# Condor-
## a Project and a System

Miron Livny
Computer Sciences Department
University of Wisconsin-Madison
miron@cs.wisc.edu

# The Condor Project (Established '85)

Distributed Computing **research** performed by a team of **~40** faculty, full time staff and students who

- face **software/middleware engineering** challenges in a UNIX/Linux/Windows/OS X environment,
- involved in national and international **collaborations**,
- interact with **users** in academia and industry,
- maintain and support a distributed **production** environment (more than 3300 CPUs at UW),
- and educate and train **students**.

**Funding** – DoE, NASA, NIH, NSF, EU, INTEL, Micron, Microsoft and the UW Graduate School

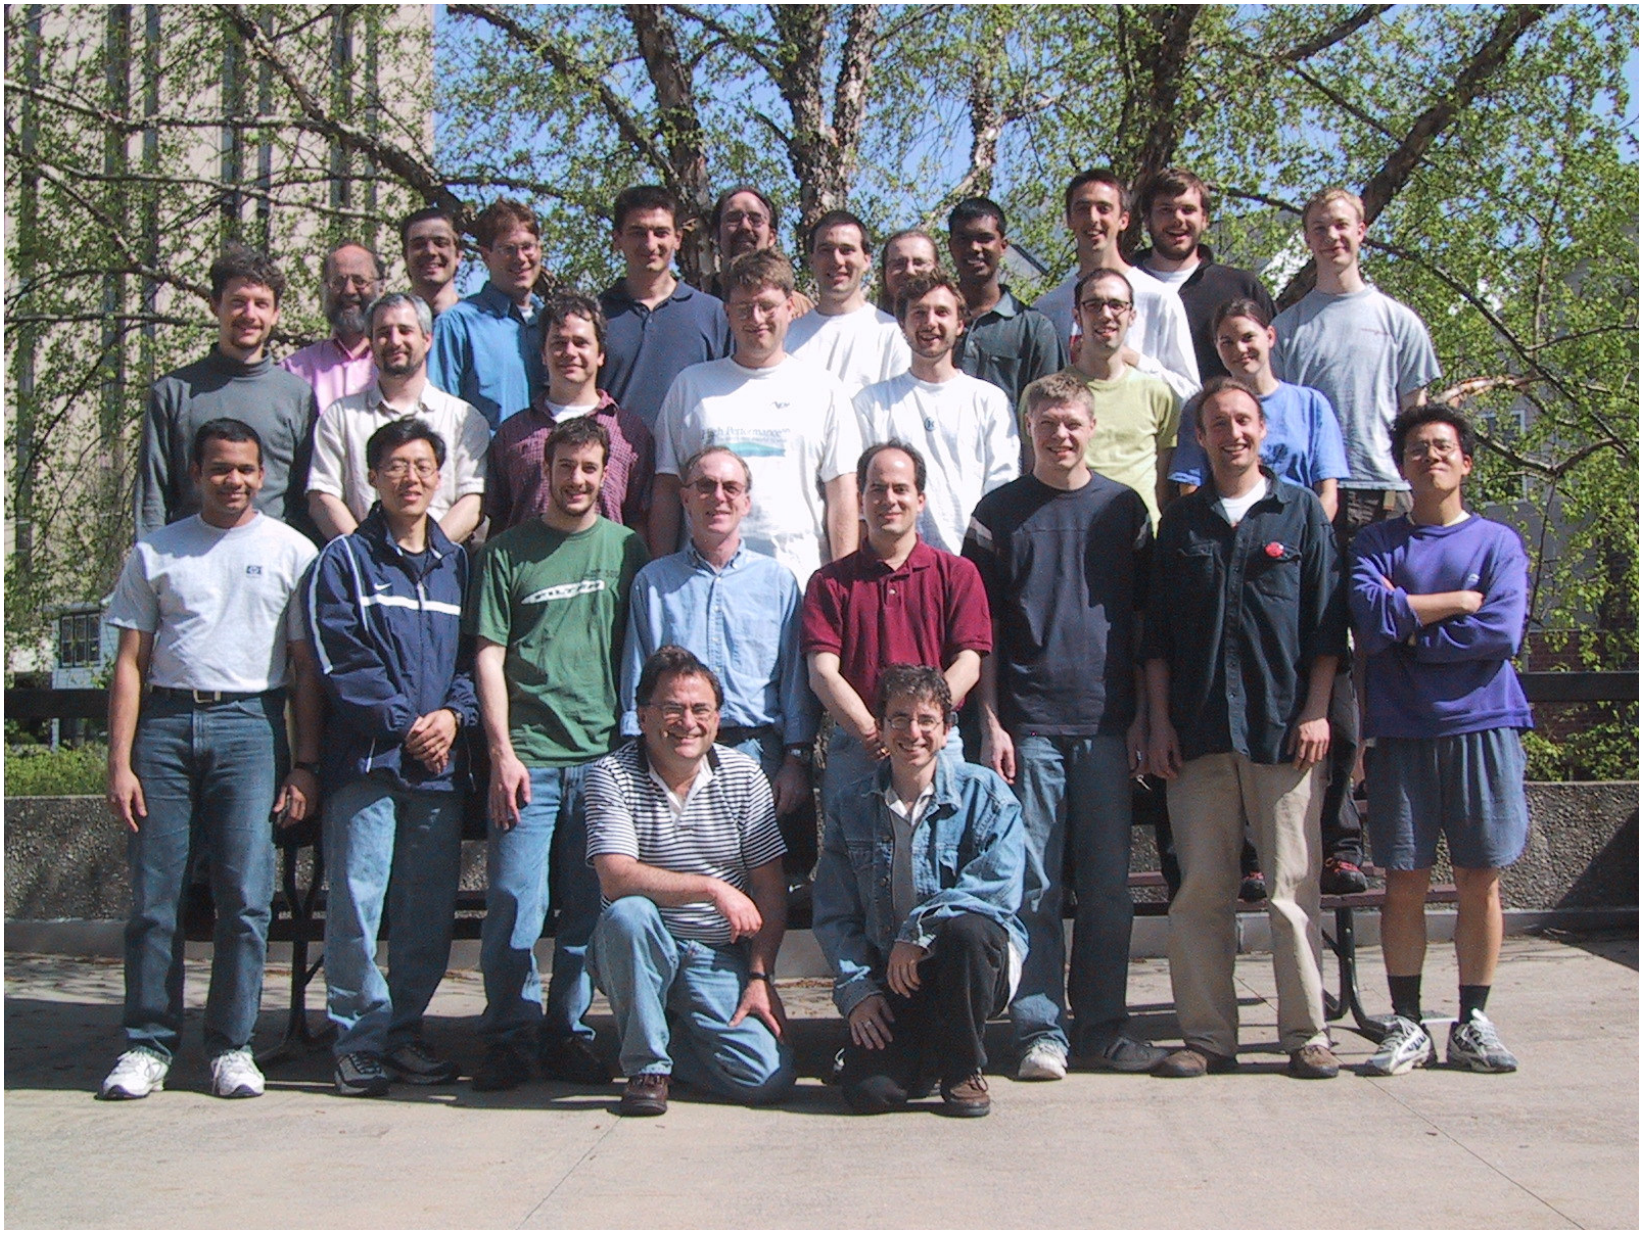

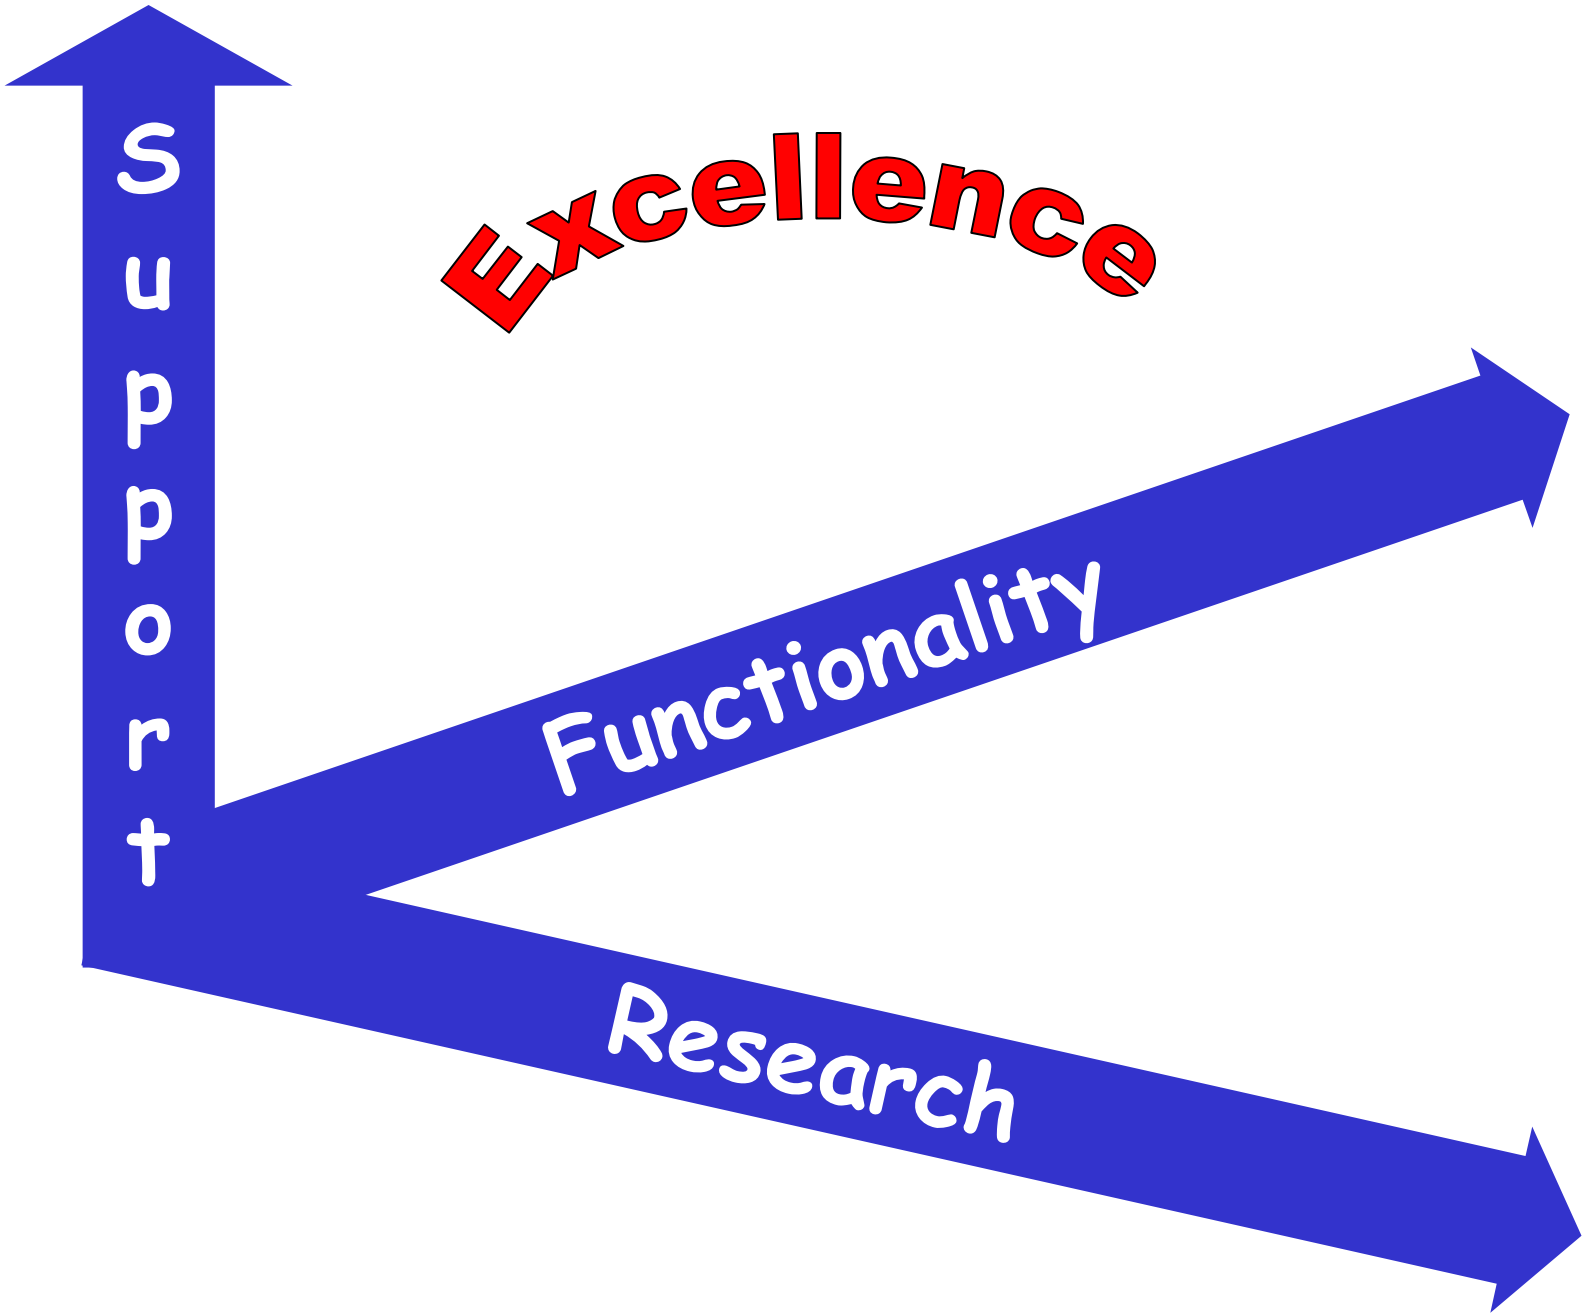
Excellence
Support
Functionality
Research


*our*
*answer to*
*High Throughput MW Computing*
*on commodity resources*

# Novel

# The Layers of Condor

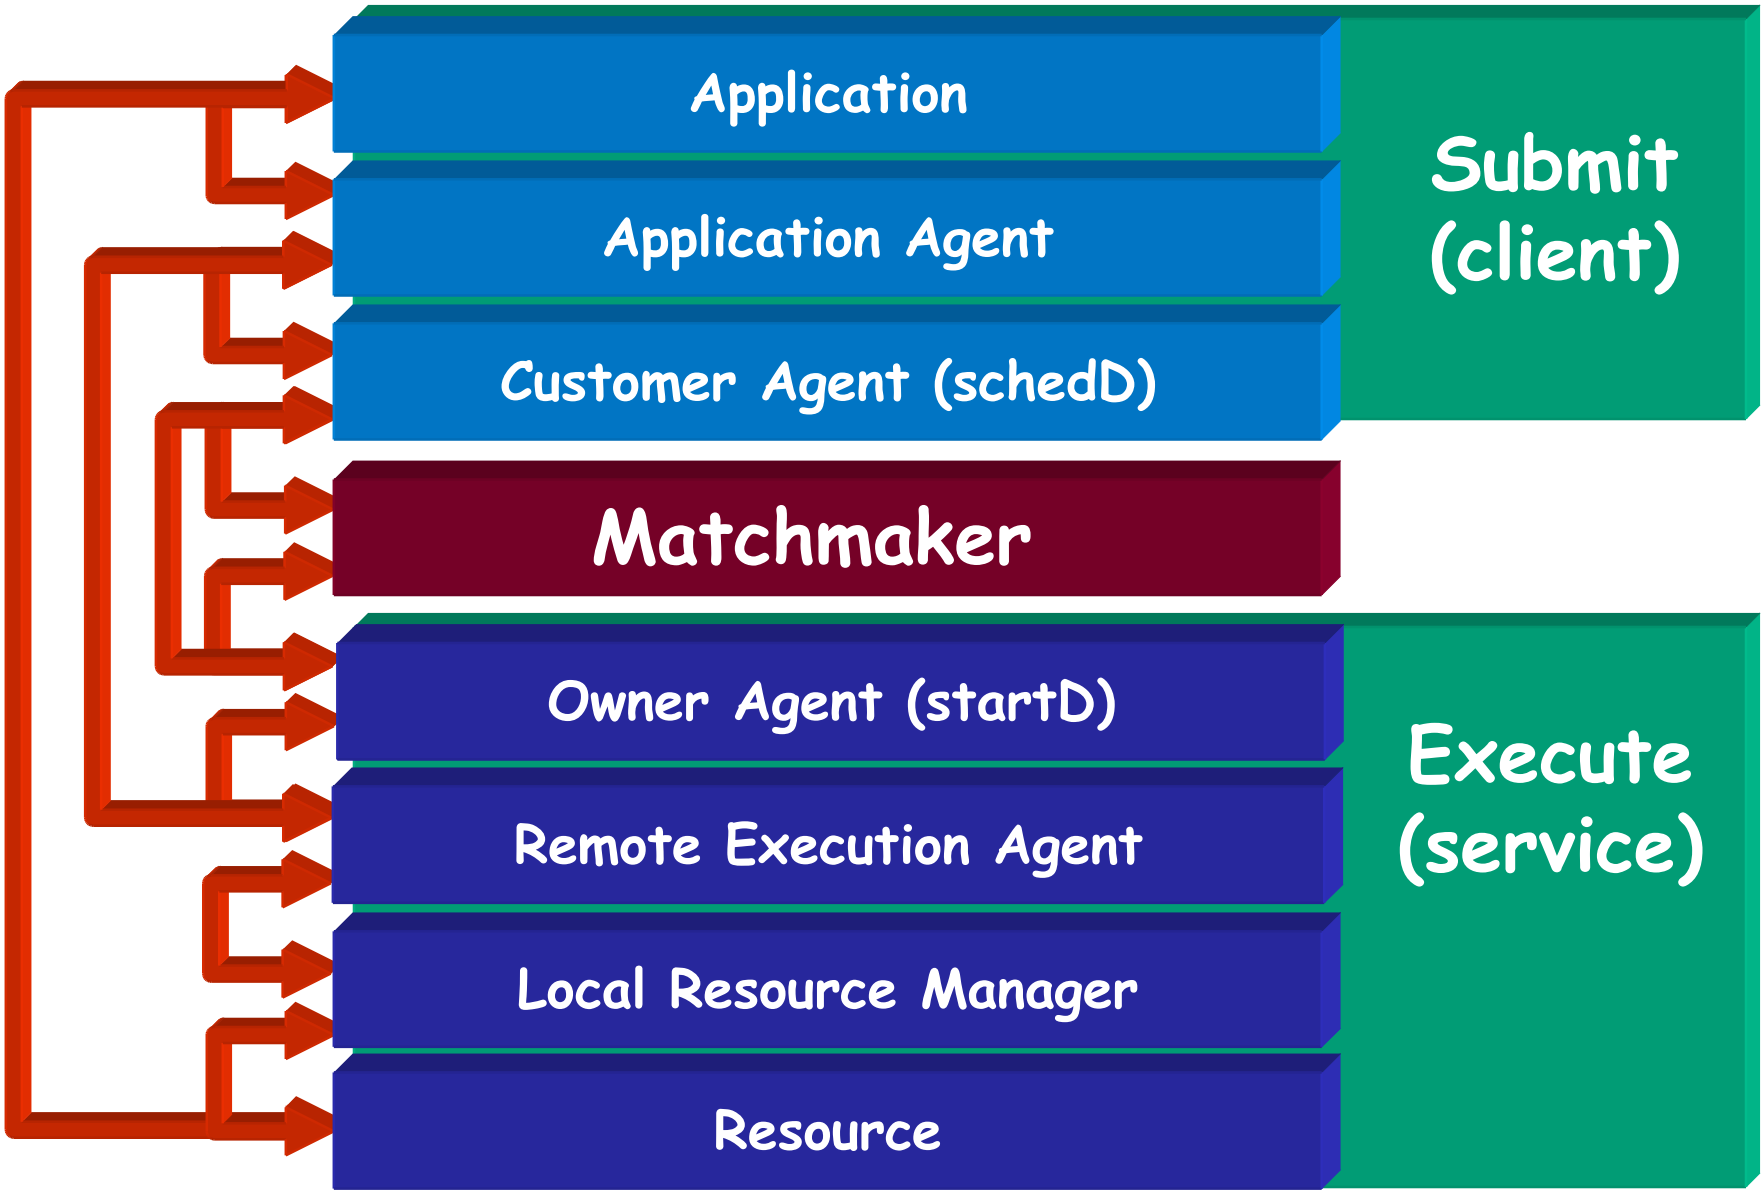


desktop

# Yearly Condor usage at UW-CS

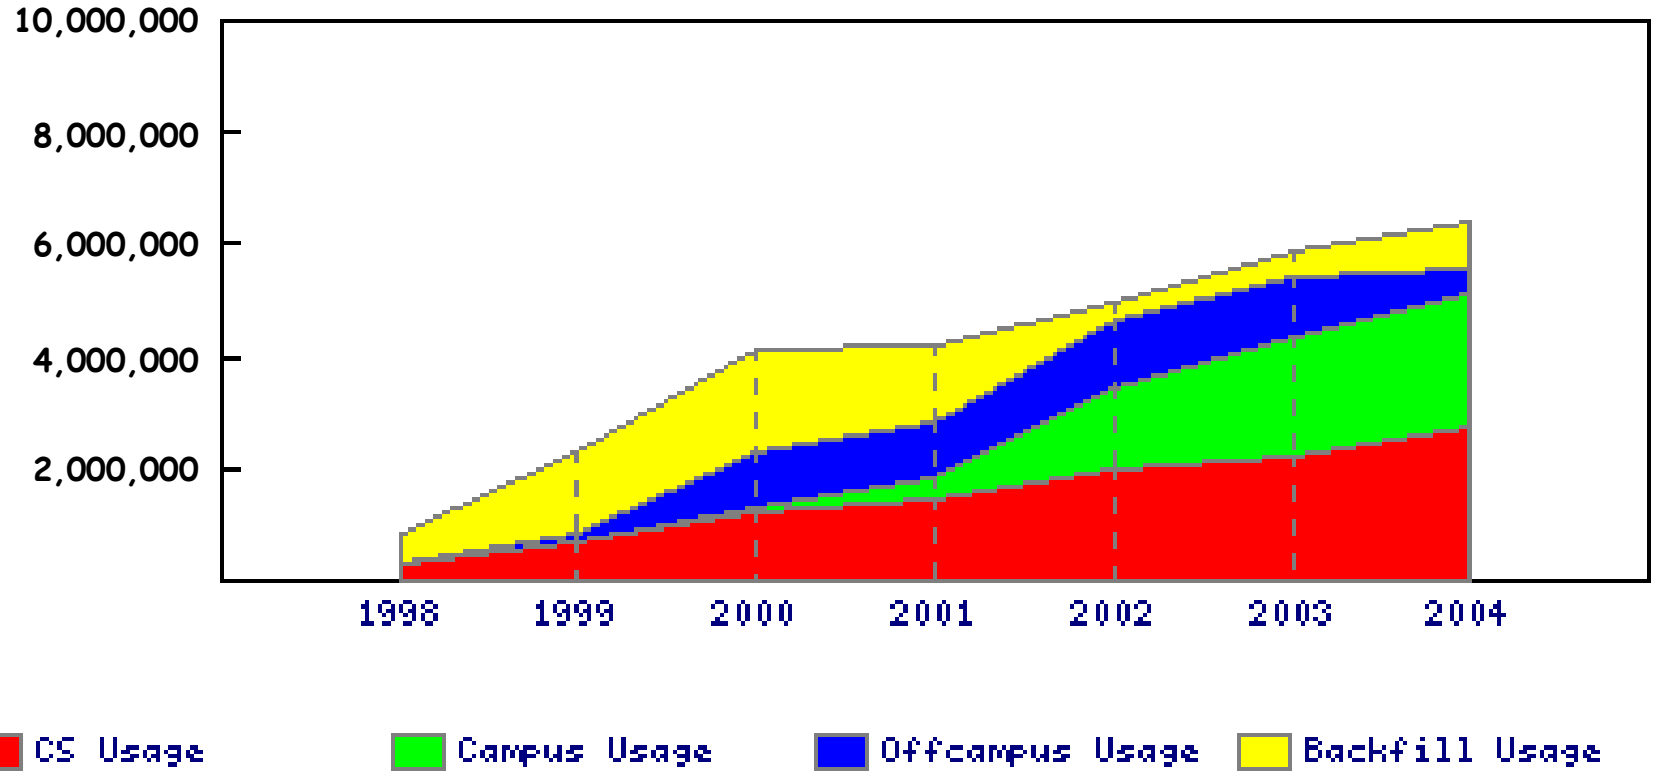

# Yearly Condor CPUs at UW

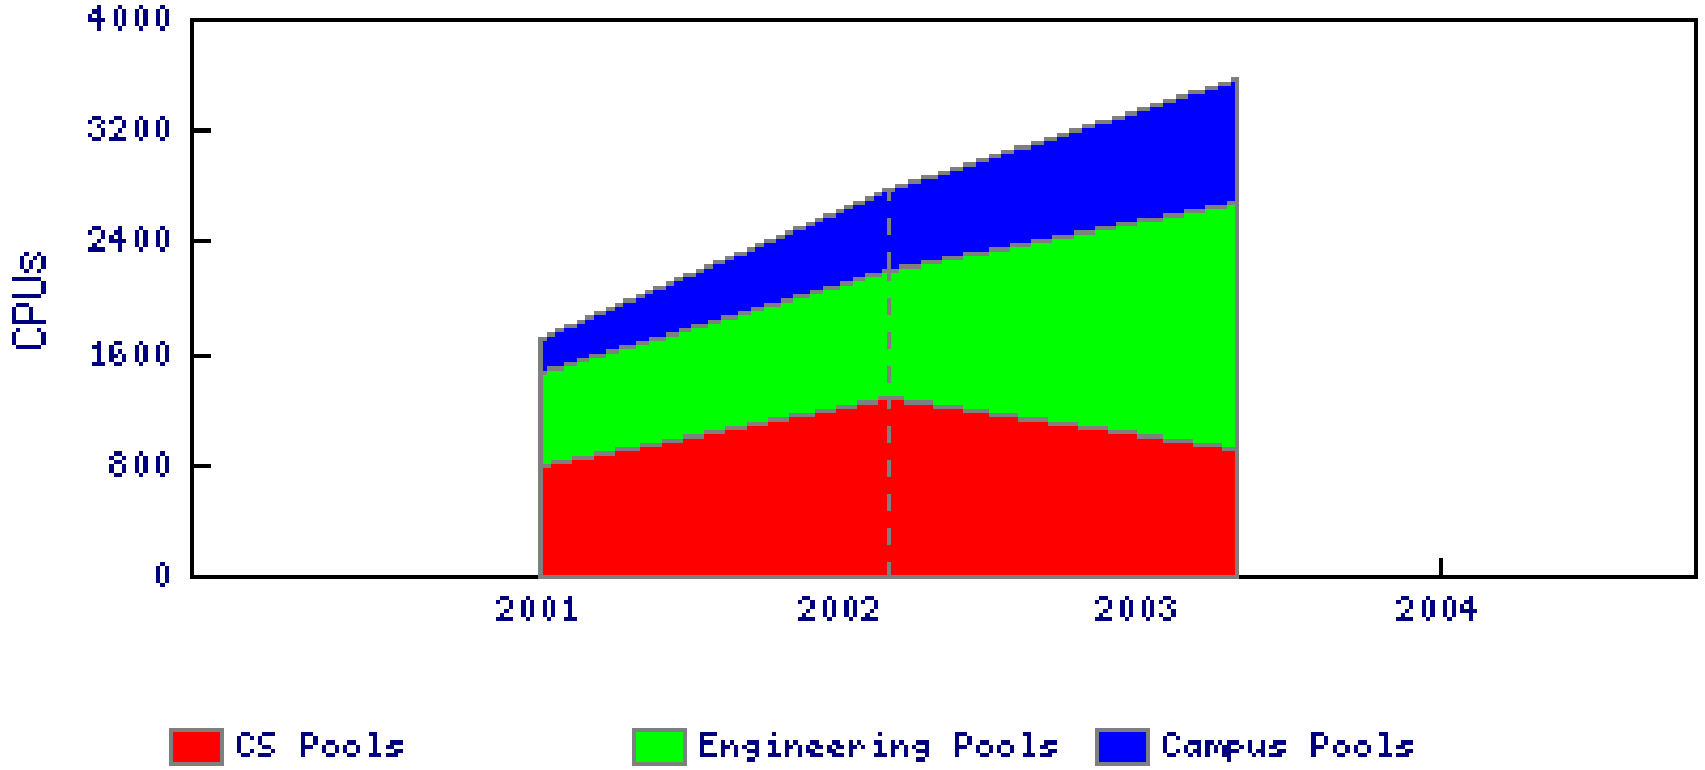

# Flexible

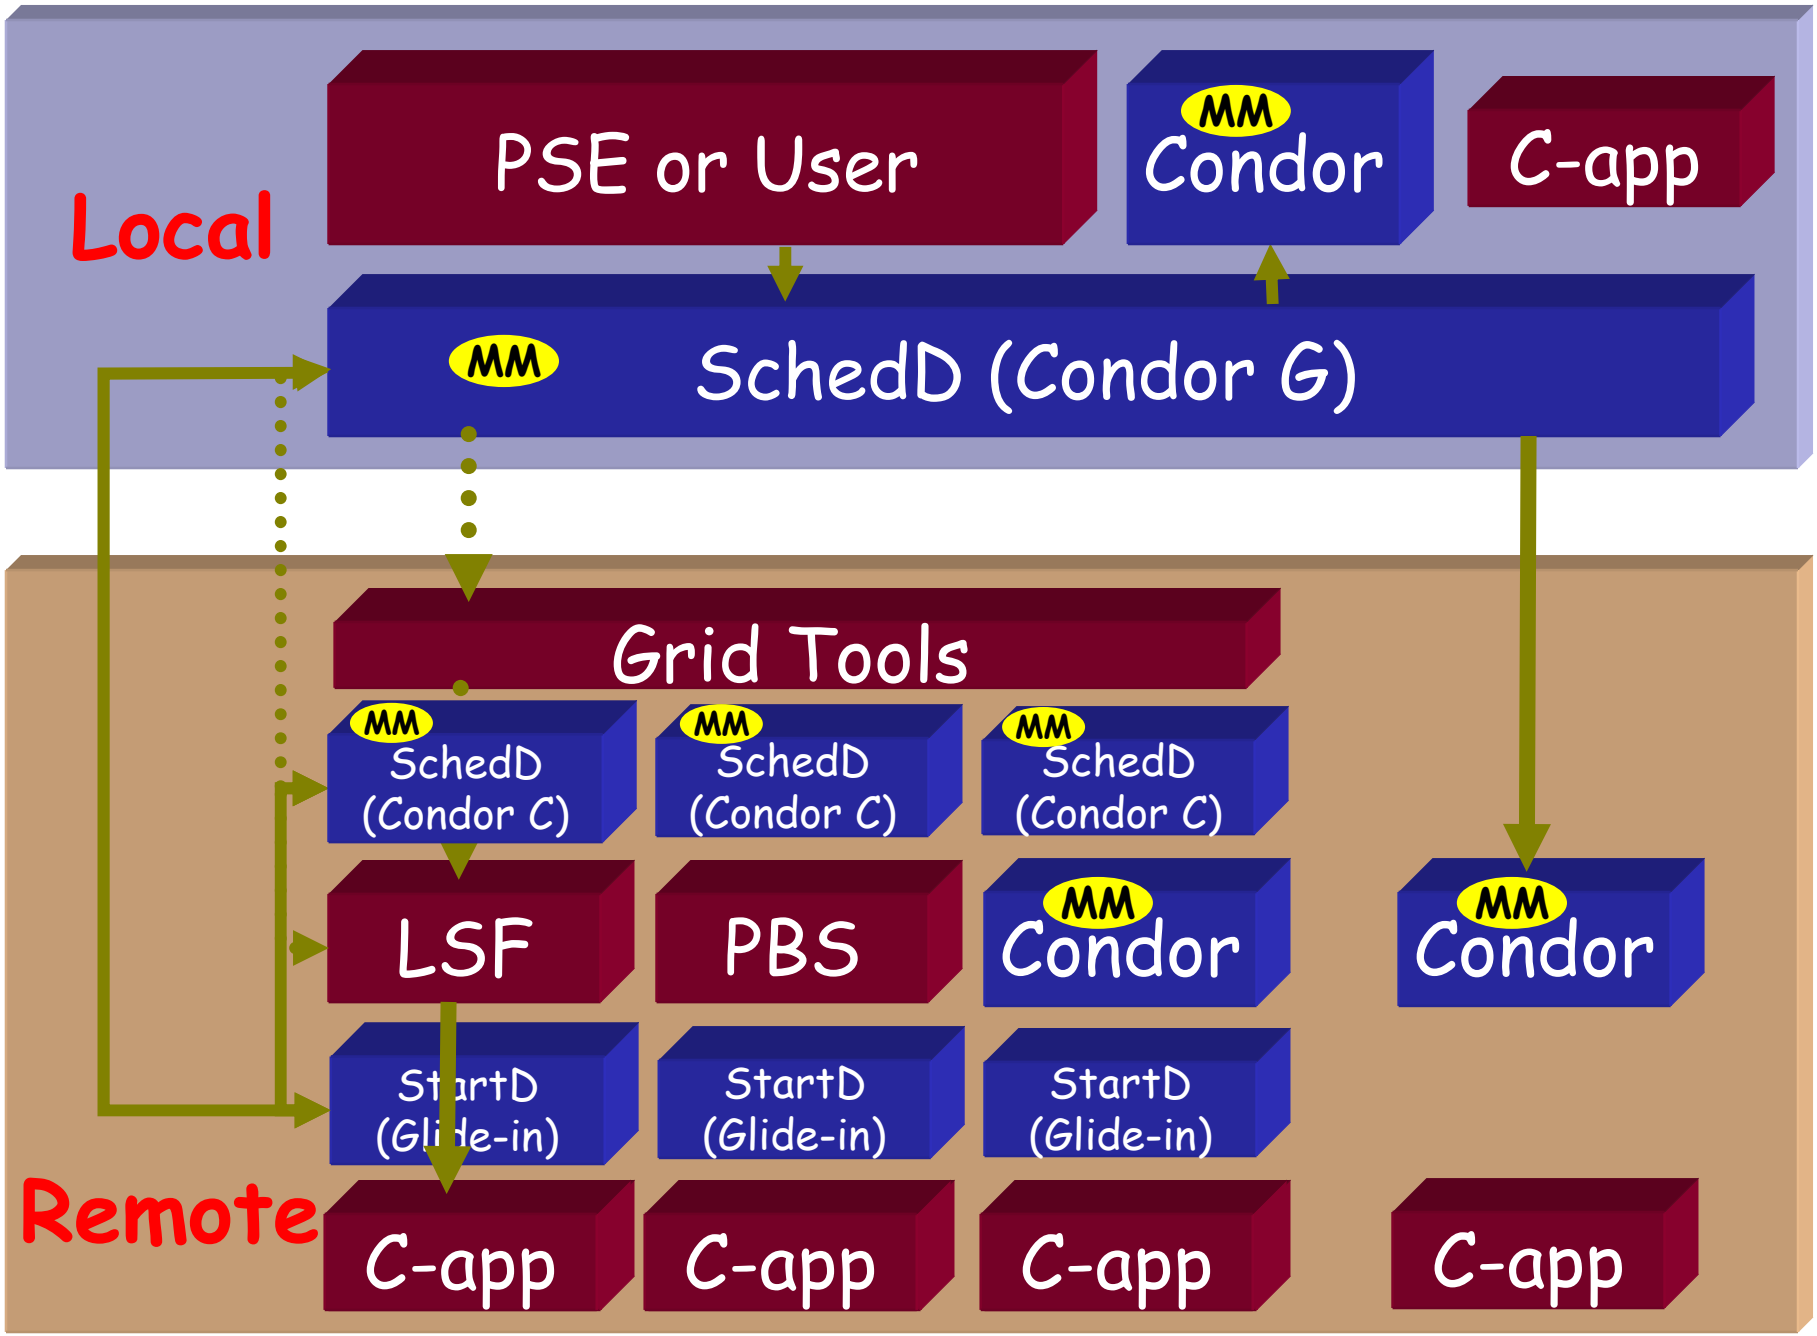
Local
PSE or User
MM
Condor
C-app
MM
SchedD (Condor G)
Grid Tools
MM
SchedD
(Condor C)
MM
SchedD
(Condor C)
MM
SchedD
(Condor C)
LSF
PBS
MM
Condor
MM
Condor
StartD
(Glide-in)
StartD
(Glide-in)
StartD
(Glide-in)
Remote
C-app
C-app
C-app
C-app

# Robust

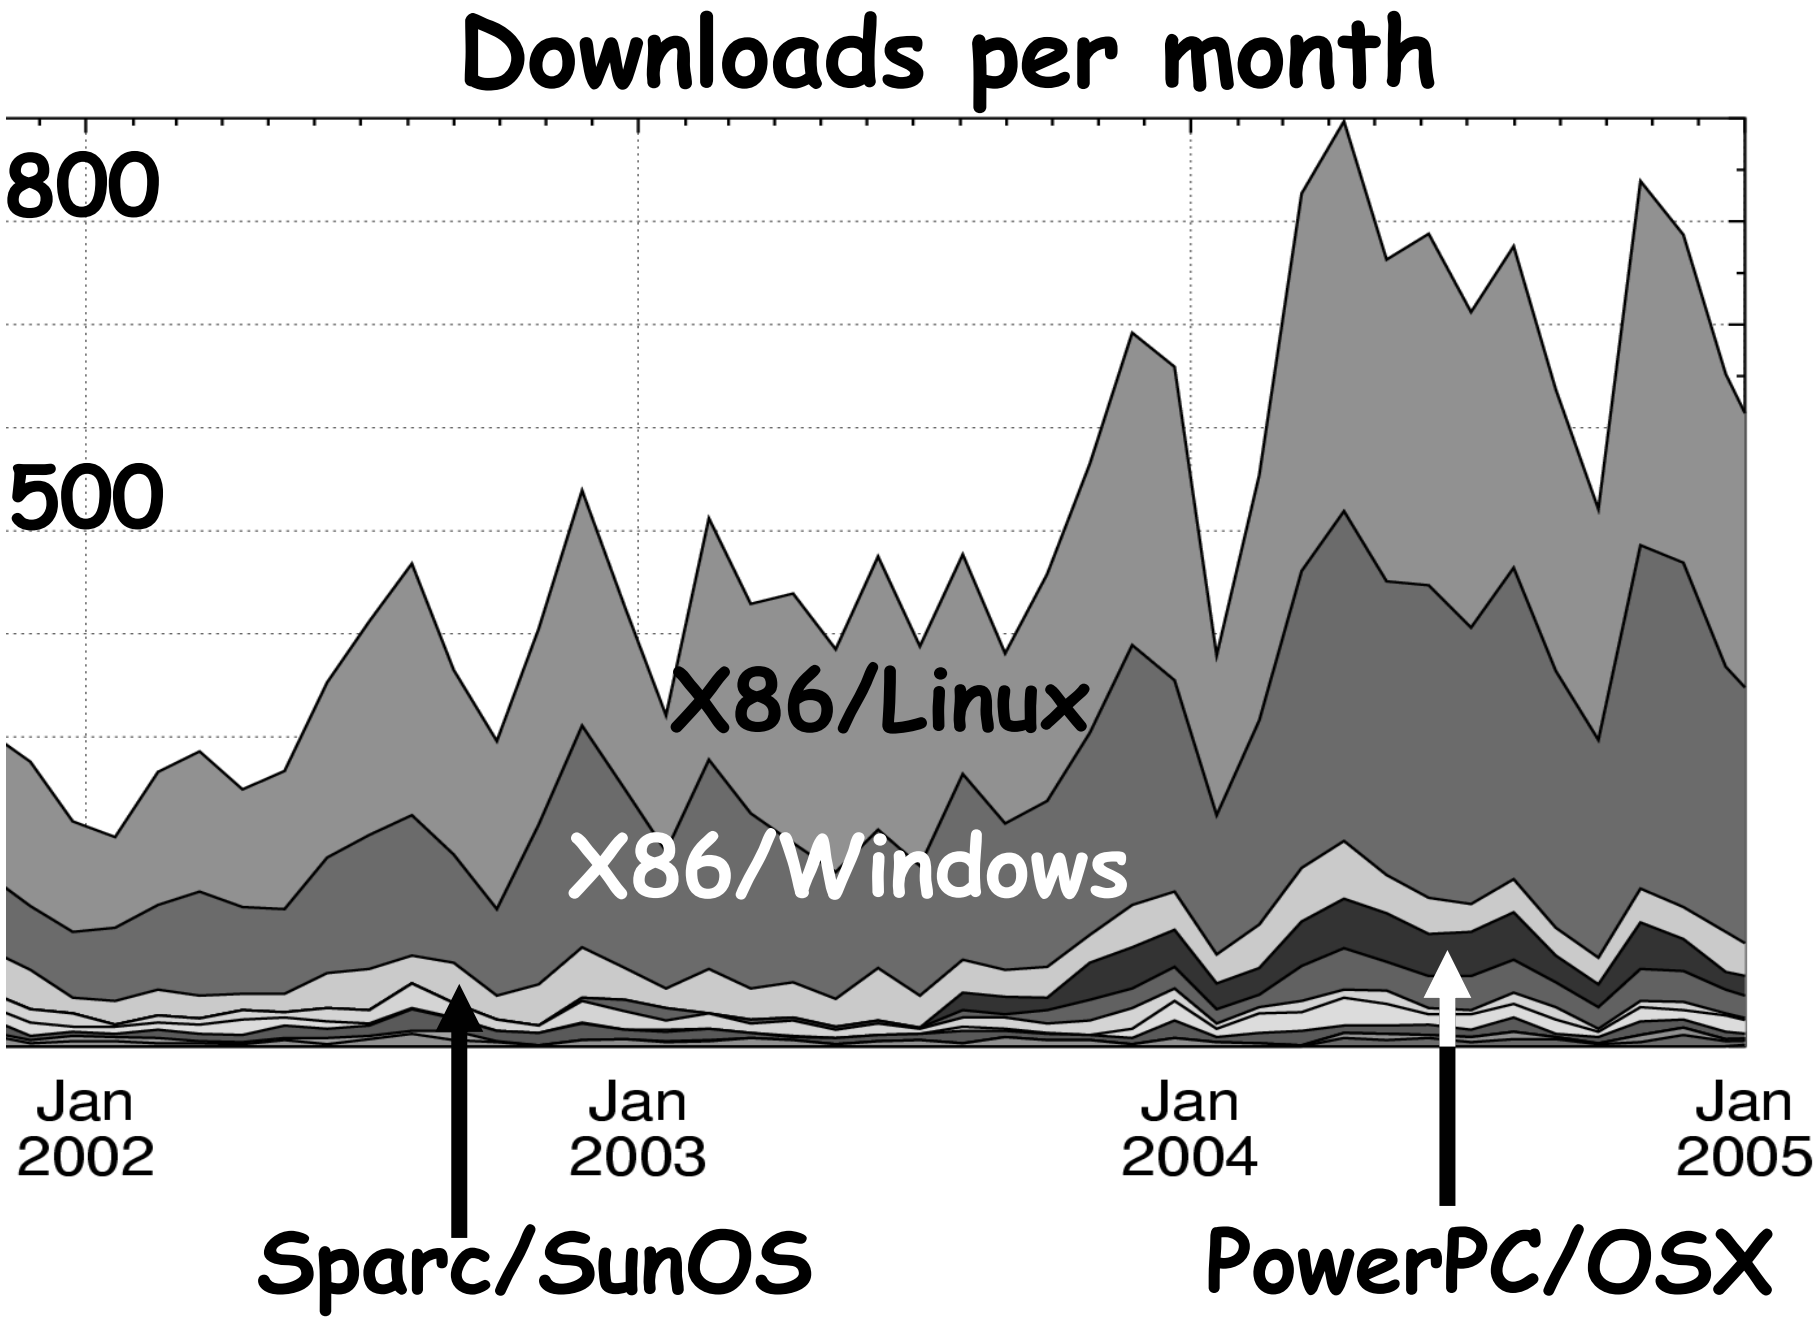
Downloads per month
800
500
X86/Linux
X86/Windows
Jan
2002
Jan
2003
Jan
2004
Jan
2005
Sparc/SunOS
PowerPC/OSX

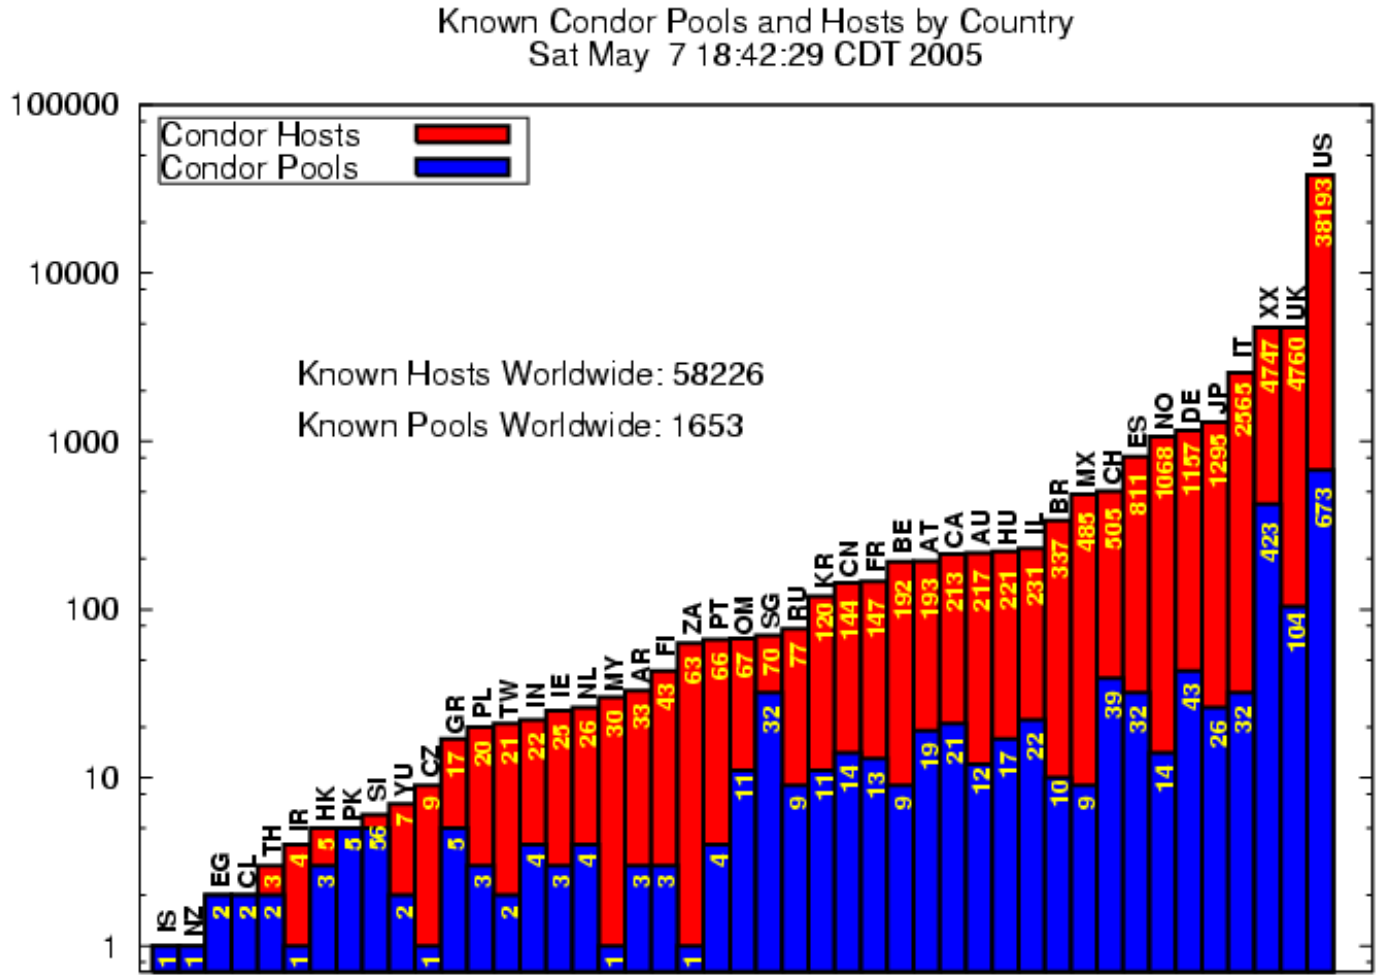
Known Condor Pools and Hosts by Country
Sat May 7 18:42:29 CDT 2005
Condor Hosts
Condor Pools
Known Hosts Worldwide: 58226
Known Pools Worldwide: 1653
100000
10000
1000
100
10
1
IS 1 1
NZ 2 2
EG 2 2
CL 2 3
TH 1 4
IR 3 5
HK 5 5
PK 5 5
SI 5 6
YU 2 7
CZ 1 9
GR 5 17
PL 3 20
TW 2 21
IN 4 22
IE 3 25
NL 4 26
MY 1 30
AR 3 33
FI 3 43
ZA 1 63
PT 4 66
OM 11 67
SG 32 70
RU 9 77
KR 11 120
CN 14 144
FR 13 147
BE 9 192
AT 19 193
CA 21 213
AU 12 217
HU 17 221
IL 22 231
BR 10 337
MX 9 485
CH 39 505
ES 32 811
NO 14 1068
DE 43 1157
JP 26 1295
IT 32 2565
XX 423 4747
UK 104 4760
US 673 38193

Seeking the massive computing power needed to hedge a portion of its book of annuity business, **Hartford Life**, a subsidiary of The Hartford Financial Services Group (Hartford; $18.7 billion in 2003 revenues), has implemented a grid computing solution based on the **University of Wisconsin's (Madison, Wis.) Condor open source software**. Hartford Life's SVP and CIO Vittorio Severino notes that the move was a matter of necessity. "It was the necessity to hedge the book," owing in turn to a **tight reinsurance market that is driving the need** for an alternative risk management strategy, he says. The challenge was to support the risk generated by clients opting for income protection benefit riders on popular annuity products.

***Resource: How did you complete this project—on your own or with a vendors help?***

**Severino**: We completed this project very much on our own. As a matter of fact it is such a new technology in the insurance industry, that others were calling us for assistance on how to do it. So it was interesting because we were breaking new ground and vendors really couldn't help us. We eventually chose grid computing software from the University of Wisconsin called Condor; it is open source software. **We chose the Condor software because it is one of the oldest grid computing software tools around; so it is mature. We have a tremendous amount of confidence in the Condor software**

# Condor at Micron

10,000+ processors in 12 “pools”
Linux, Solaris, Windows
<50th Top500 Rank
3+ TeraFLOPS

Centralized governance
Distributed management

16+ applications
Self developed

Micron’s Global Grid

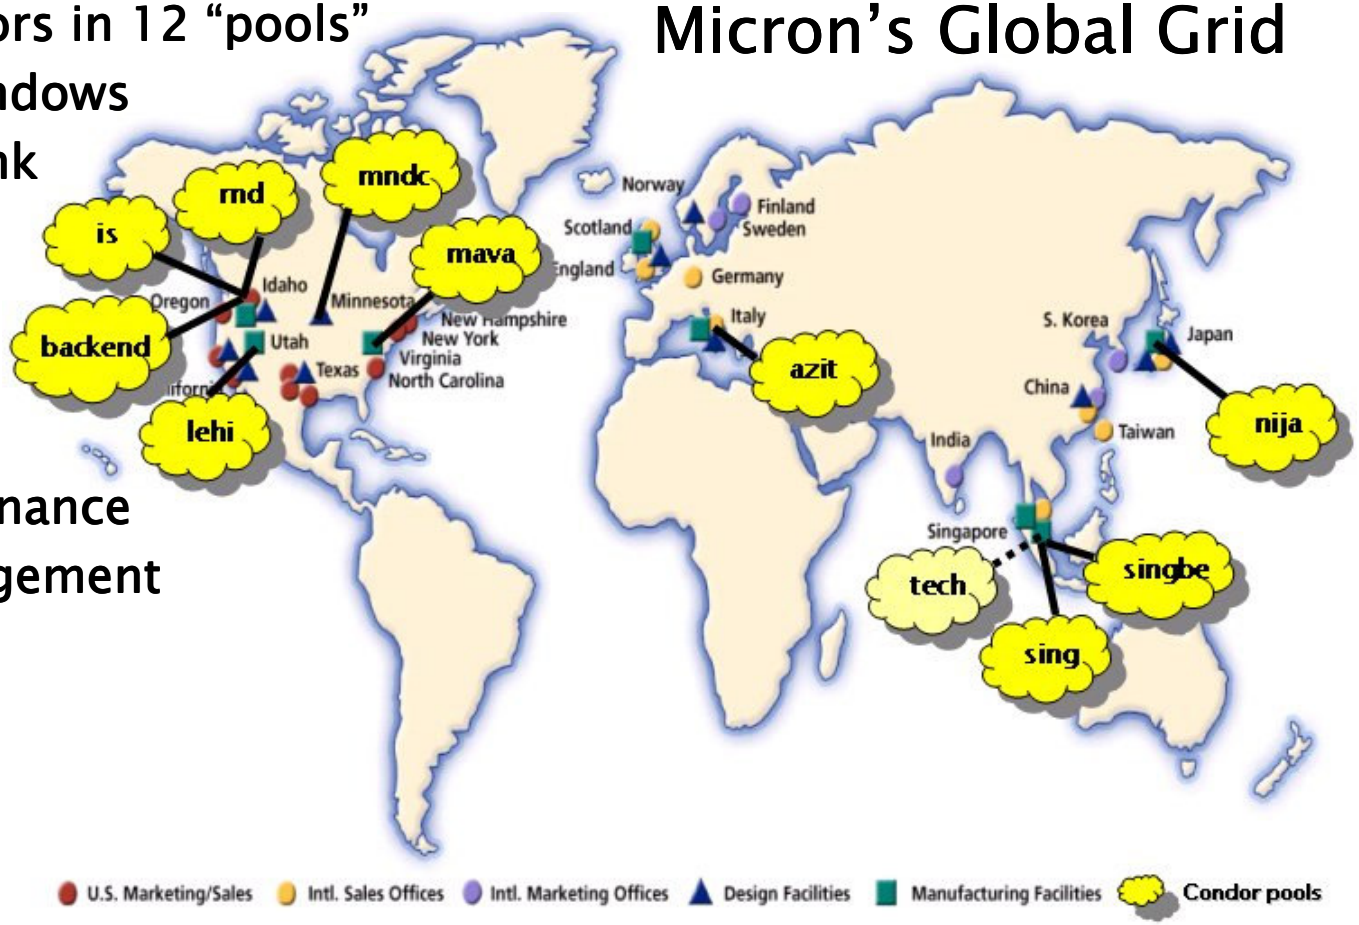

# Condor at Oracle

Condor is used within Oracle's Automated Integration Management Environment (AIME) to perform automated build and regression testing of multiple components for Oracle's flagship Database Server product.

Each day, nearly 1,000 developers make contributions to the code base of Oracle Database Server. Just the compilation alone of these software modules would take over 11 hours on a capable workstation. But in addition to building, AIME must control repository labelling/tagging, configuration publishing, and last but certainly not least, regression testing. Oracle is very serious about the stability and correctness about their products. Therefore, the AIME daily regression test suite currently covers 90,000 testable items divided into over 700 test packages. The entire process must complete within 12 hours to keep development moving forward.

About five years ago, Oracle selected Condor as the resource manager underneath AIME because they liked the maturity of Condor's core components. In total, over 3,500 machines at Oracle are managed by Condor.

# Laboratory of Molecular and Computational Genomics
## University of Wisconsin-Madison

Our research laboratory focuses on the chemistry, biology and physics of single DNA molecules as a means of genomic analysis.


www.cs.wisc.edu/condor

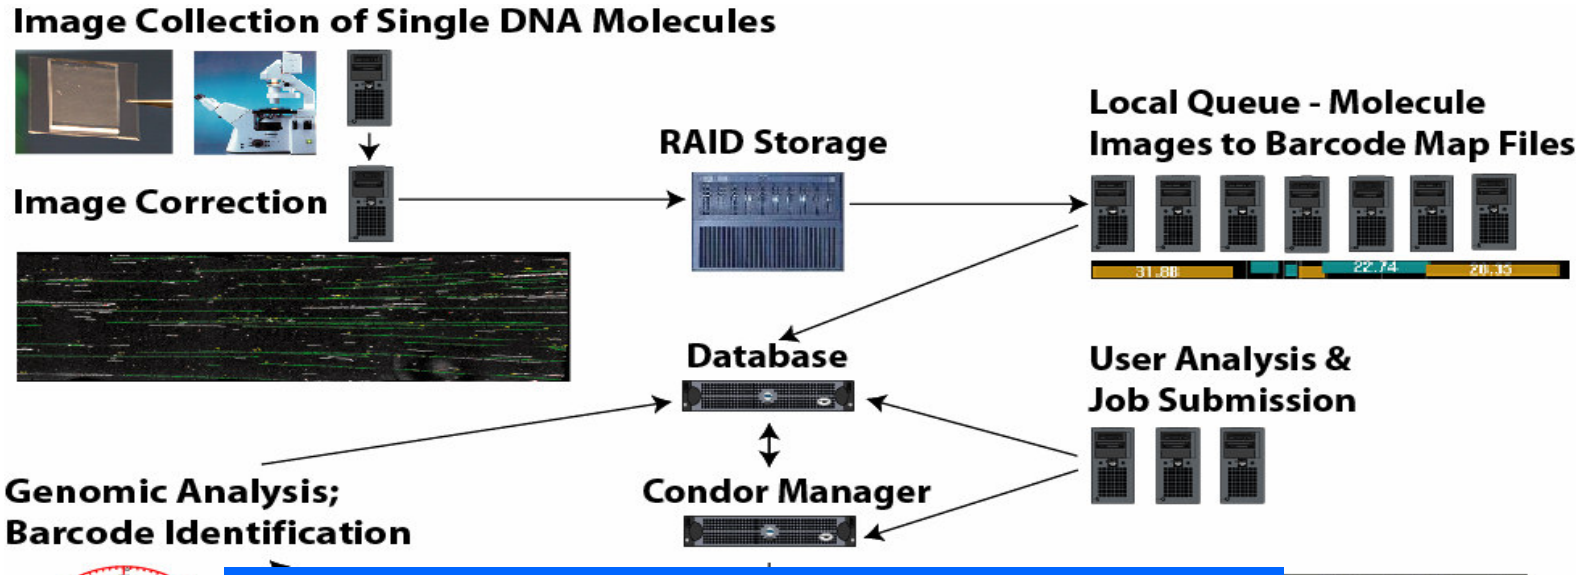


# Grid Laboratory Of Wisconsin (GLOW)

6 disciplines

~1000 CPUs

~80 TB of disk

Local GLOW CS

www.cs.wisc.edu/condor

Condor

**Session 4: Reports from the Field, Part One**

| | |
|---|---|
| Semiconductor Manufacturing (and other stuff) with Condor | Boorklin Gore, M…ogy |
| Risk Modeling with Condor at The Hartford | Bob Nord… |
| Large, Fast, and Out of Control: Tuning Condor for Film Production | Jaso… …ure |
| Optena: Enterprise Condor | …Optena |
| Introduction to gridMatrix and Condor | …ipineni, Cadence Design …ms |

**Session 5: Reports from the Field, Part Two**

| | |
|---|---|
| The Use of Condor in the gL… | Erwin Laure, EGEE |
| CMS Data Grid, Ope… …ndor-C | Ian Fisk, Fermi National Laboratory |
| Condor Usage… …al Lab | Brookhaven National Laboratory |
| Data r… …n the SAM-Grid | Gabriele Garzoglio, Fermi National Laboratory |
| …ge Scale Data Analysis within the …llaboration | Duncan Brown, LIGO |
| … for On-line Data Analysis within the …ientific Collaboration | Kipp Cannon, LIGO |

# Powerful

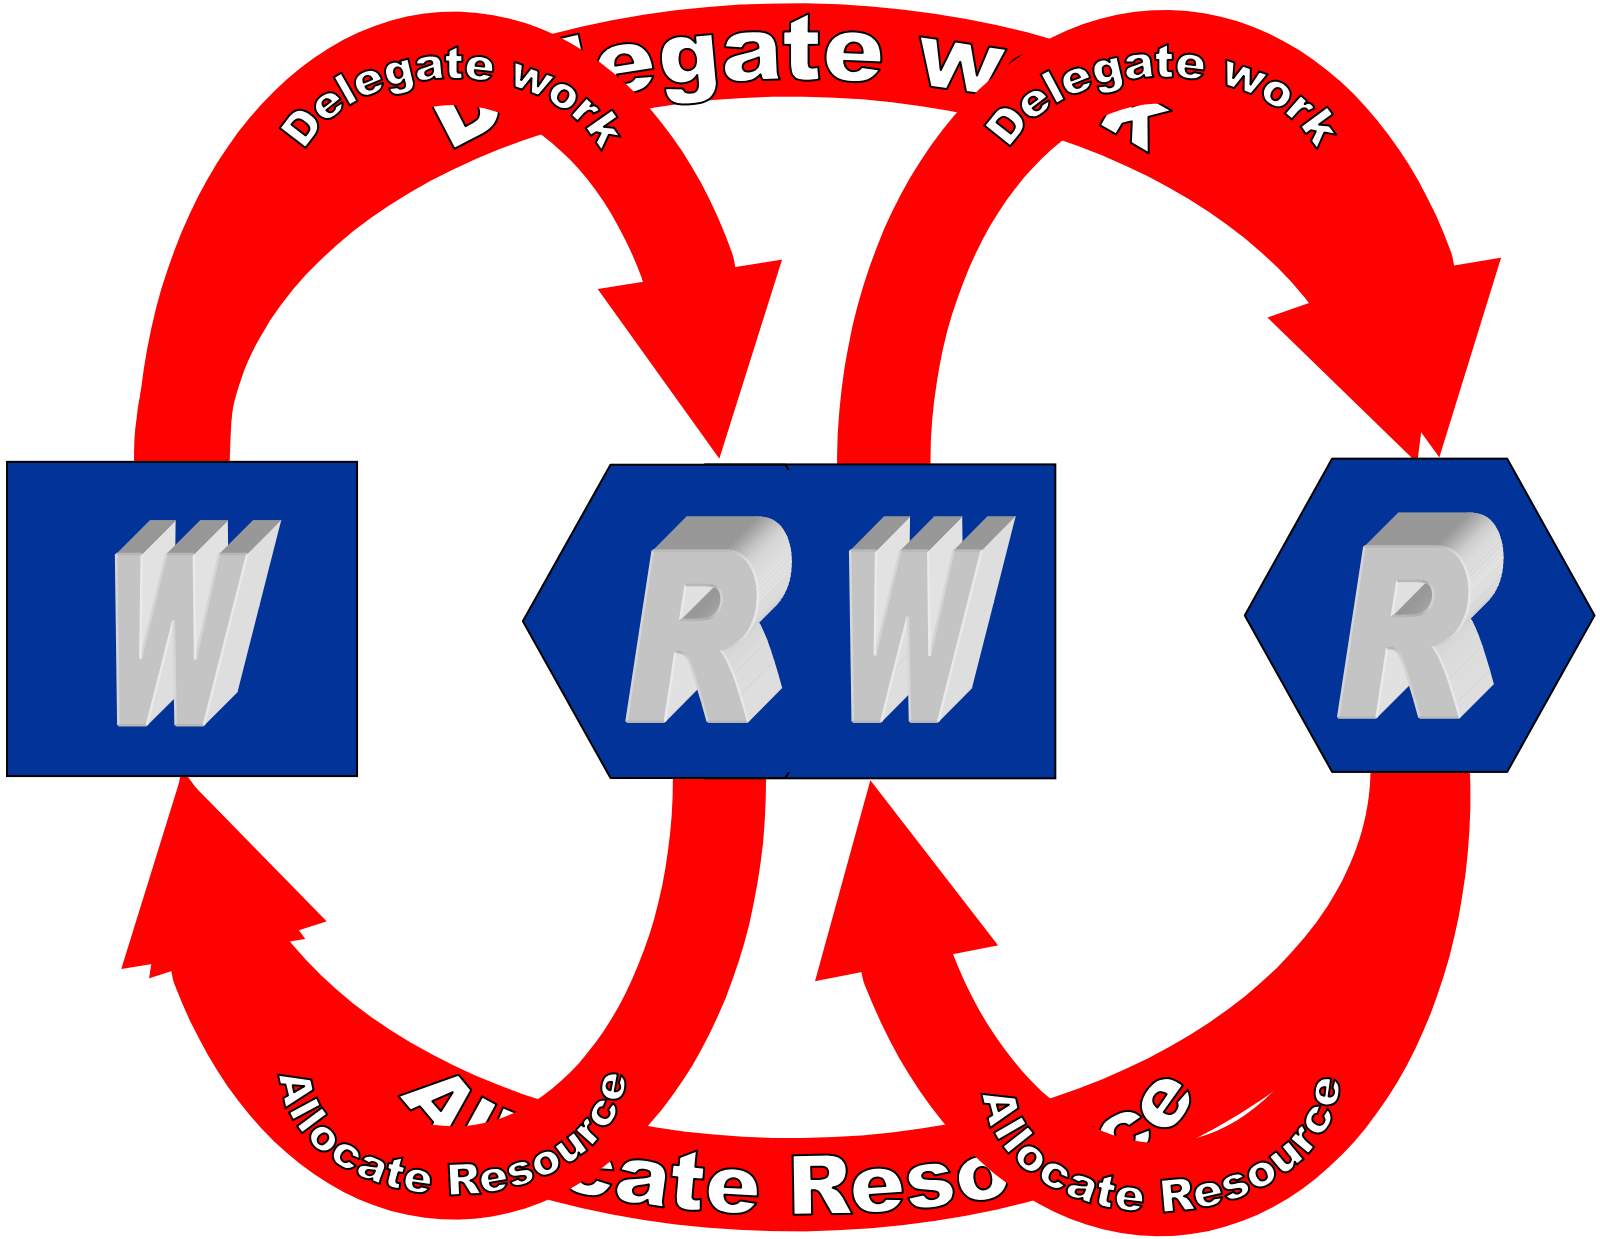
Delegate work
Delegate work
Allocate Resource
Allocate Resource
W
R W
R

# Resource Allocation

A limited assignment of the "ownership" of a resource

- Owner is charged for allocation regardless of actual consumption
- Owner can allocate resource to others
- Owner has the right and means to revoke an allocation
- Allocation is governed by an "agreement" between the client and the owner
- Allocation is a "lease"
- Tree of allocations



"We present some principles that we believe should apply in any compute resource management system. The first, P1, speaks to the need to avoid "resource leaks" of all kinds, as might result, for example, from a monitoring system that consumes a nontrivial number of resources.

P1 - It must be possible to monitor and control *all* resources consumed by a CE—whether for "computation" or "management."

Our second principle is a corollary of P1:

P2 - A system should incorporate circuit breakers to protect both the compute resource and clients. For example, negotiating with a CE consumes resources. How do we prevent an eager client from turning into a denial of service attack? "

Ian Foster & Miron Livny, *"Virtualization and Management of Compute Resources: Principles and Architecture ", A working document (February 2005)*

# Work Delegation

A limited assignment of the responsibility to perform the work

- Delegation involved a definition of these "responsibilities"
- Responsibilities my be further delegated
- Delegation consumes resources
- Delegation is a "lease"
- Tree of delegations

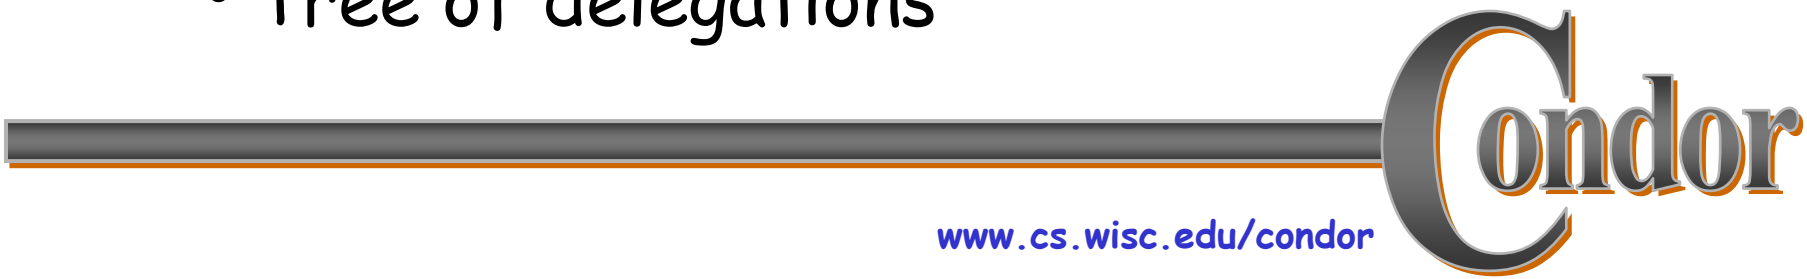


HawkEye
DAGMan
NeST
Stork
Parrot
BirdBath
MW
Condor-G
Condor-c
GCB
chirp