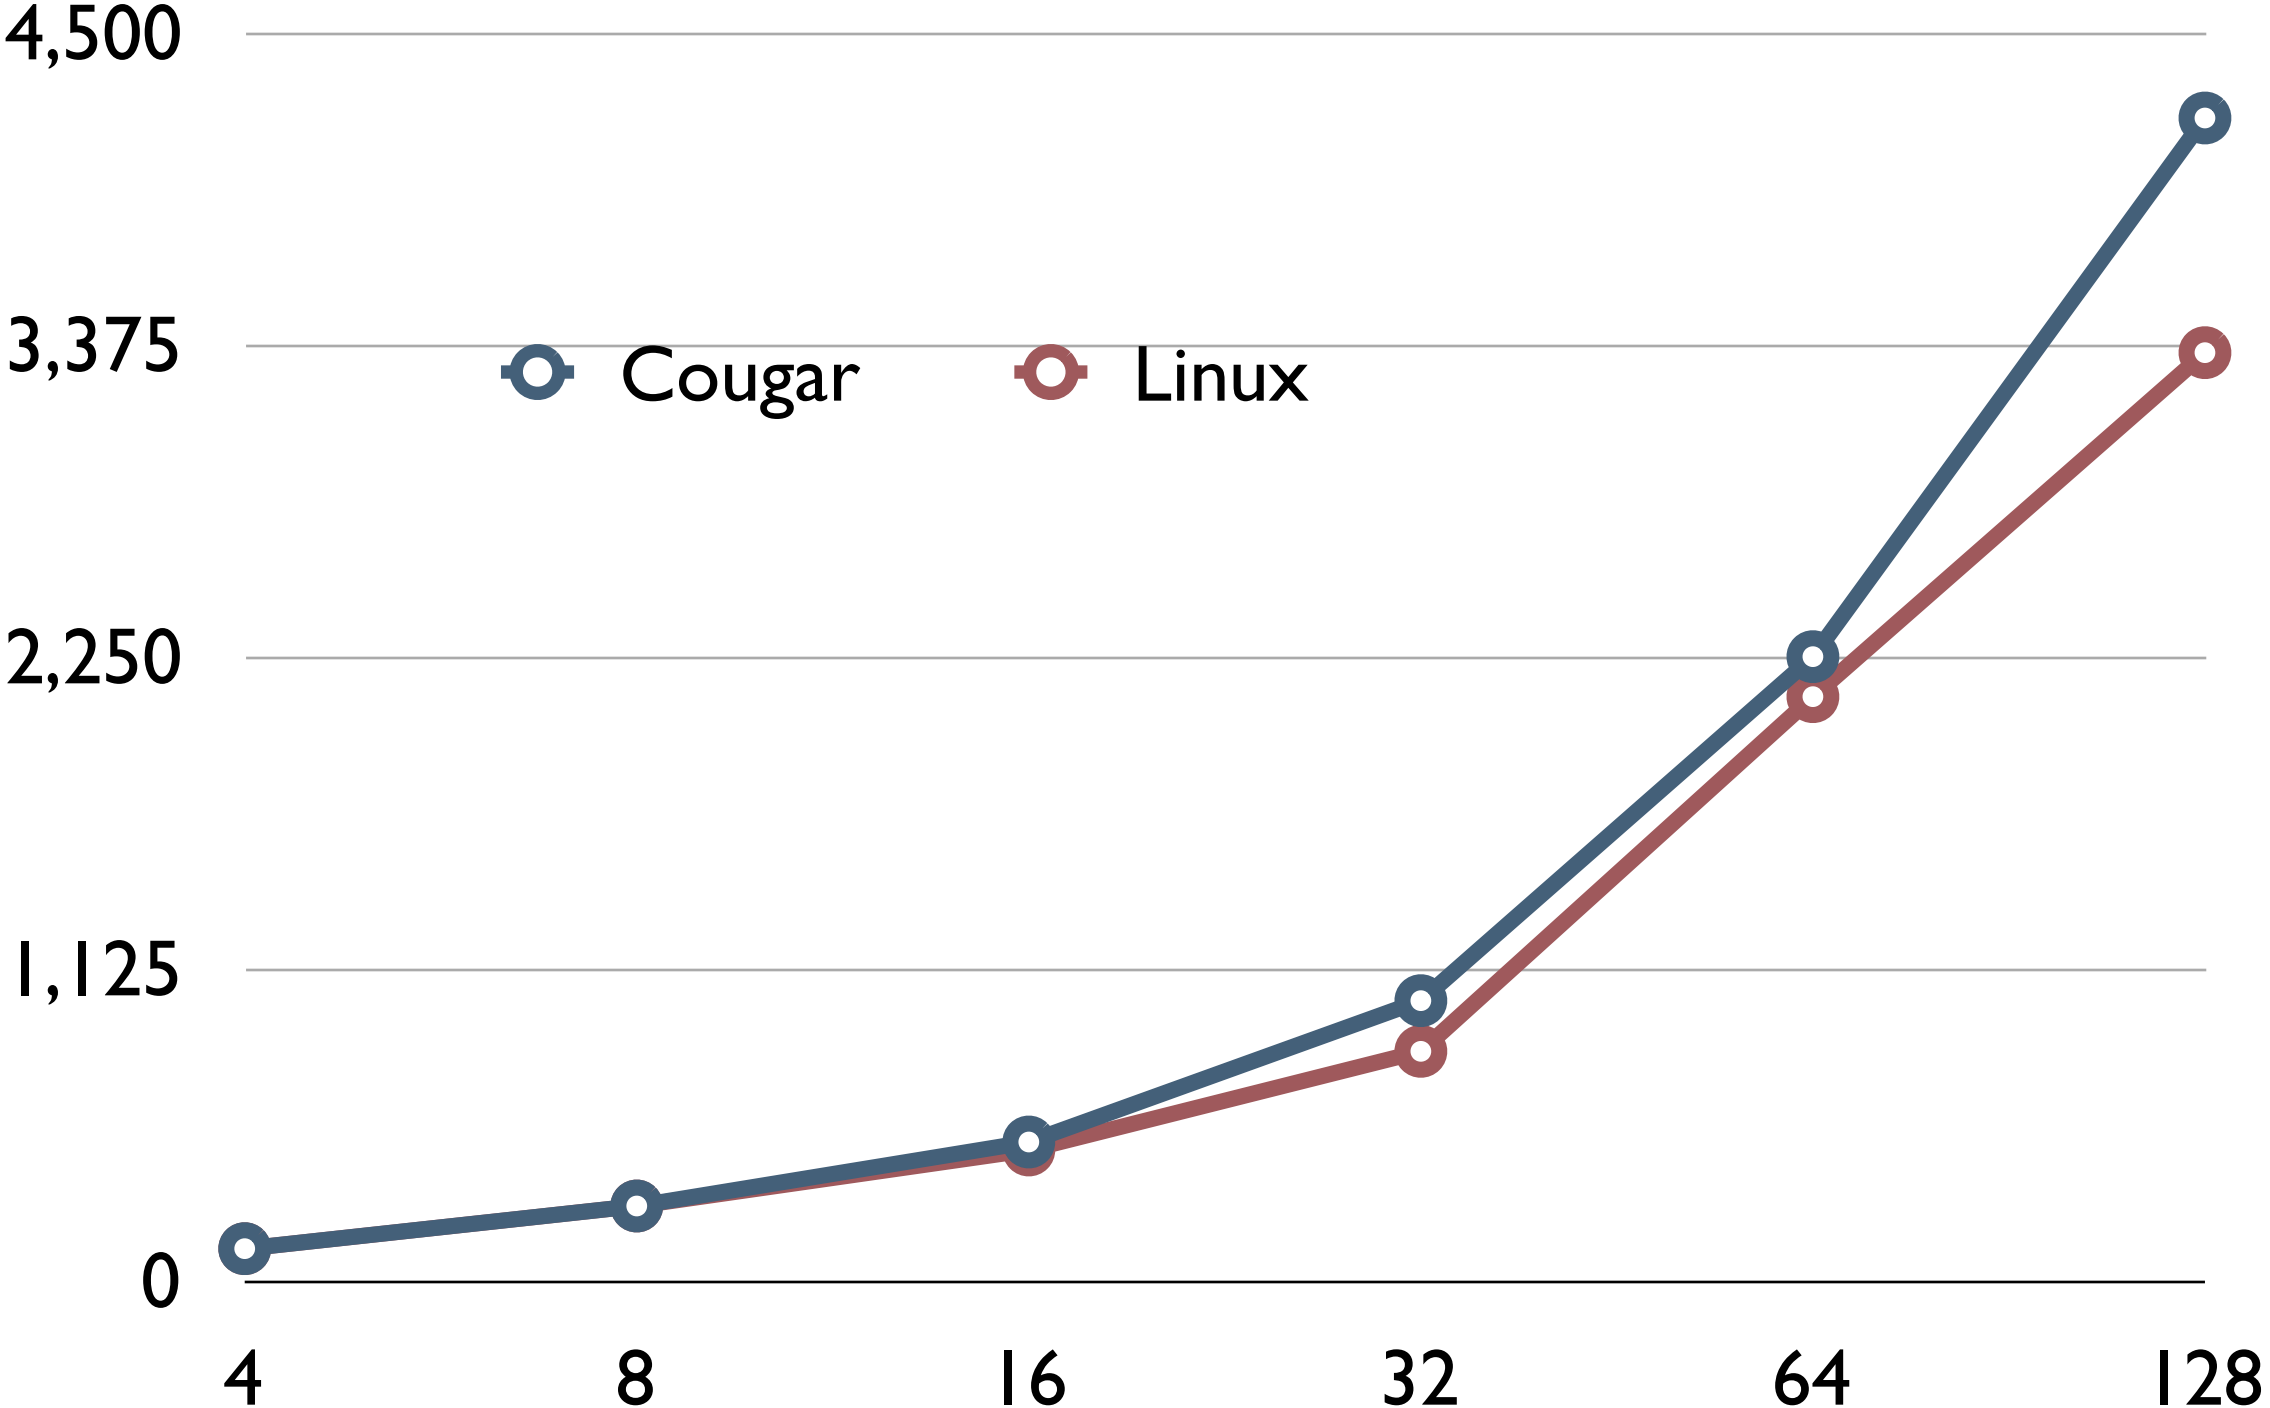
- There are lots of other things that need to be done
 - ◆ better support for single system image on clusters
 - ◆ clean up internal Linux structure

- Testbed systems
 - ◆ systems is an experimental discipline
 - ◆ need to be able to perform true apples to apples comparisons





NAS B: mg



NAS B: cg

