CDF Computing Strategy and Progress

Igor Sfiligoi
for the CDF Offline Group

International Finance Committee
October 22, 2005

Outline

- Mission
- Software and Algorithms
- Operations
- Infrastructure Development
- dCAF status report
- Future Plans
- Summary

Mission Summary

- Rapid turnaround from taking collider data to physics results and publications
- Use distributed, shared GRID computing resources to meet growing needs for MC and analysis
  - Use standard, supported tools for computing and data-handling
  - Minimize disruption in user interfaces to maximize physics productivity
- Finalize reconstruction and simulation software to allow collaboration to focus on data analysis
- Streamline operations to reduce personnel needs

CDF Physics Goals

- "1 fb^{-1} challenge for 2005-2006"
  - Present results using 1/fb of analyzed data at winter conferences, and bulk of results using (1+) fb^{-1} at summer 06 conferences
- Follow with "2 fb^{-1} challenge" for 2006-2007

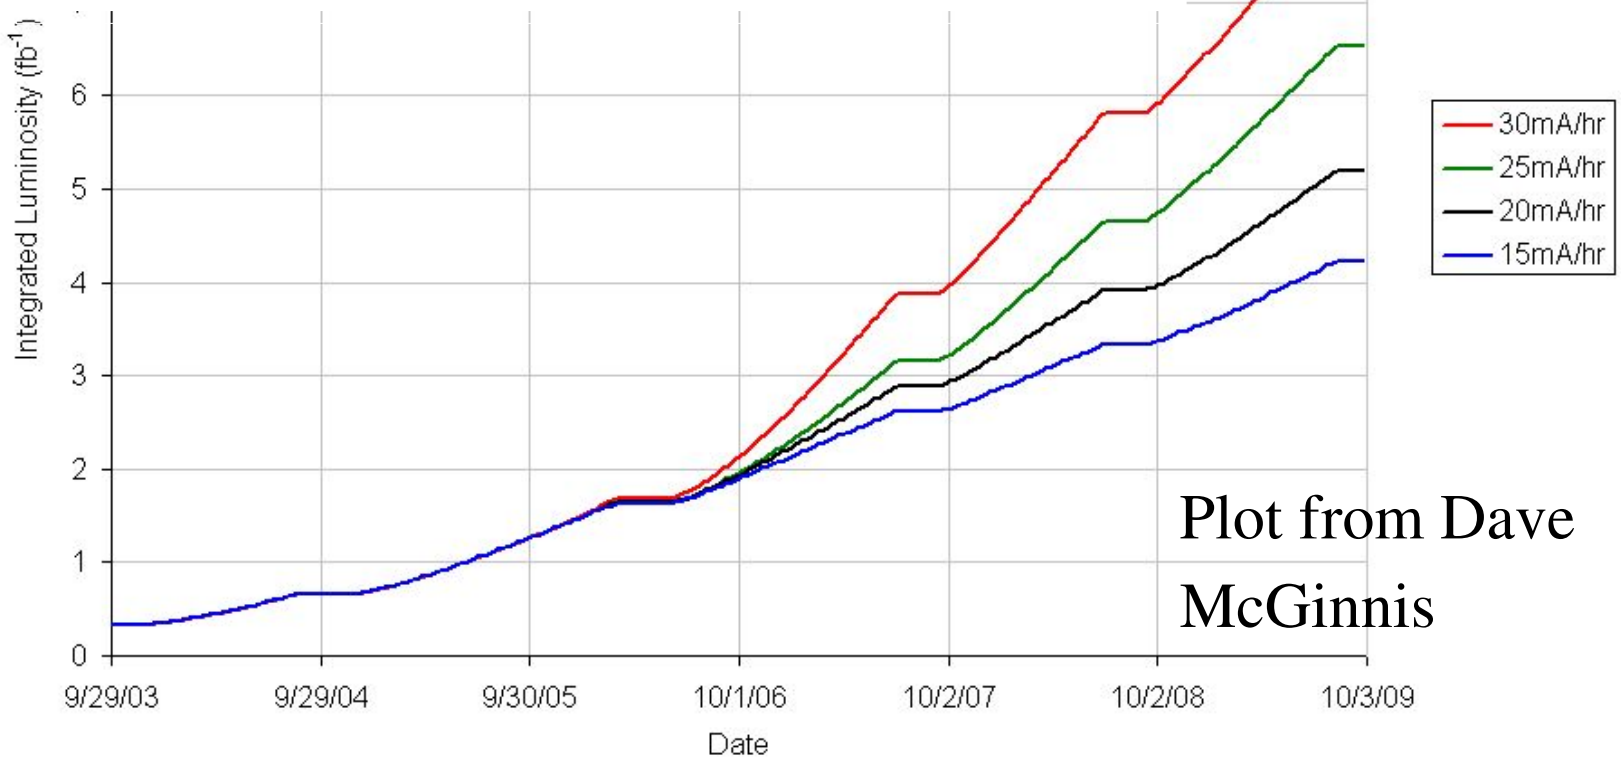

Plot from Dave McGinnis

Software Status

- Reconstruction code has achieved high level of physics performance and operational stability, and is essentially stable
  - No major changes anticipated for remaining years
- Simulation code reaching high level of maturity
  - Good agreement with data
  - Simulation machinery incorporates run-dependent
    - detector configurations
    - Multiple interactions as a function of instantaneous luminosity
  - Allows straight-forward extension of Monte Carlo datasets to match collider data
- Conclusion: software is close to being in stable, maintenance mode

Data Processing

- Goal announced last year: move away from "built-in data reprocessing" mode (process all data for ~ 1 year, finalize calibrations for physics, reprocess all the data)
- Achieved "one-pass processing mode" in 2005
  - Preprocess calibration datasets only, turn around all calibrations in 4-6 weeks, use for official dataset production
  - Successfully delivered calibrated, reconstructed data for physics analysis within 6 weeks of raw data-taking.

Period	Period Lumi	Total Lumi	Available for Physics
Dec. 04 - Mar 19, 2005	130	680	July 2005
Mar. 19 - May 20, 2005	130	810	Aug 2005
May 20 - Jul 20, 2005	100	910	Sept 2005
Jul 20 - Aug 30, 2005	50	960	Oct 2005

  - Plan to maintain this operational mode in the future

Computing Infrastructure Development

- Goal: create a uniform, GRID-enabled computing platform for all computing
  - Data processing
  - Monte Carlo production
  - User analysis
- Accomplishments in 2004-2005
  - SAM deployment (distributed data-handling)
  - New reconstruction farm deployment
  - Glide-CAF installations on LHC computing pools
  - Access to Fermi-Grid

New Reconstruction Farm

- Reconstruction farm is now integrated into a CDF-wide common computing platform
  - *"Just another CAF"*
- Successfully deployed and commissioned during January-May 2005 – Now in production mode
- Demonstrated
  - High-efficiency, high-throughput performance (25 million events / day)
  - Flexibility of reassigning CPU between analysis farm and reconstruction farm*

* Better efficiency of CPU usage

Use of Grid resources

- Successful use of “condor glide-in” technique
  - Submitting CDF jobs to a shared pool, using a dedicated head node
  - User interface identical to the other dCAFs
  - Standard Grid Gatekeeper used under the hood to get to the resources
  - Three successful installations:
    - Italian CNAF in production mode
    - FermiGrid and San Diego in final beta stage

New Direction in CDF Computing

- Existing dCAF sites moving to shared pools (away from dedicated CDF resources)
  - Italian dCAF at CNAF no longer has any CDF-specific CPU
    - All CPUs in a LCG pool
  - San Diego dCAF has all the CPUs in an OSG pool
- New dCAFs starting with OSG / LCG shared pools
  - Short-term method: condor glide-in using one headnode per site
  - New sites coming up or under negotiation
    - Paris group (Lyon center)
    - Wisconsin GRID Laboratory
    - Chicago ATLAS Tier 2 site
- Fermilab CAF will go the shared route, joining the FermiGrid

Going beyond dCAFs

- dCAFs are working fine, but we want fewer entry points
  -> We are trying to group at least some of them
  - Ongoing efforts (mainly by CDF Italian collaborators) to make the CAF interface operable with the LCG
    - Fruitful collaboration between CDF and LCG efforts in Italy
    - Goal: user jobs submitted to “LCG-CAF” would be routed by LCG Resource Broker to the whole LCG
  - Similar effort (a joint project with the UCSD CMS group) to make an “OSG-CAF”
    - Extends the GlideCAF model
    - Goal: single point of submission for the whole OSG

dCAF status report

Offsite CPU resources

All offsite resources

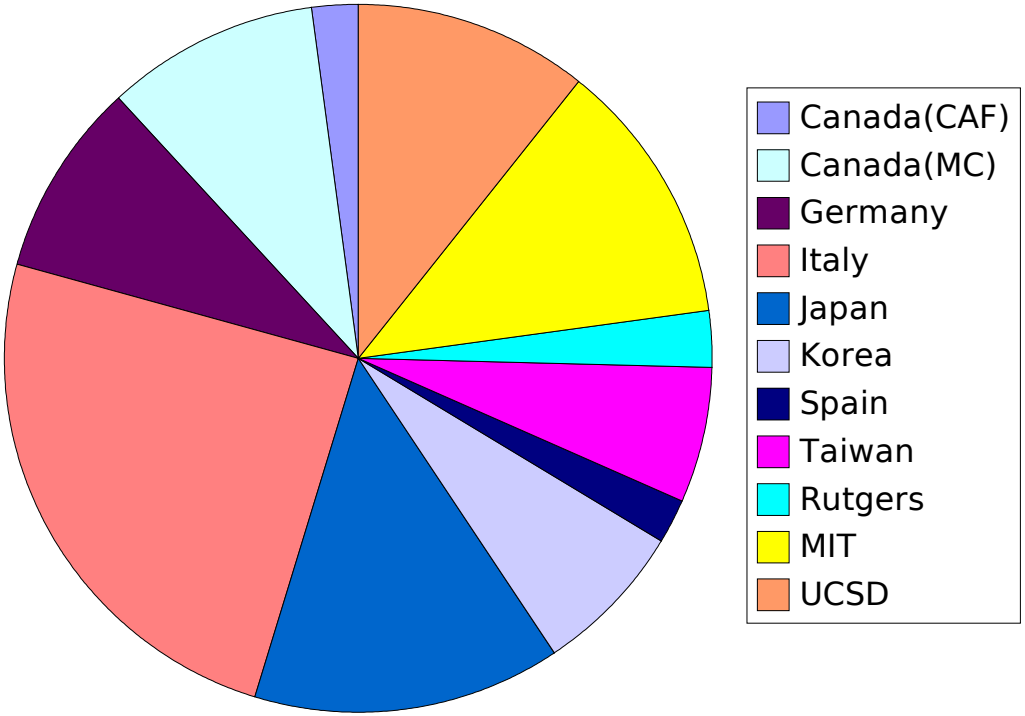

2.5 THz

dCAFs only

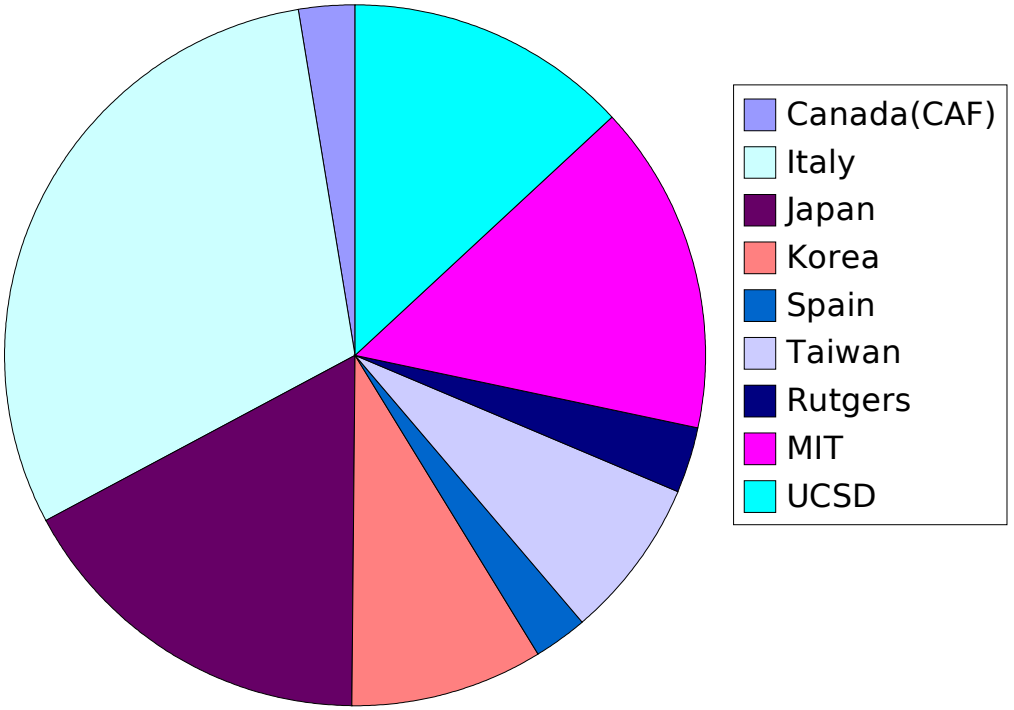

2.0 THz

CPU resources well used

- At Fermilab all CPUs always in use, could use more

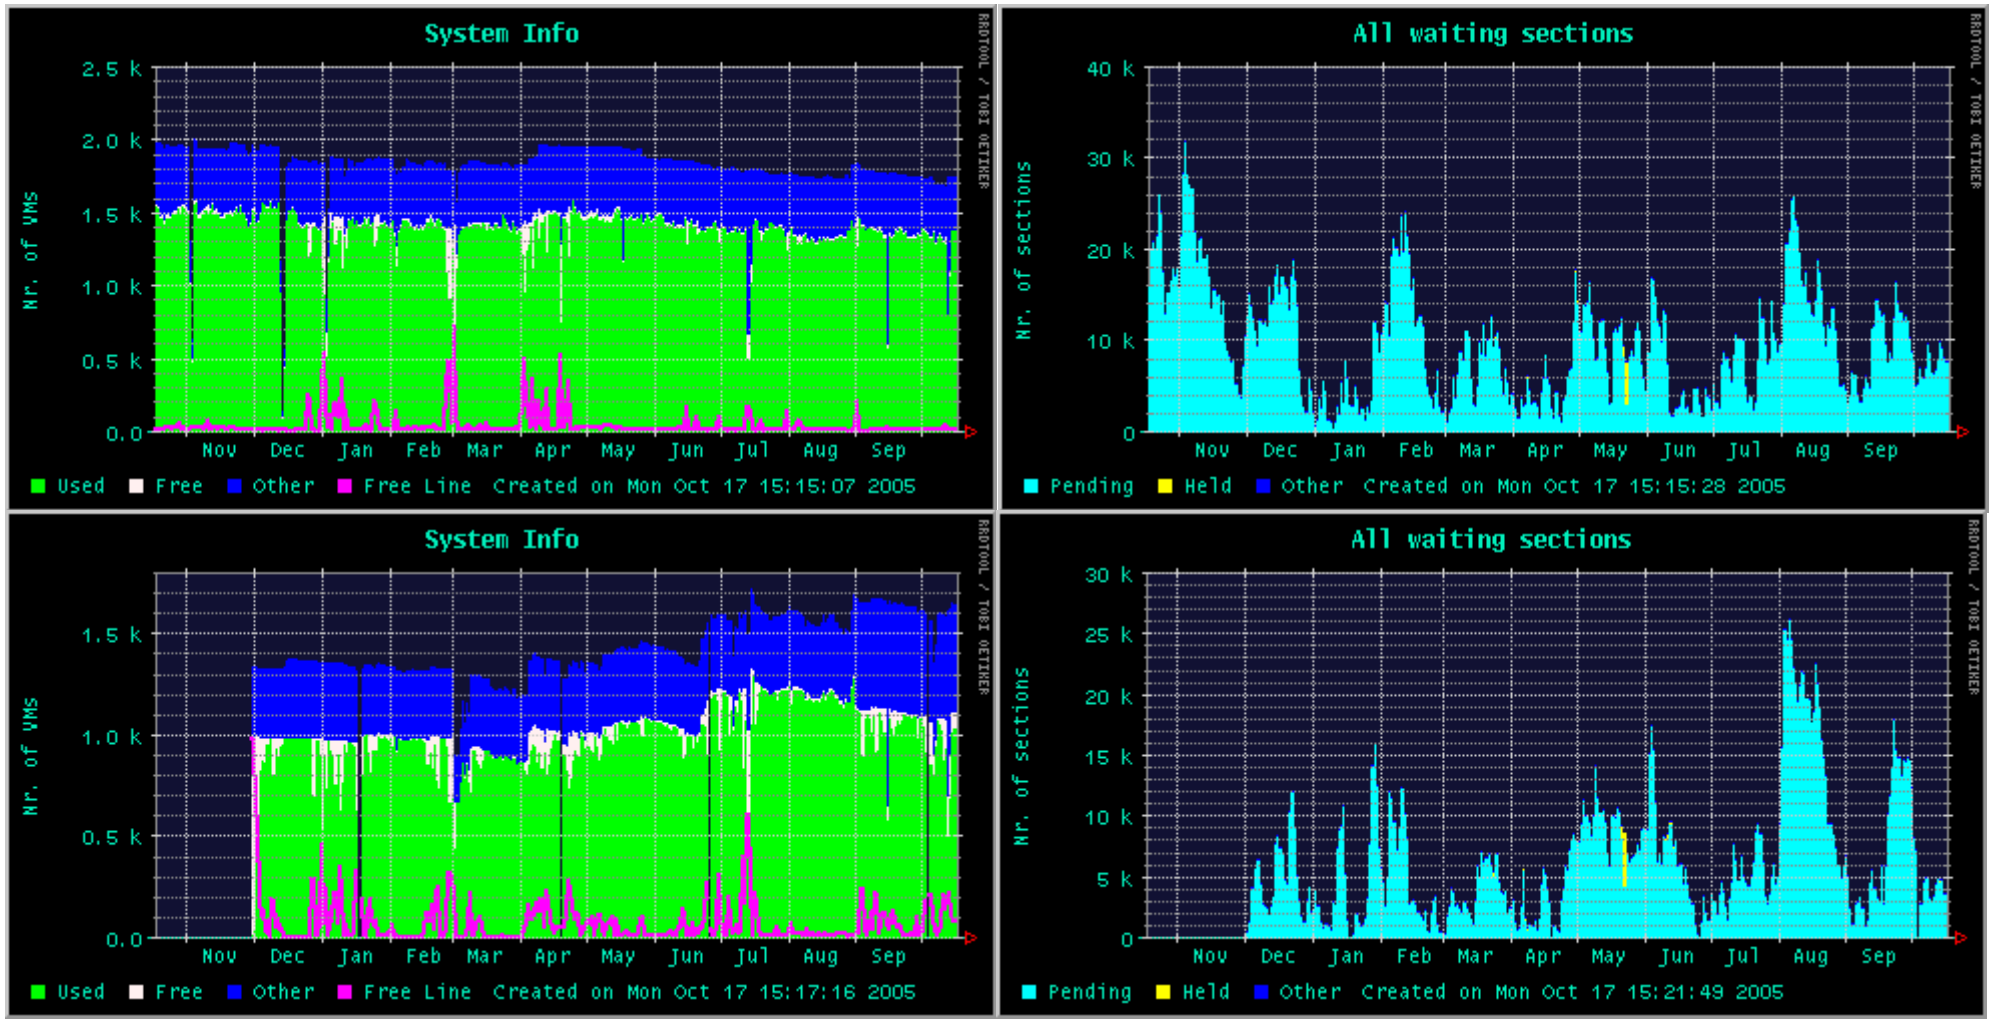

CPU resources well used(2)

- Offsite dCAFs fully loaded most of the time

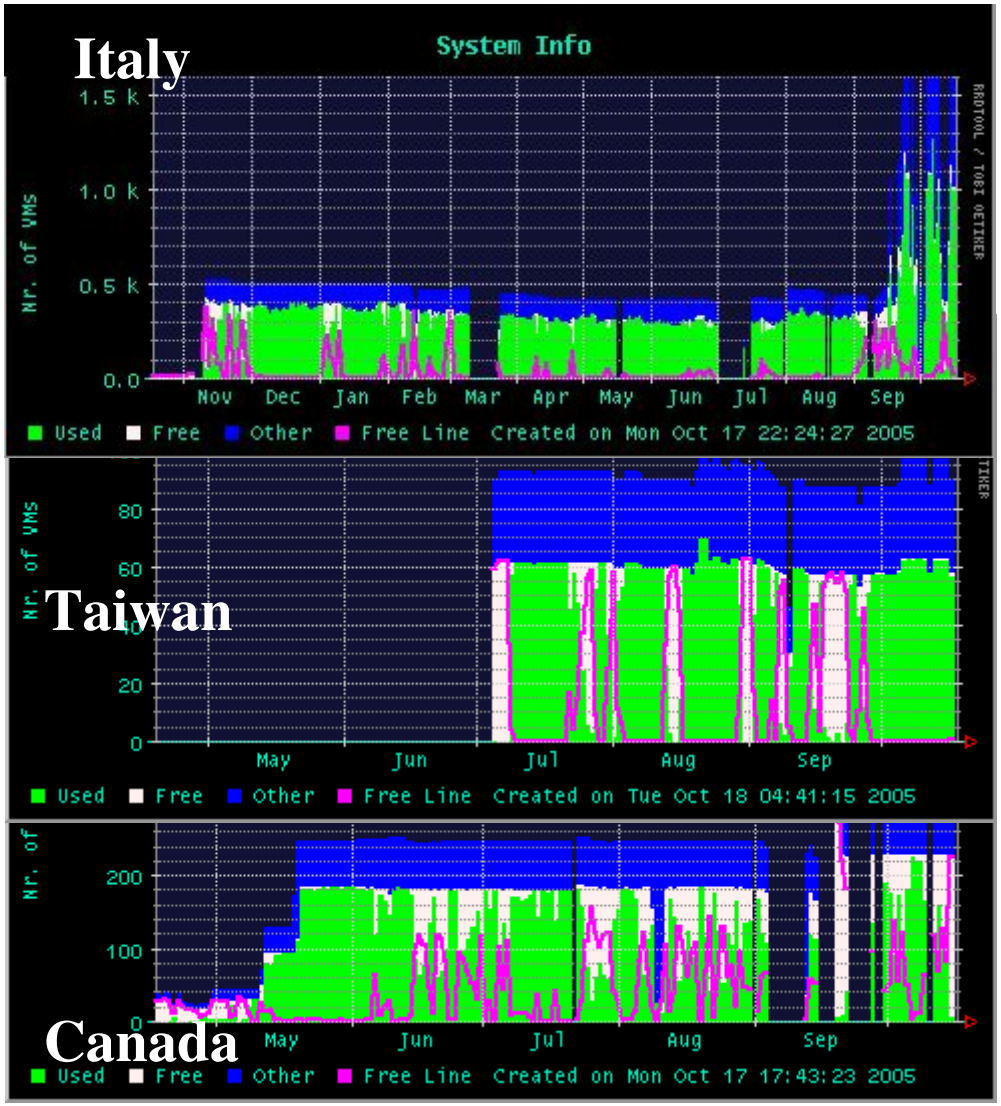

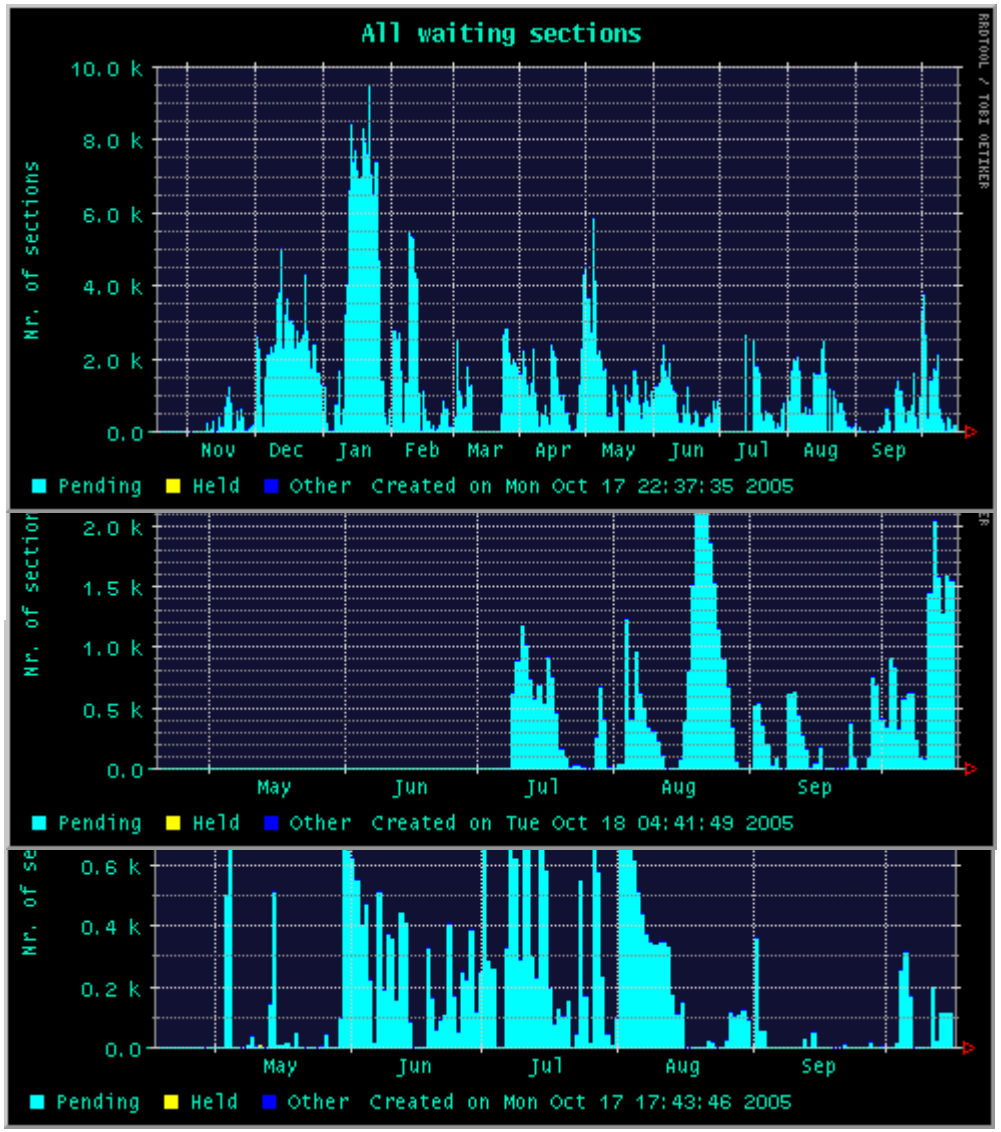

CPU resources well used(2b)

- Offsite dCAFs fully loaded most of the time

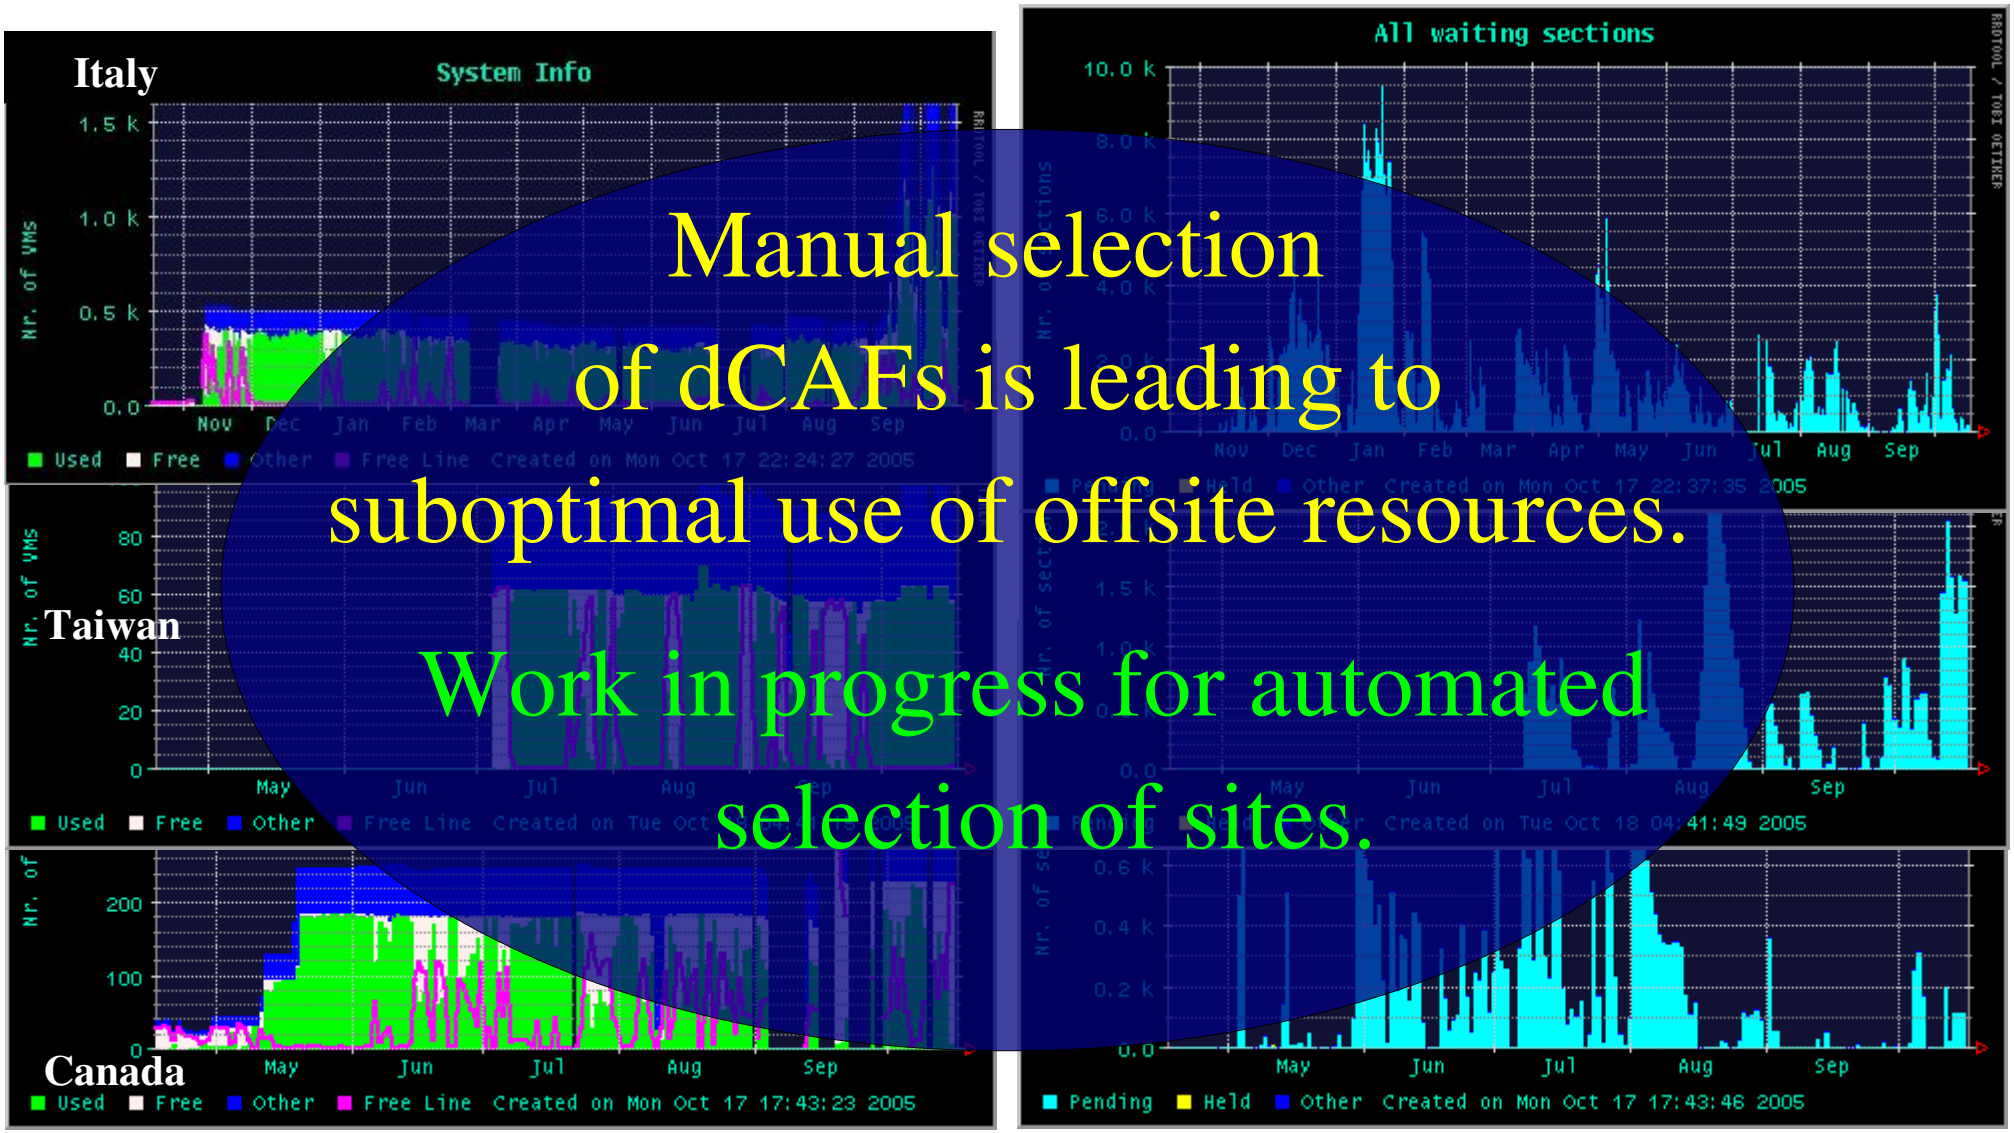

Manual selection of dCAFs is leading to suboptimal use of offsite resources.

Work in progress for automated selection of sites.

dCAF CPU usage by type

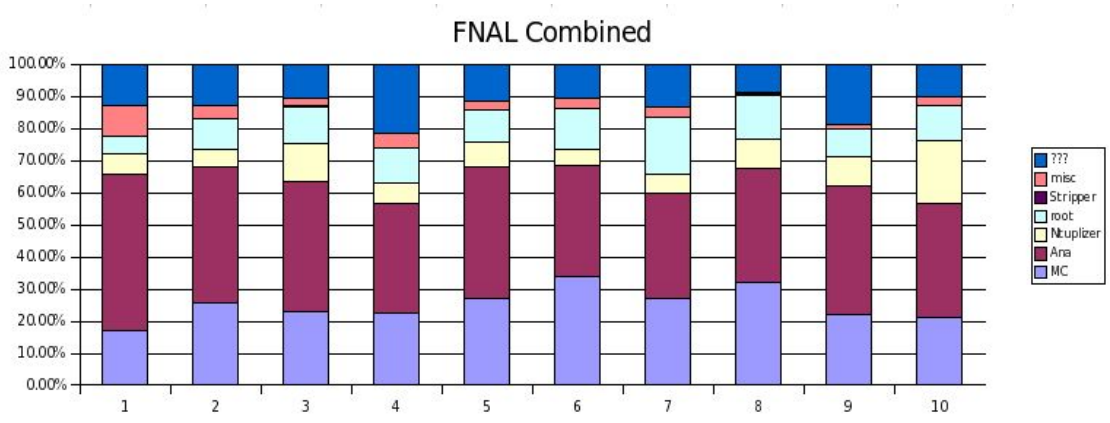

- Fermilab mostly analysis oriented

- Offsite more MC, but data-oriented jobs not negligible

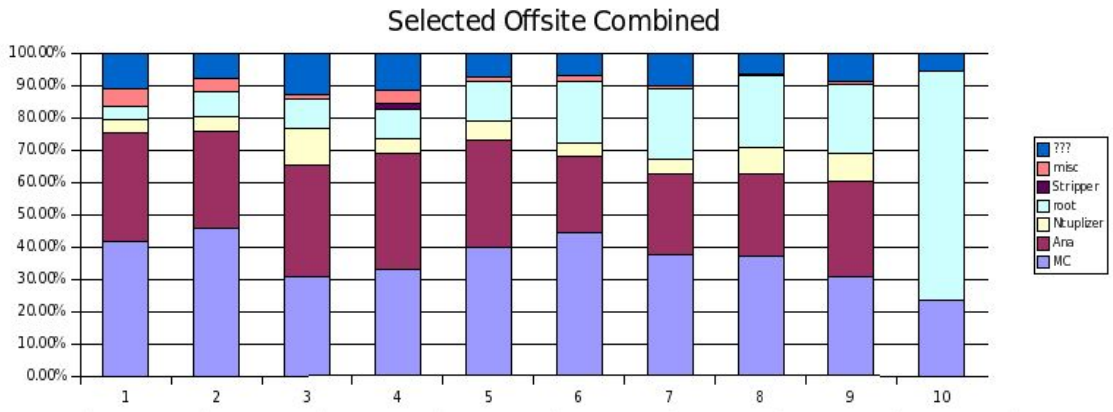

The above figures do not include MC factories

Top papers blessed

dCAF CPU usage by type(2)

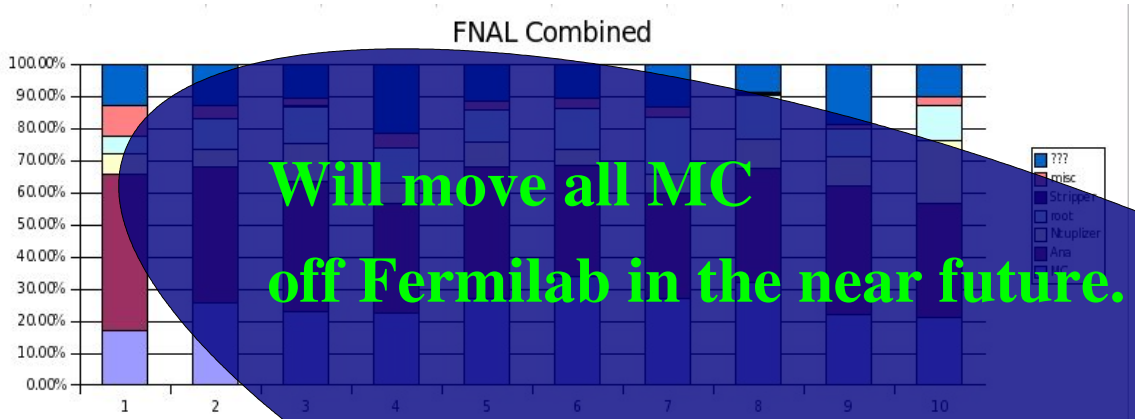

Will move all MC off Fermilab in the near future.

- Fermilab mostly analysis oriented

- Offsite more MC, but data-oriented jobs not negligible

But a good part of analysis jobs need to stay offsite!

The above figures do not include MC factories

Not all offsite dCAFs are equal

- Some dCAFs are used for MC only

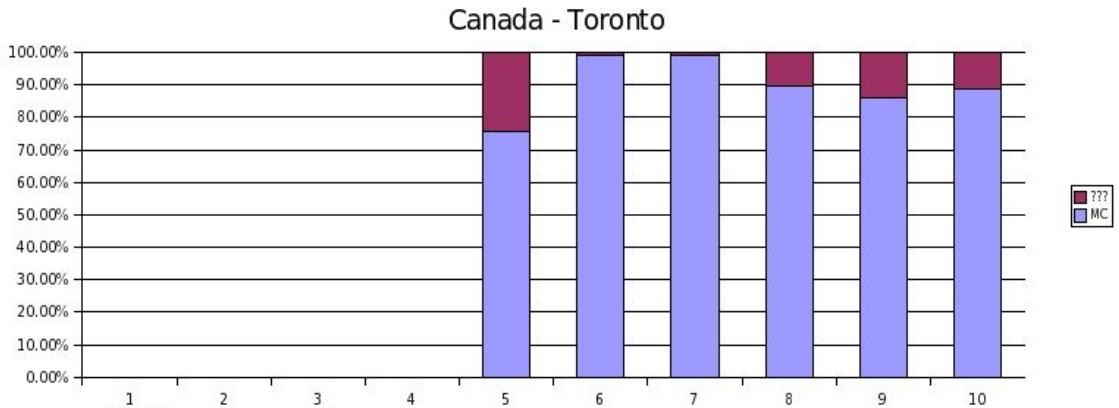

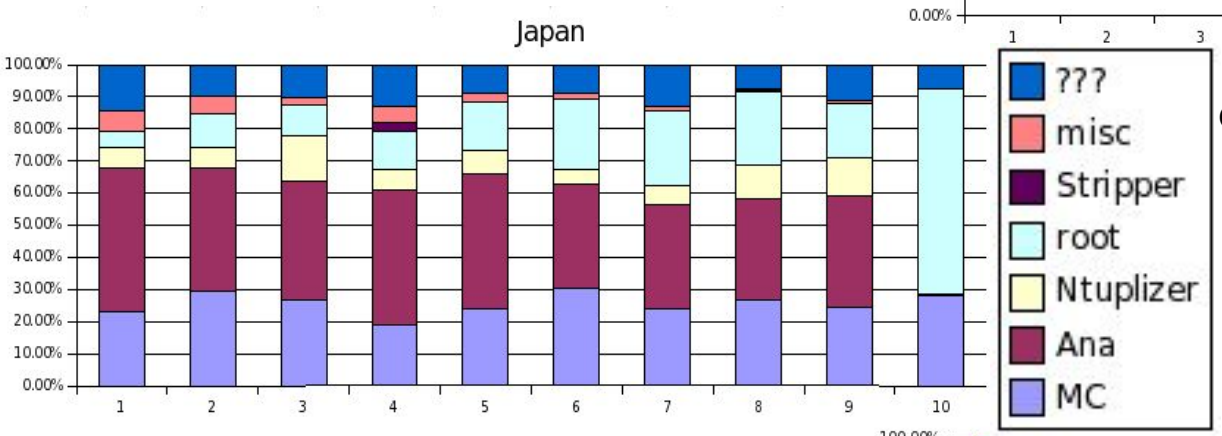

- Others are used heavily for analysis

- Some are shifting toward analysis just recently

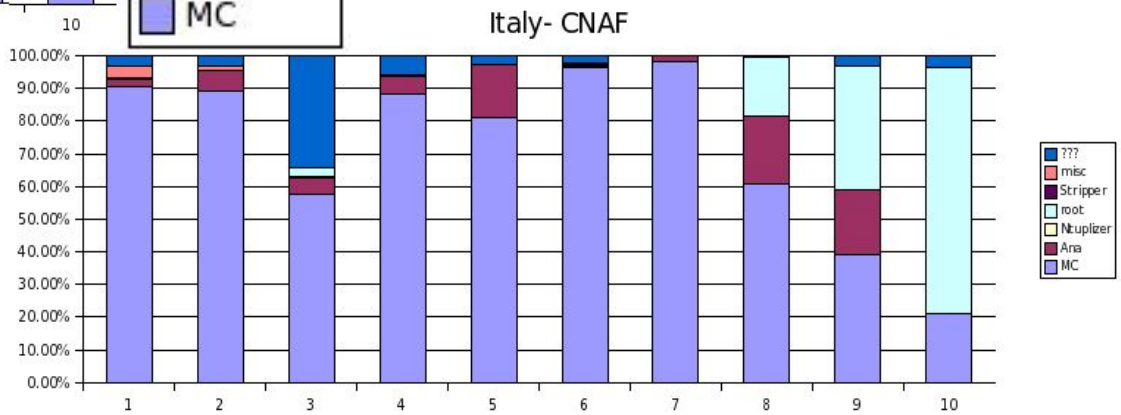

Top analysis: An example of offline use

	CPU (GHz years)	Disk (TB)
Top Mass (template l+j)	16	4.2
Top Mass (DLM, l+j)	6	0.1
Top Mass (ME, dilepton)	10	0.05
Top Mass (others dilepton)	4	1
Top XS (NN l+j)	10	0.7
Top XS (SECVTX, l+j)	8	1
Top XS (others, all)	20	4
Mttbar (ME, l+j)	10	1
	84	

CDF offsite future

- As data volume grows, so grows the CPU and disk requirements of the collaboration
  - Availability of more analysis CPU will help 1/fb challenge for winter and summer 2006
  - Both Fermilab and Offsite resources must grow
- More disk space will be needed offsite
  - With sufficient local disk cache we can target specific datasets to be served by remote dCAFs
    - Create remote physics centers
    - SAM is mature enough to make analysis as easy as if made at Fermilab
  - Similar to the LHC Tier model

CDF offsite physics centers

- A physics center must have enough CPU and disk to host all the physics analysis of those dataset(s)
- Several physics centers envisioned
  - Different in size, depending on the hosted dataset(s)
    - Disk need is approx. twice the size of the hosted dataset
    - Analysis CPU depends on the type of physics, ranges between 10GHz/TB to 30GHz/TB
    - Most would have more CPU to accommodate a reasonable amount of MC production, too
  - Affinity with local analysis needs expected (but not required)

CDF offsite physics centers(2)

- Some examples (for 1/fb, not counting MC):
  - A B physics center would need
    800GHz of CPU and 50TB of disk
  - A complete high p_T physics center would need
    200GHz of CPU and 40TB of disk
  - A jets physics center would need
    100GHz of CPU and 20TB of disk
  - A high p_T photon-only physics center would need
    50GHz of CPU and 10TB of disk

Summary

- Data processing and Monte Carlo generation schemes streamlined for fast turnaround
- Reconstruction and Simulation software mature and stable
- Distributed Computing and Data Handling working well
- Increasing access to global computing resources to match physics needs
- More analysis should move offsite
- Offsite resources are critical for the Physics Program
  - Must grow with increasing dataset size

Backup slides

Fermilab CPU usage

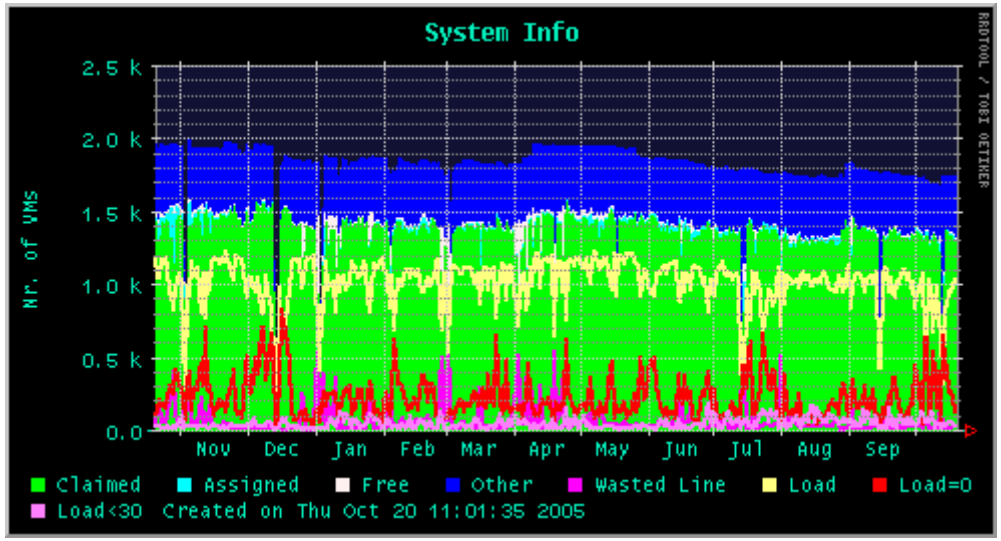

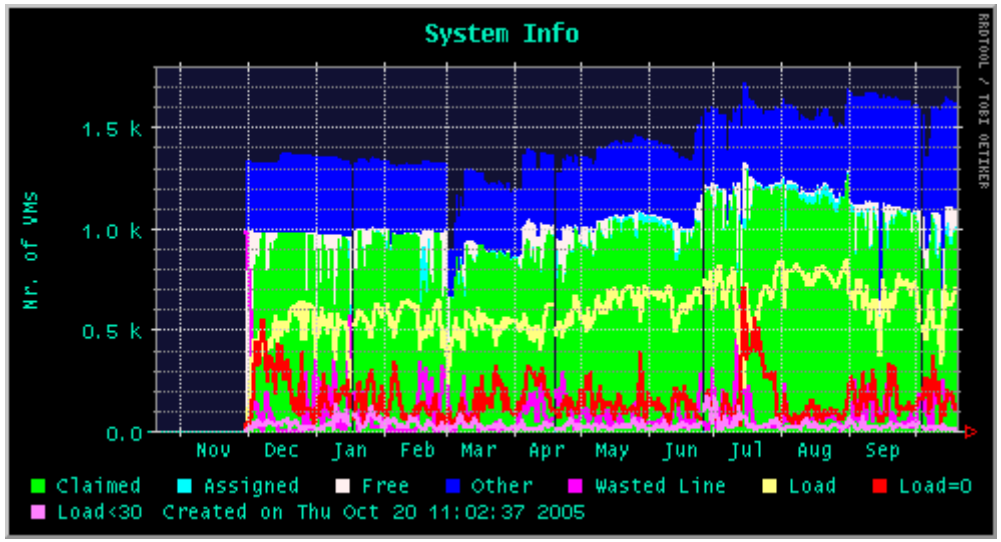

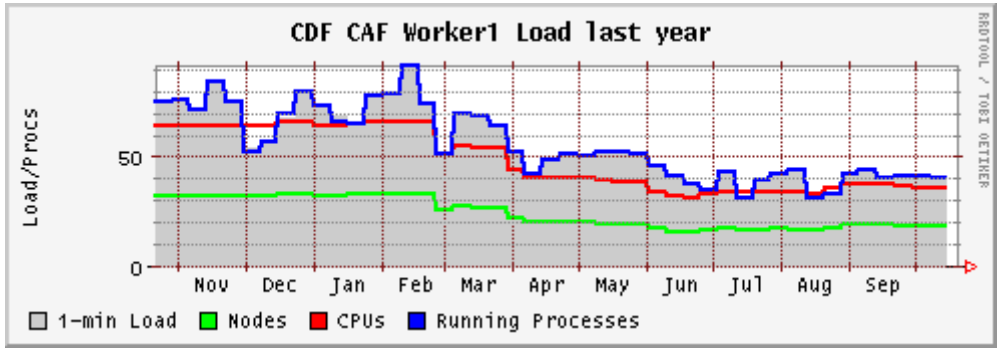

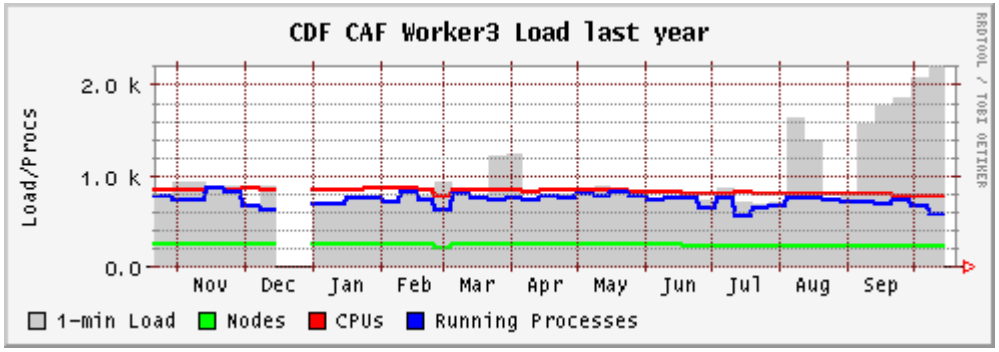

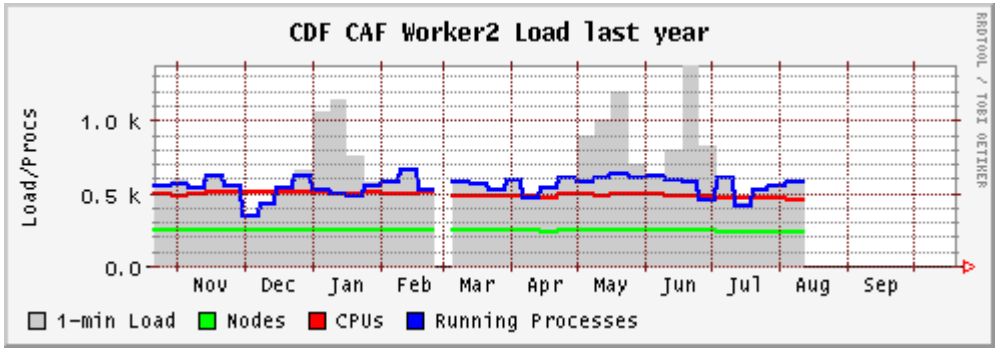

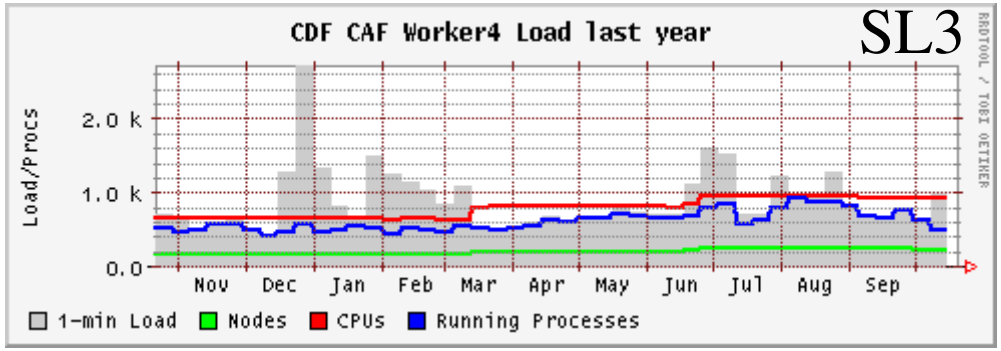
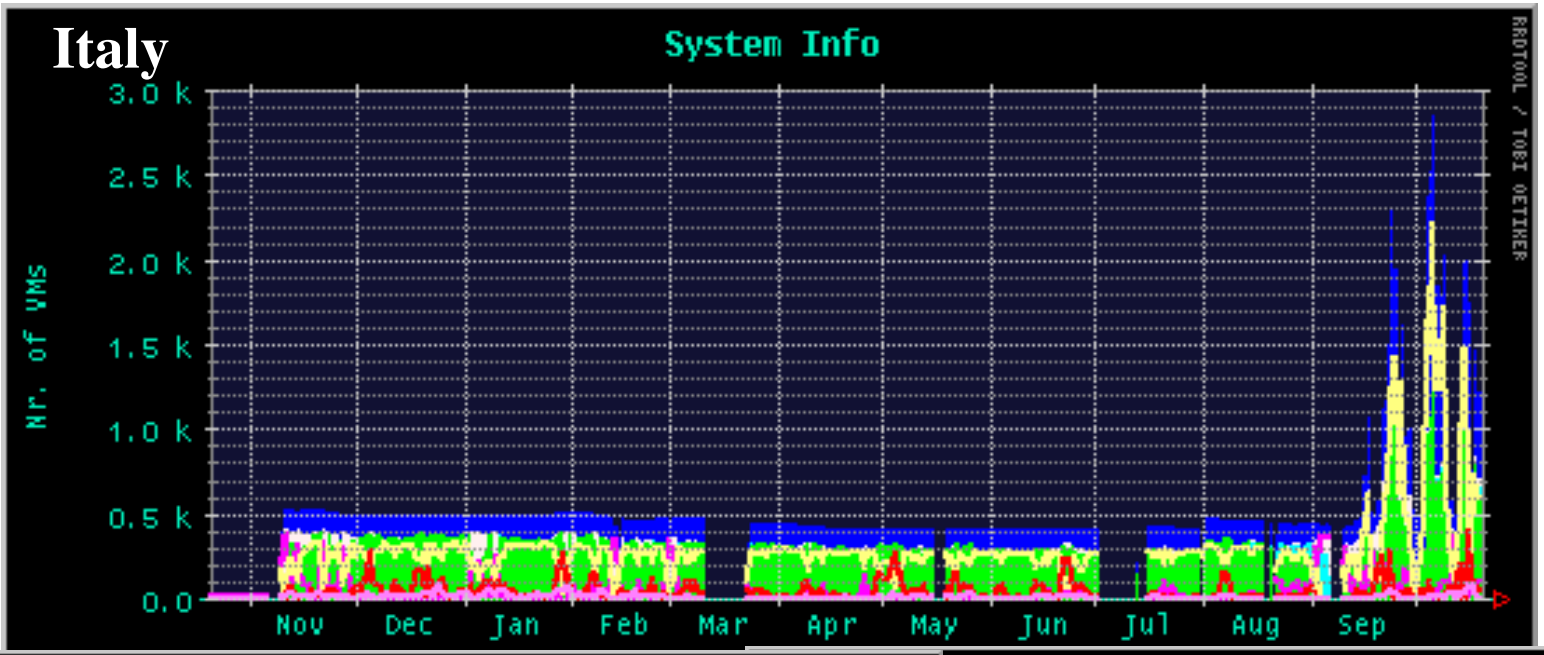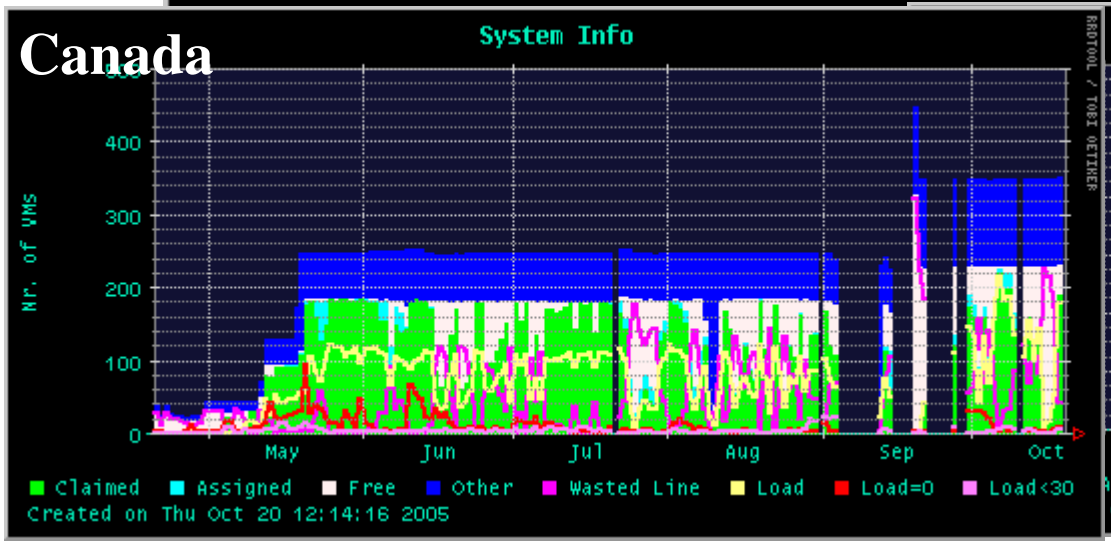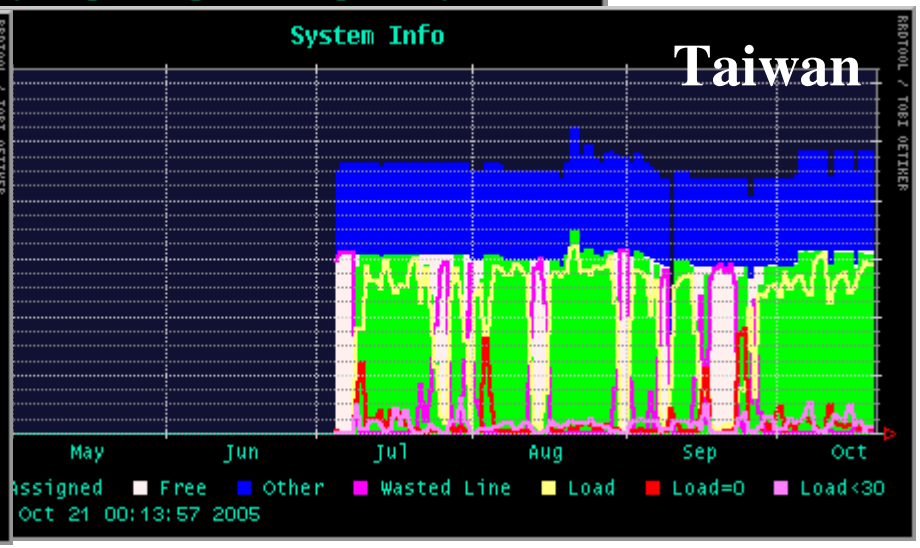
Offsite CPU load