NERSC:
Applying centers metrics

Horst Simon, Bill Kramer,
Francesca Verdier

20 February 2007

Element 1. Centers “control” metrics for performance measurement and assessment

The Panel recommends the following “control” metrics e.g. those used by PART for the centers performance:

1. User Satisfaction (overall) of provided services obtained via user surveys. *NERSC survey is recommended*
2. Scheduled Availability. Overall Availability is “observed”.
3. Response time to solve user problems as measured by the centers' trouble reporting systems.
4. Support for high capability work at LCF as per agreements; *observed and reported* distributions of jobs

User Satisfaction

NERSC has conducted an annual user survey since 1998:

- 7 point rating scale
  - 7 = very satisfied; 1= very dissatisfied
- Similar questions each year
  - Allows trend analysis and follow-up
- Each year there are 1 overall, ~10 area summary questions and ~100 detailed satisfaction questions
- NERSC analyze, reports and responds to results
- Many free form comments are collected and analyzed as well
- Current and past years surveys are at http://www.nersc.gov/news/survey/
  - Typical response is 10-15% of all NERSC users, across all areas

User Satisfaction: Overall facility score should be above 5.25

Scores above 5.25 on a 7-point scale are considered satisfactory.

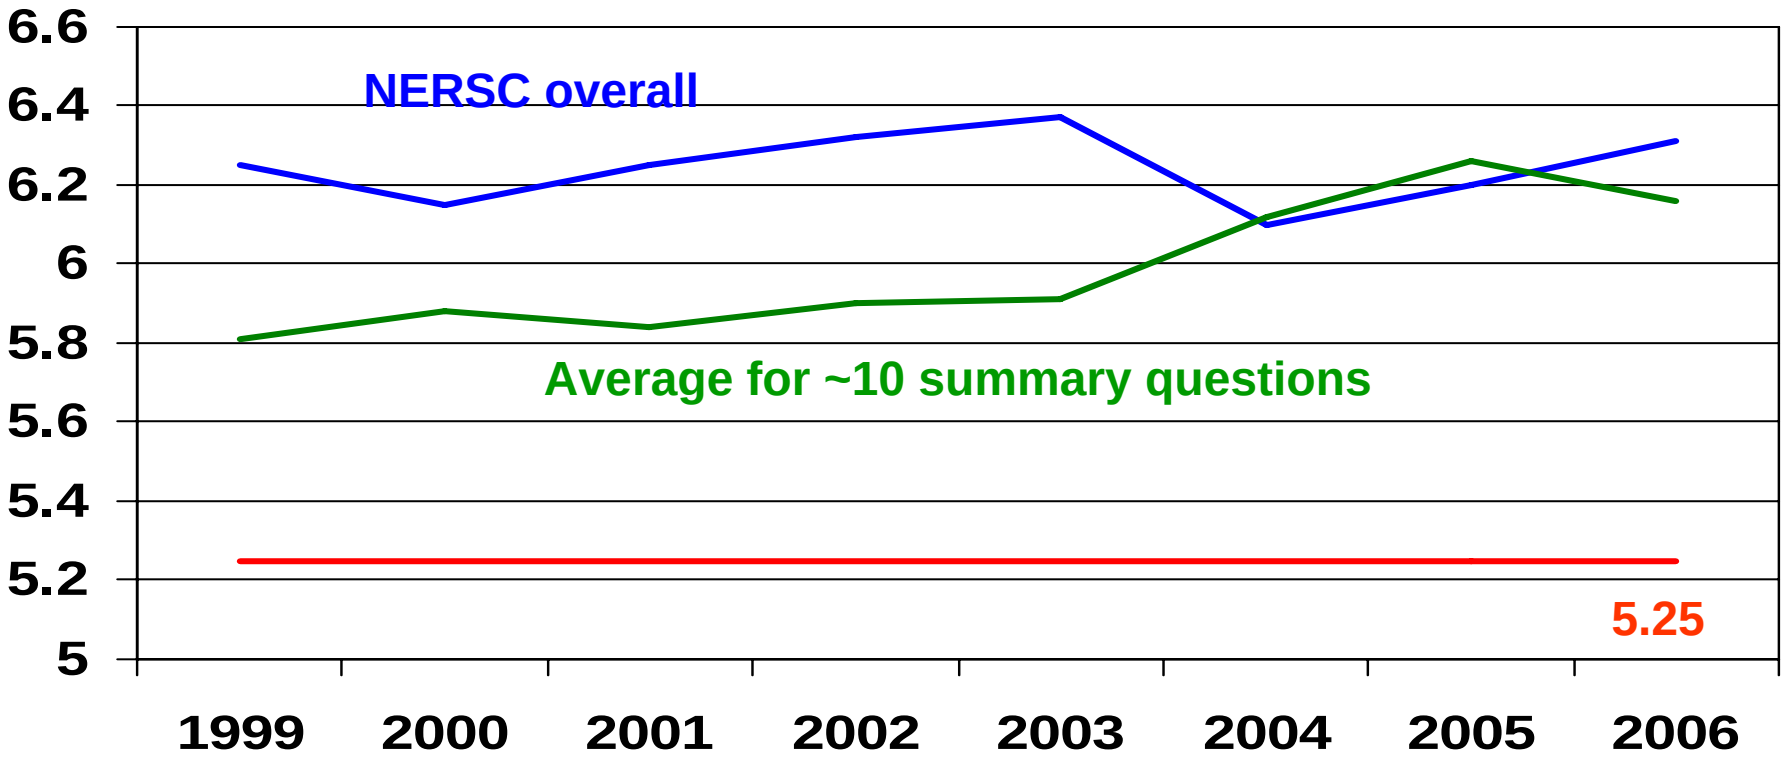

User Satisfaction
Examples of Area Summary Responses

Scores above 5.25 on a 7-point scale are considered satisfactory.

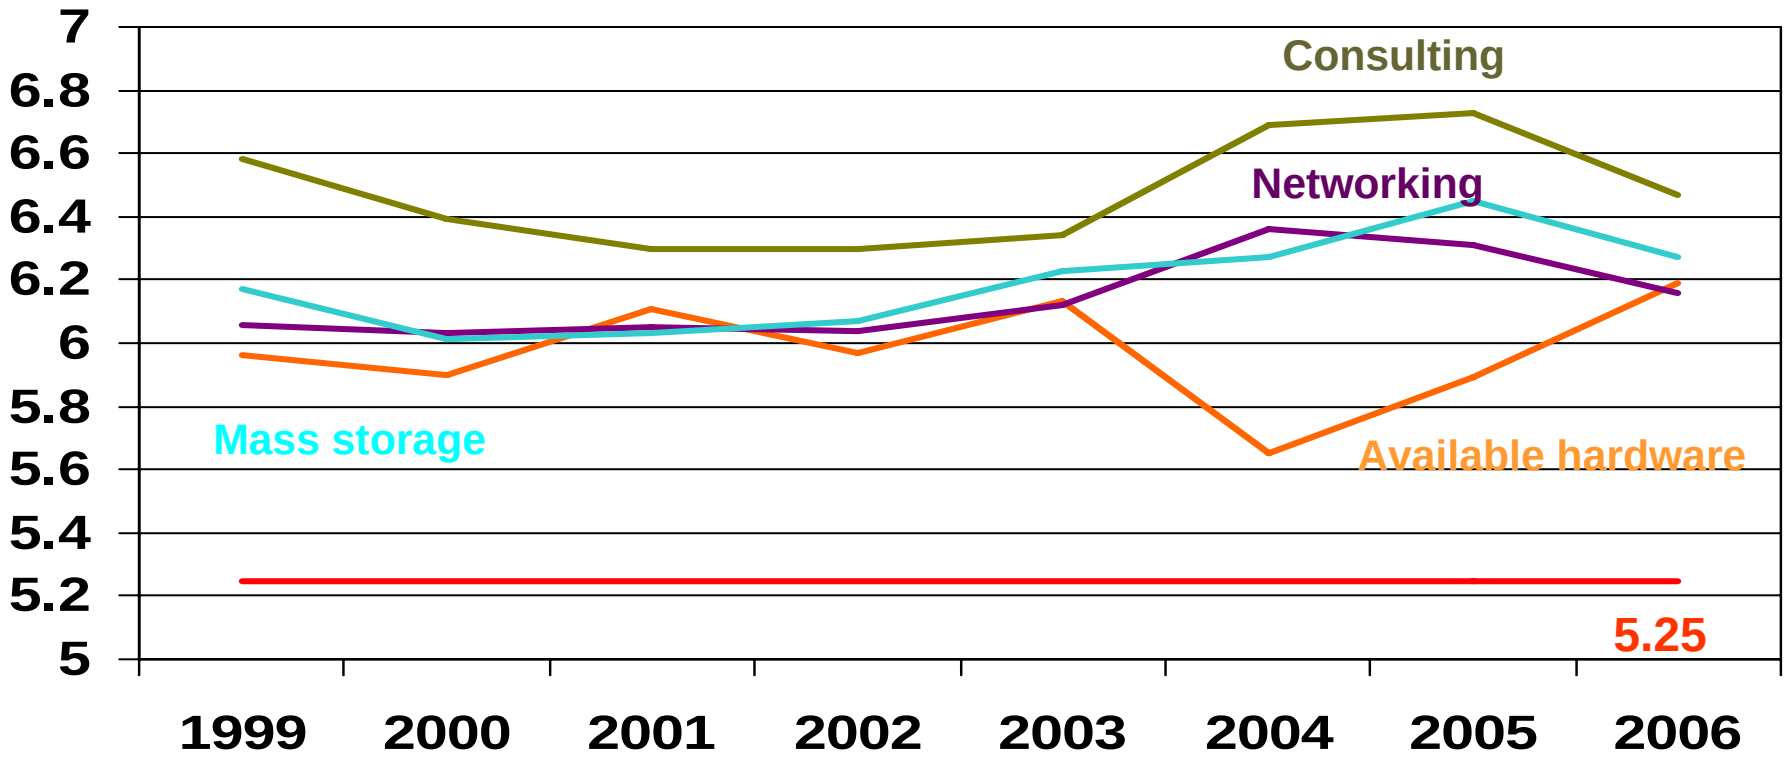

User Satisfaction: Improved Rating Next Year for Previous Low Scores

There should be an improved rating in at least half the areas for which the previous year's rating was below 5.25.

Note: for 2005 only 4 user satisfaction scores were < 5.25. By 2006, three of these had significantly improved. <4% are below 5.25

Topic	2005 score	2006 score	Significant change?	What NERSC did
Jacquard batch wait	5.16	5.87	**+0.71**	**Used Maui scheduler for better queue mgt; fixed Infiniband problems with large jobs**
PDSF disk and I/O	5.14	5.10	No change	**Switched from NFS to GPFS, with most of the improvement between 2004 and 2005.**
Seaborg queue config	5.06	5.77	**+0.72**	**Provided fairer scheduling between premium, midrange, and large jobs.**
Seaborg batch wait	3.95	4.94	**+0.99**	**Adjust duty cycle: better balance of throughput versus overall utilization**

Systems Availability/Reliability

Scheduled availability should be at least 95% for machines in their first 18 months of deployment; 97% thereafter.

System	FY04 Sched Avail	FY04 Overall Avail	FY04 MTBF d:h	FY05 Sched Avail	FY05 Overall Avail	FY05 MTBF d:h	FY06 Sched Avail	FY06 Overall Avail	FY06 MTBF d:h
Seaborg	97.57	96.68	17:15	98.53	97.27	10:18	98.71	98.30	16:02
Bassi							97.0	95.64	12:01
Jacquard							99.04	97.97	8:16
Analytics Servers	99.0	98.94	21:01	99.99	99.99	29:17	99.02	98.35	11:20
HPSS	99.47	98.90	9:17	99.27	98.09	8:12	99.86	98.21	6:04
NGF							99.37	99.24	23:15

Overall Availability is 24*7*365 and is less than or equal to scheduled availability

Compute Hours used by Science Projects Out of Theoretical Max Time

Seaborg POWER3 cycles to science users

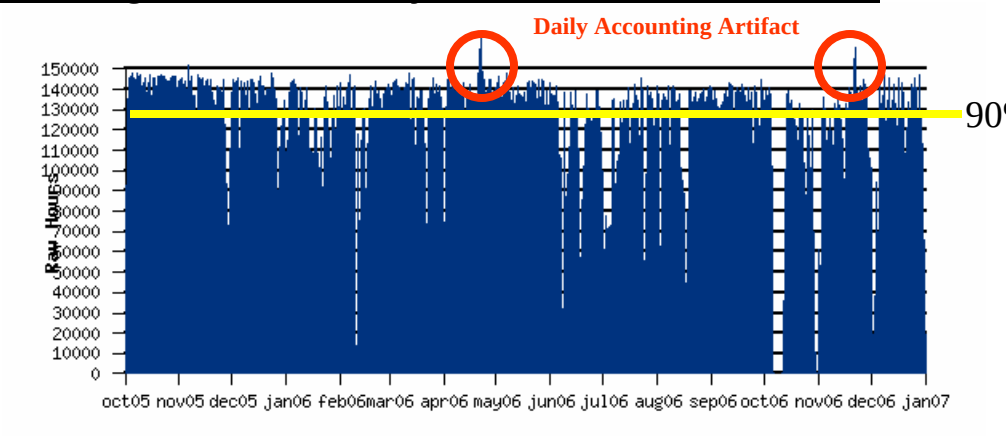

90%

Percent of overall (24*7*365*number of Compute CPU) time used by science users on Seaborg	
AY 2004	90.0%
AY 2005	93.5%
AY 2006	87.5%

Bassi POWER5 cycles to science users

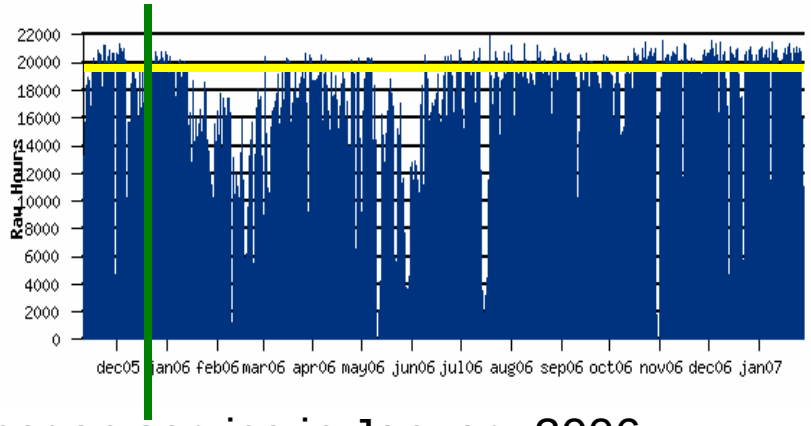

90%

Note: Bassi began service in January 2006

Response Time for Assistance

80% of user problems are addressed within 3 working days, either by resolving them or providing the user with a plan for resolution

We show days to closure which often is significantly great than days to a plan for resolution which we do not currently measure.

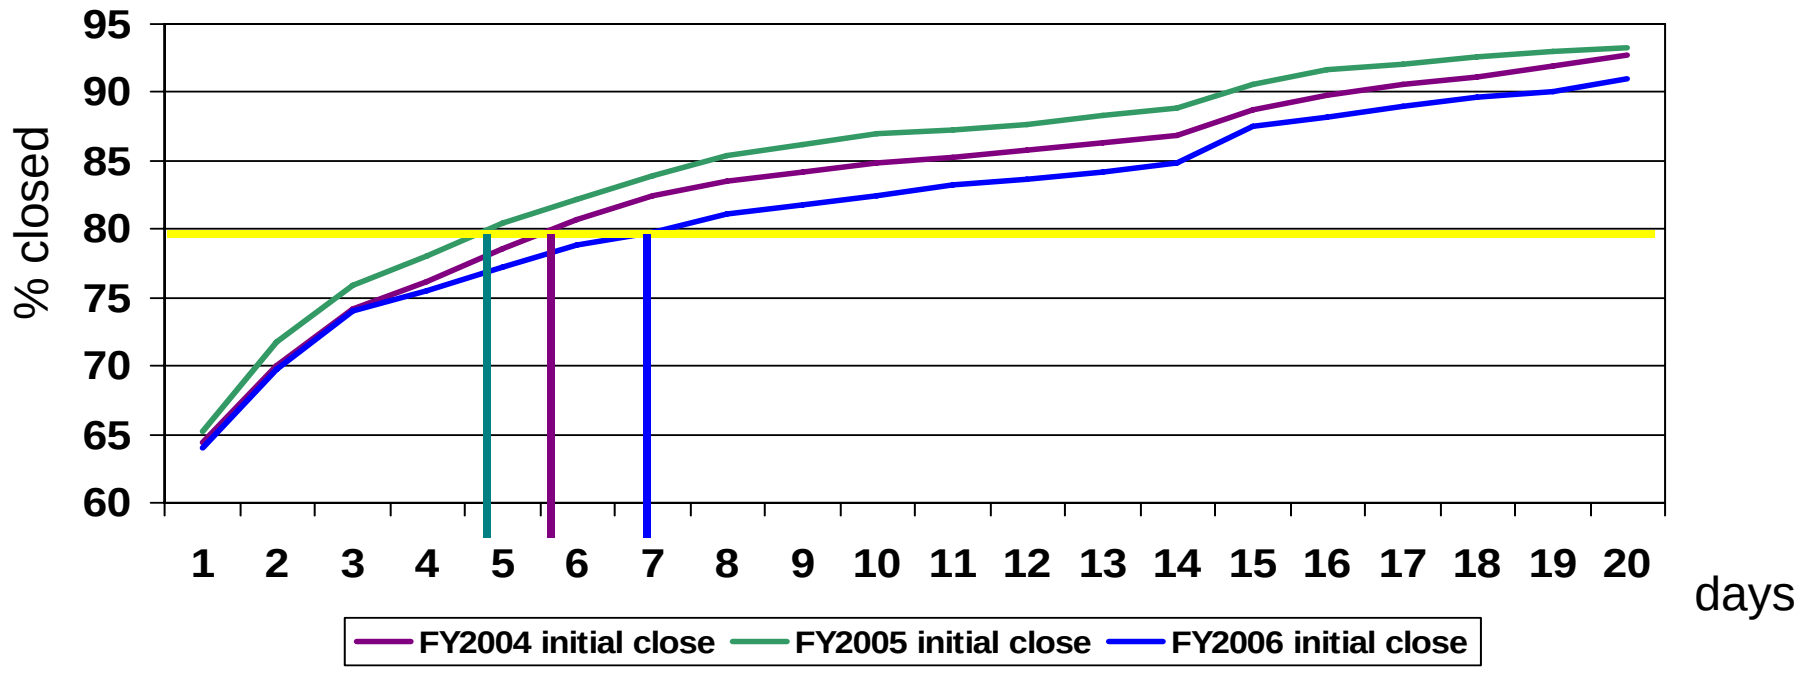

Applying the Metrics to NERSC – Priority Service to Capability Users

On leadership platforms at least 40% of the cycles should be used by jobs running on 1/8th or more of the processors.

The graph show the percent of Seaborg cycles run on 1/8th or more of the processors. Half of these "big cycles" were provided by the DOE allocation; half by incentive programs.

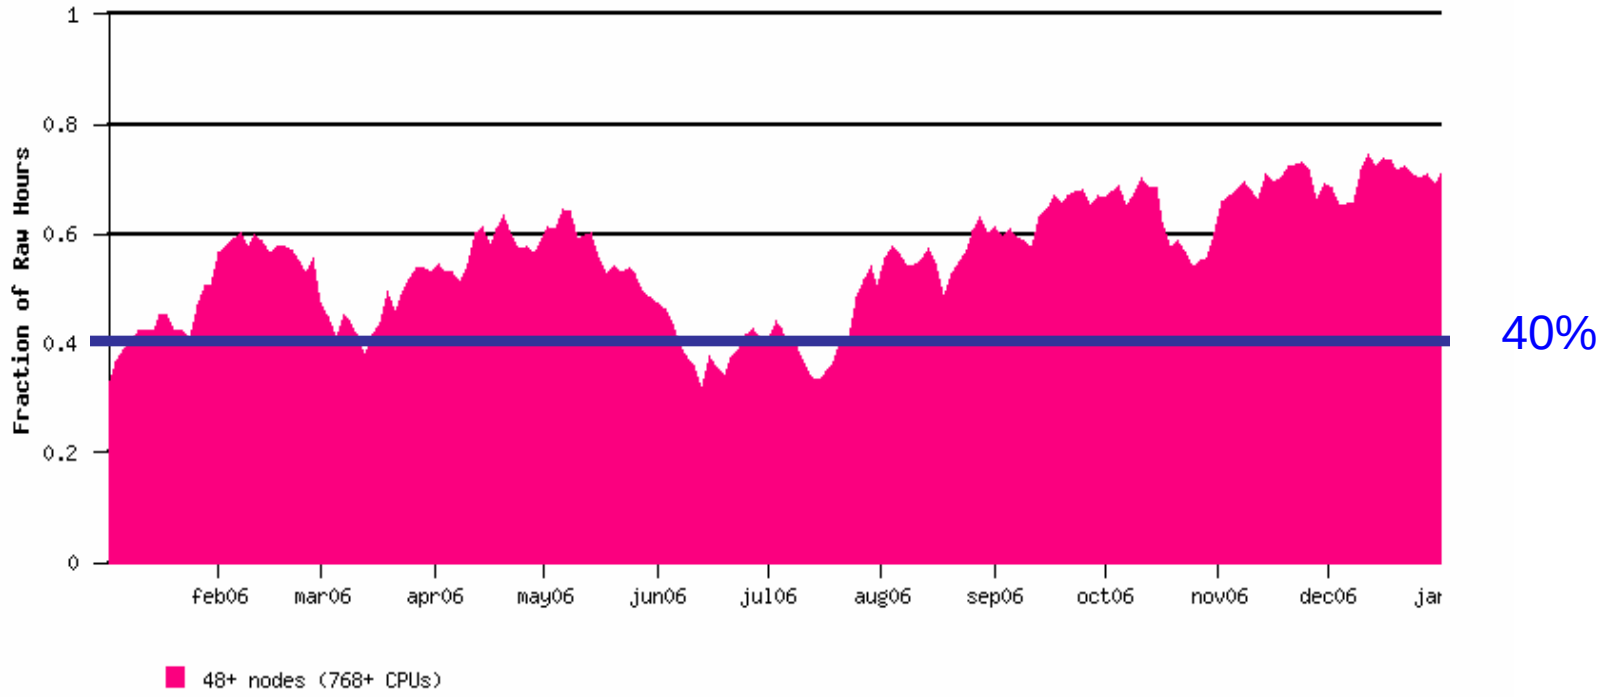

Priority Service to Capability Users

On leadership platforms the queue expansion factor for capability jobs should be monitored.

The table below shows the expansion factor (EF) for Seaborg's regular priority jobs.

	Allocation pressure	512-1008 procs EF	1024 – 2032 procs EF	1048+ procs EF
FY04	INCITE pressure	4.67	7.21	6.98
FY05	Significant over-allocated	6.51	8.84	13.68
FY06	Average	3.58	4.13	4.96

Many factors must be considered when computing the EF. Seaborg's EF:

- FY05/FY06 uses (wait + request / request)
- FY04 uses (wait + run time / run time) – request time data was not preserved in FY 04
- Individual job EF weighted by total hours used by that job
- Wait times are based on max of 6 jobs queued per user
    - Users can submit any number of jobs at any time
- Batch systems may not capture needed data and other subtle
- Some wait times are inflated due to having been put on user hold

4 Year Seaborg Queue Wait Statistics

(expansion factors depend on many conditions and might not be directly related to proposed metric)

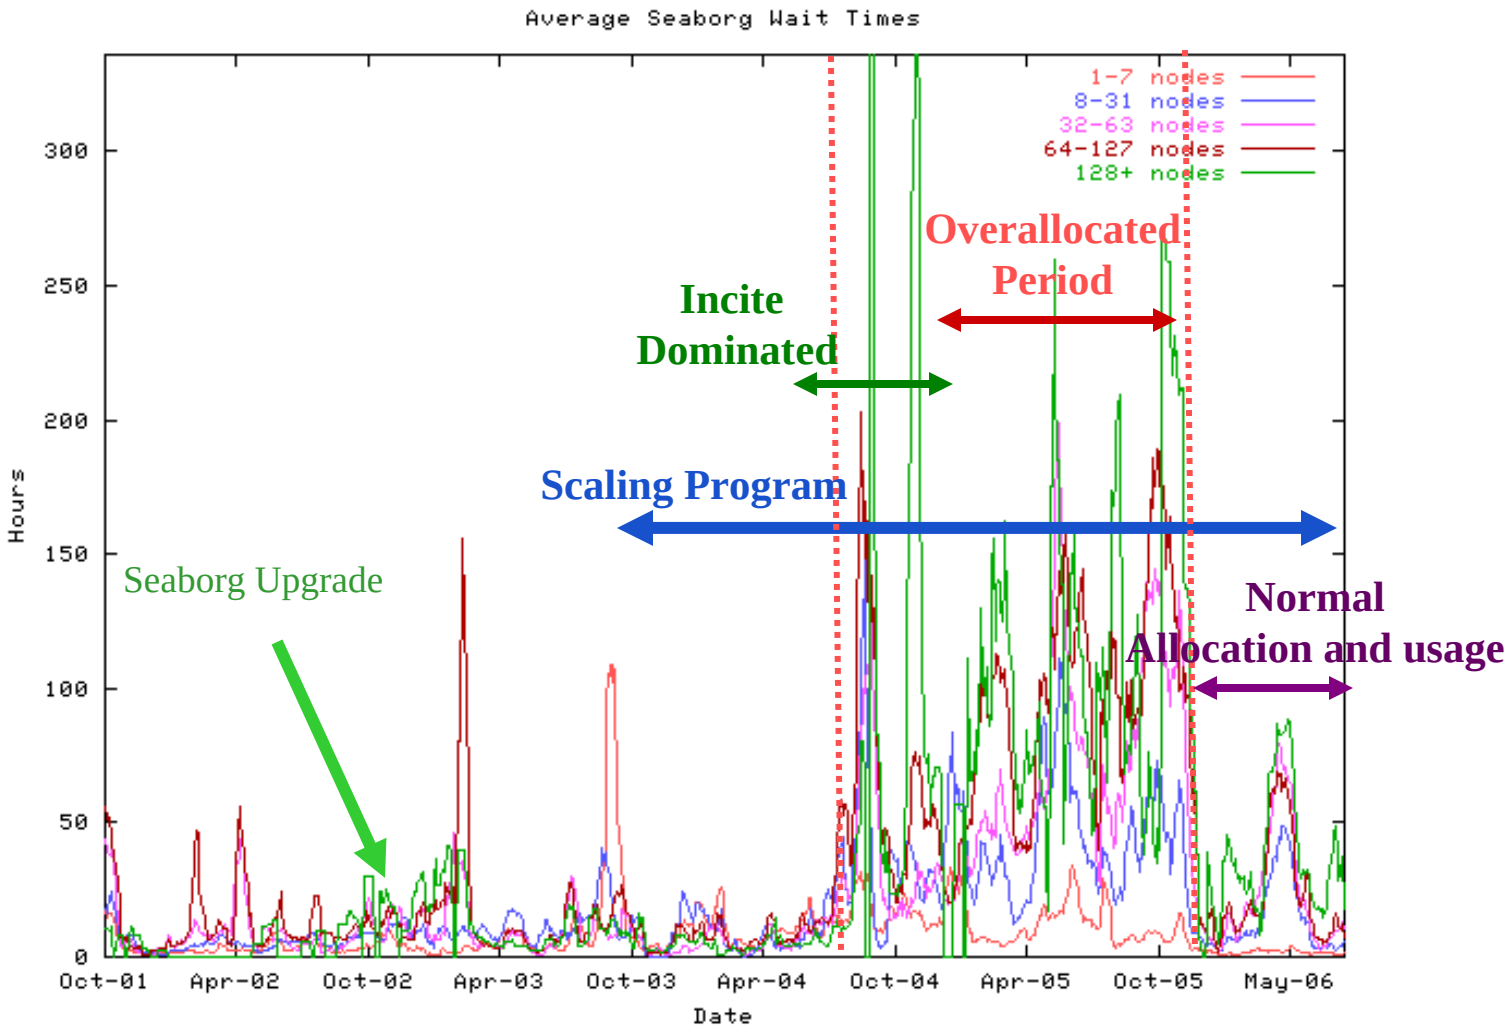

Yearly Publications

Each year on their allocation renewal form, PIs indicate how many refereed publications their project had in the previous 12 months.

Year of request renewal	Number of refereed publications
2007	1,437
2006	1,448
2005	1,270