Utilization & Availability Metrics

http://wk122.nas.nasa.gov/HPCC/Metrics/metrics.html

http://www-parallel.nas.nasa.gov/Parallel/Metrics/

Number of Unscheduled Interrupts

Counts the number of unscheduled service interruptions. An "interrupt" in this context could be a software failure (crash), hardware failure, forced reboot to clear a problem, facility problem, or any other unscheduled event that prevented a significant portion of the machine from being useful. The maximum acceptable value for each system is one per day, and the goal is less than one per week.

Number of Unscheduled Interrupts

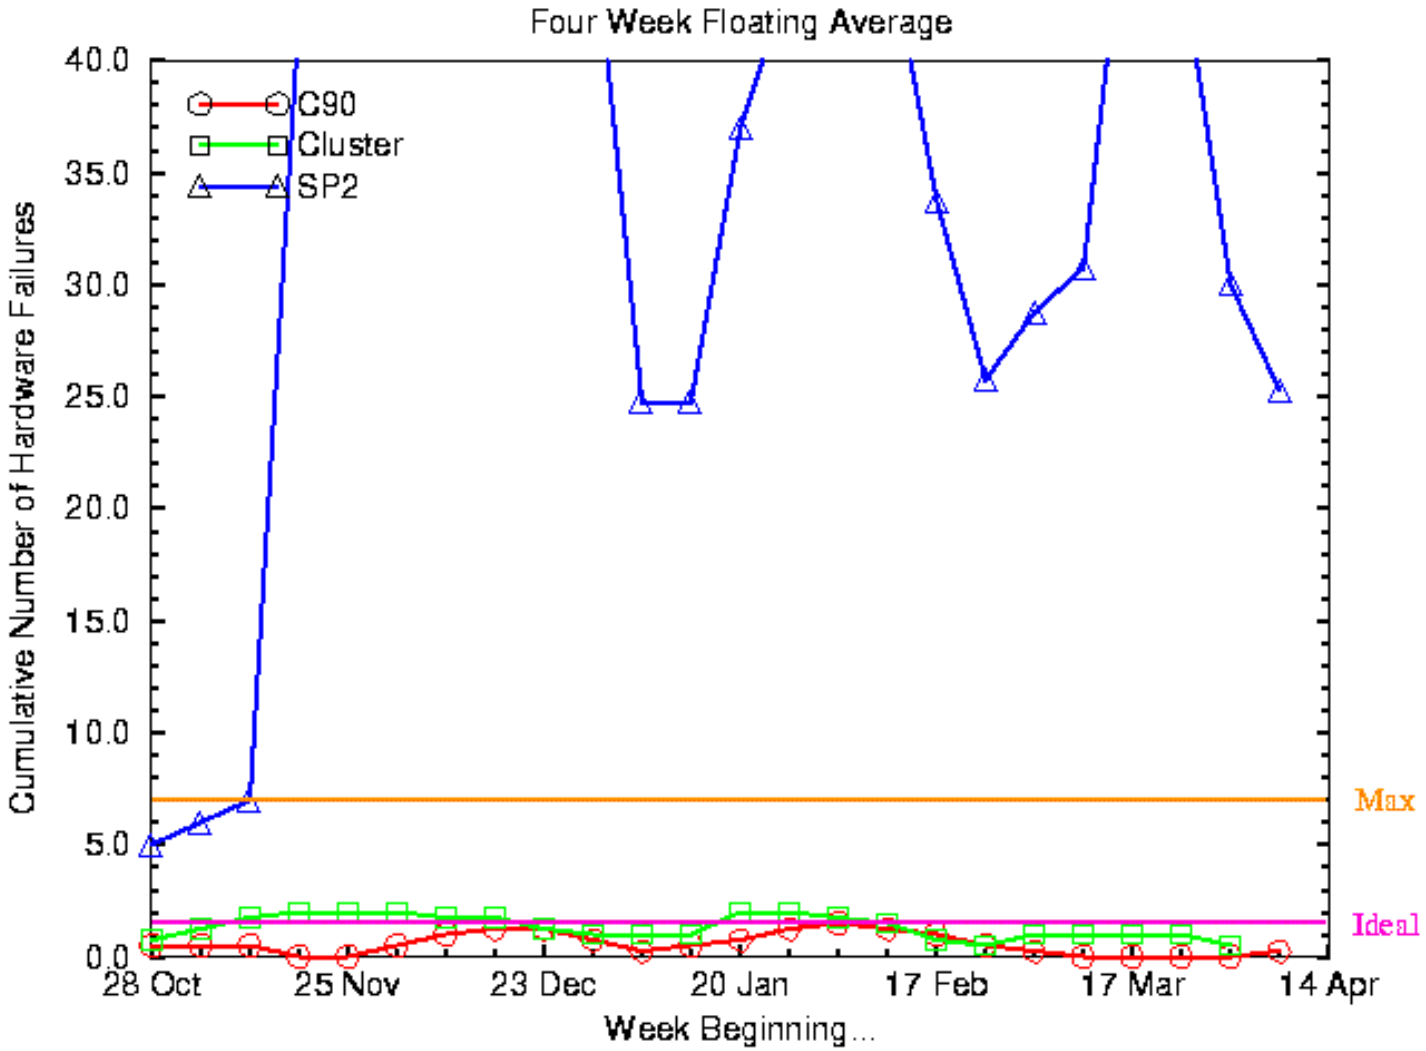

Gross Availability

Shows the fraction of time that the machine was available, regardless of whether it was scheduled to be up or down. Dedicated time reduces this number but has no effect on "net availability". Acceptable values and goals are not defined.

Gross Availability

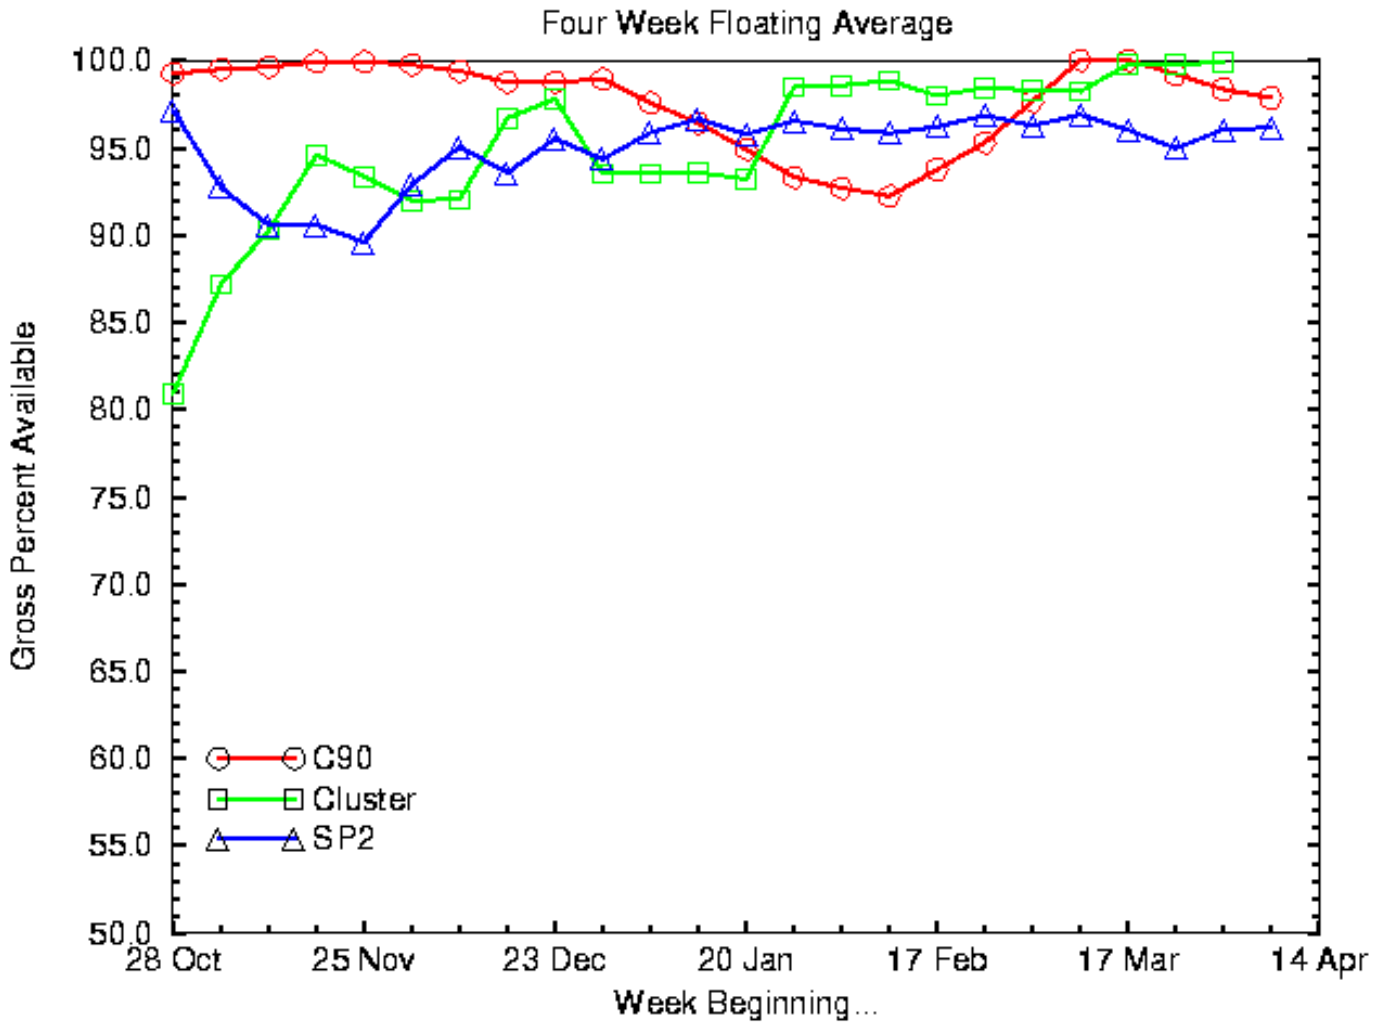

Hardware Failures

Count of hardware faults that required physical repair or replacement. Hardware failures do not necessarily result in an interrupt, because some machines (C90) perform error correction and can continue to run with damaged hardware. The replacement would then take place in scheduled maintenance time, counting as a hardware failure but not as an interrupt. The maximum acceptable value is one hardware failure per week, and the goal is less than one per month.

Hardware Failures

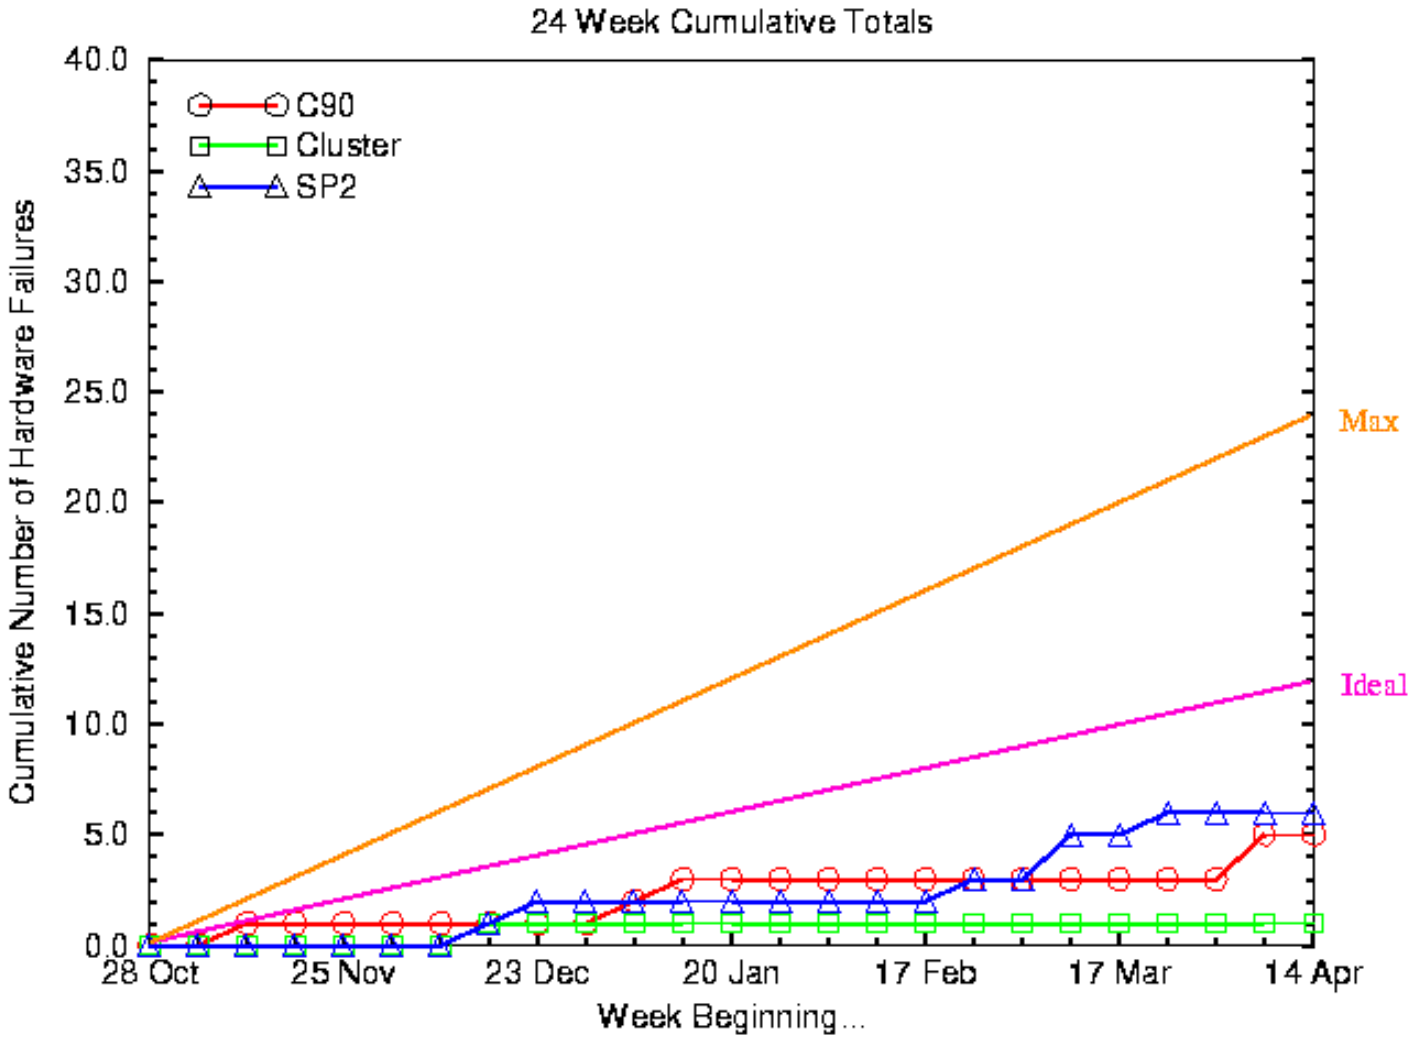

Load Average

Mean "load average" for a week. Sampled every fifteen minutes throughout the week and averaged. "Load average" is the average number of processes that are ready to run, reflecting the busy-ness of the machine. There are no goals yet set, and for most parallel systems, time-sharing is still a bad idea, hence the ideal load average should be close to 1 (or 1 times the number of compute nodes, if the load average is not normalized).

Load Average

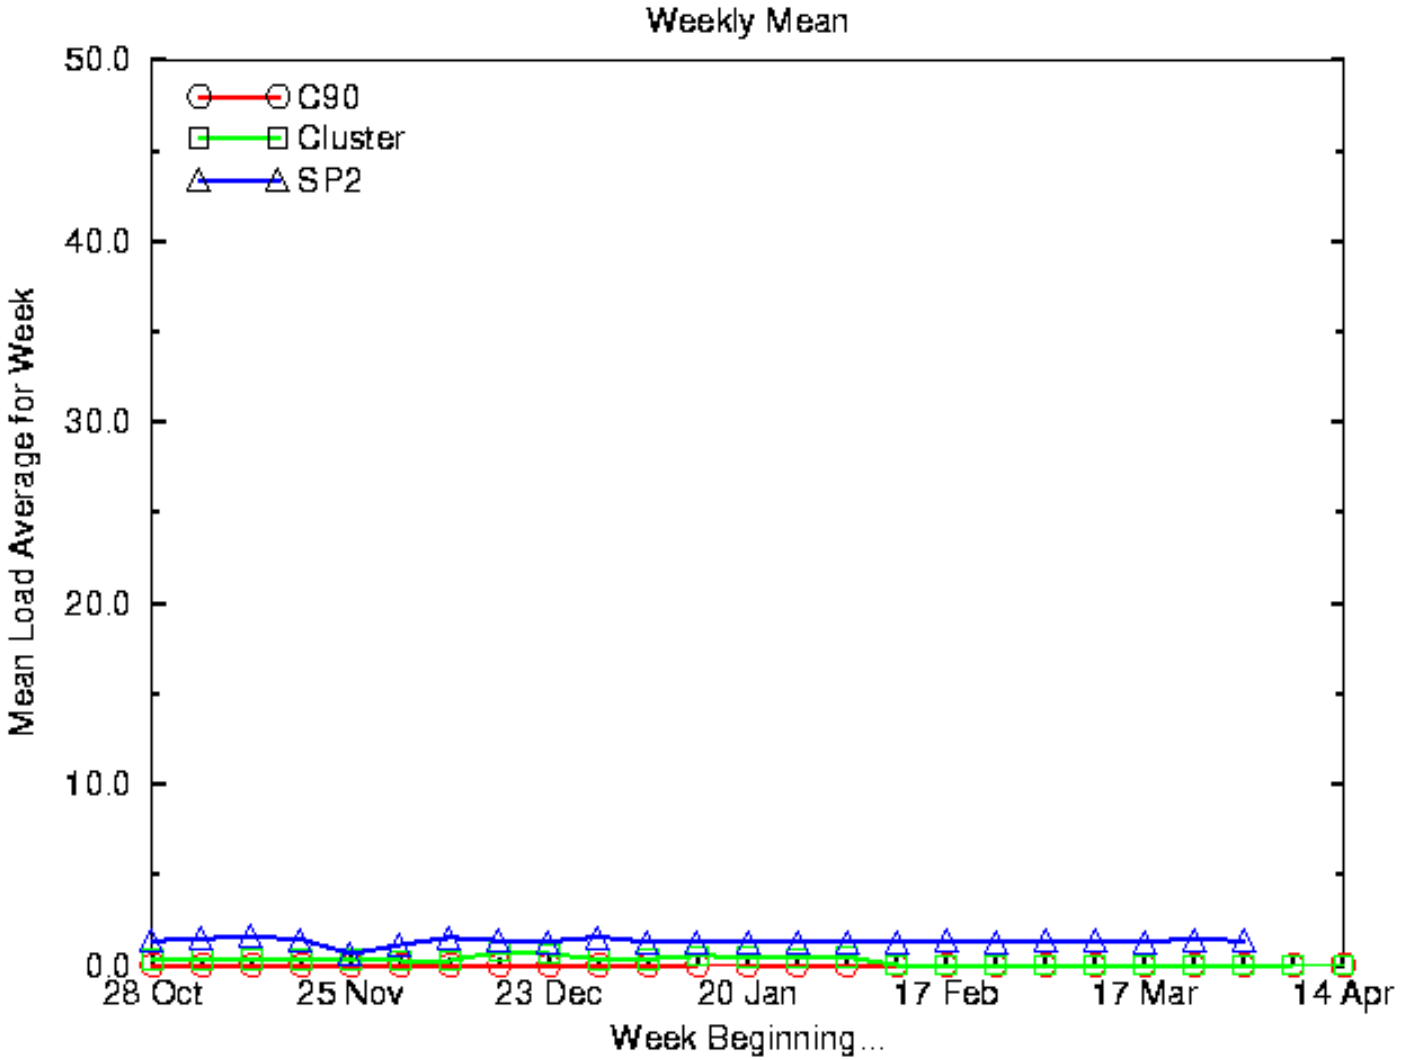

Mean Time Between Interrupts

Mean time between unscheduled service interruptions. Calculated by dividing the uptime by the number of interruptions plus one. Acceptable values and goals for this metric are not separately defined, but are dependent on the values chosen for the "Number of Unscheduled Interrupts" and "Mean Time To Repair" metrics.

Mean Time Between Interrupts

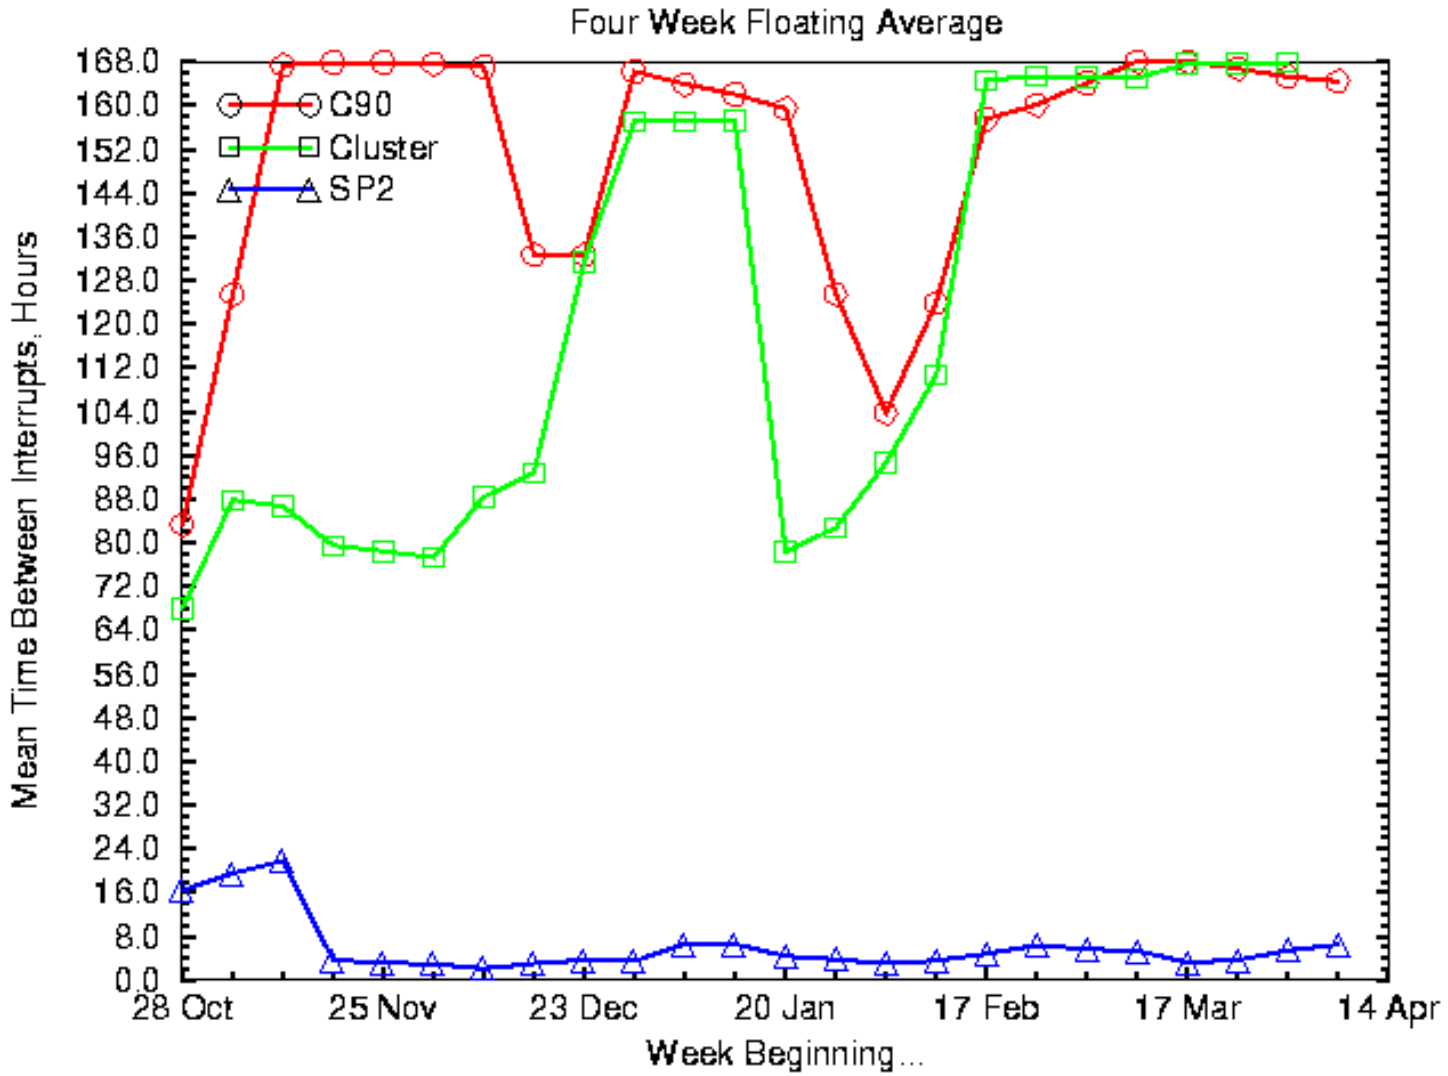

Mean Time To Repair

Average time required to bring the machine back up after an unscheduled interrupt. The outage begins when the machine is recognized to be down and ends when it begins processing again; it is not a reflection of the vendor FE response time. Maximum acceptable value is 4 hours, and the goal is less than one hour.

Mean Time to Repair

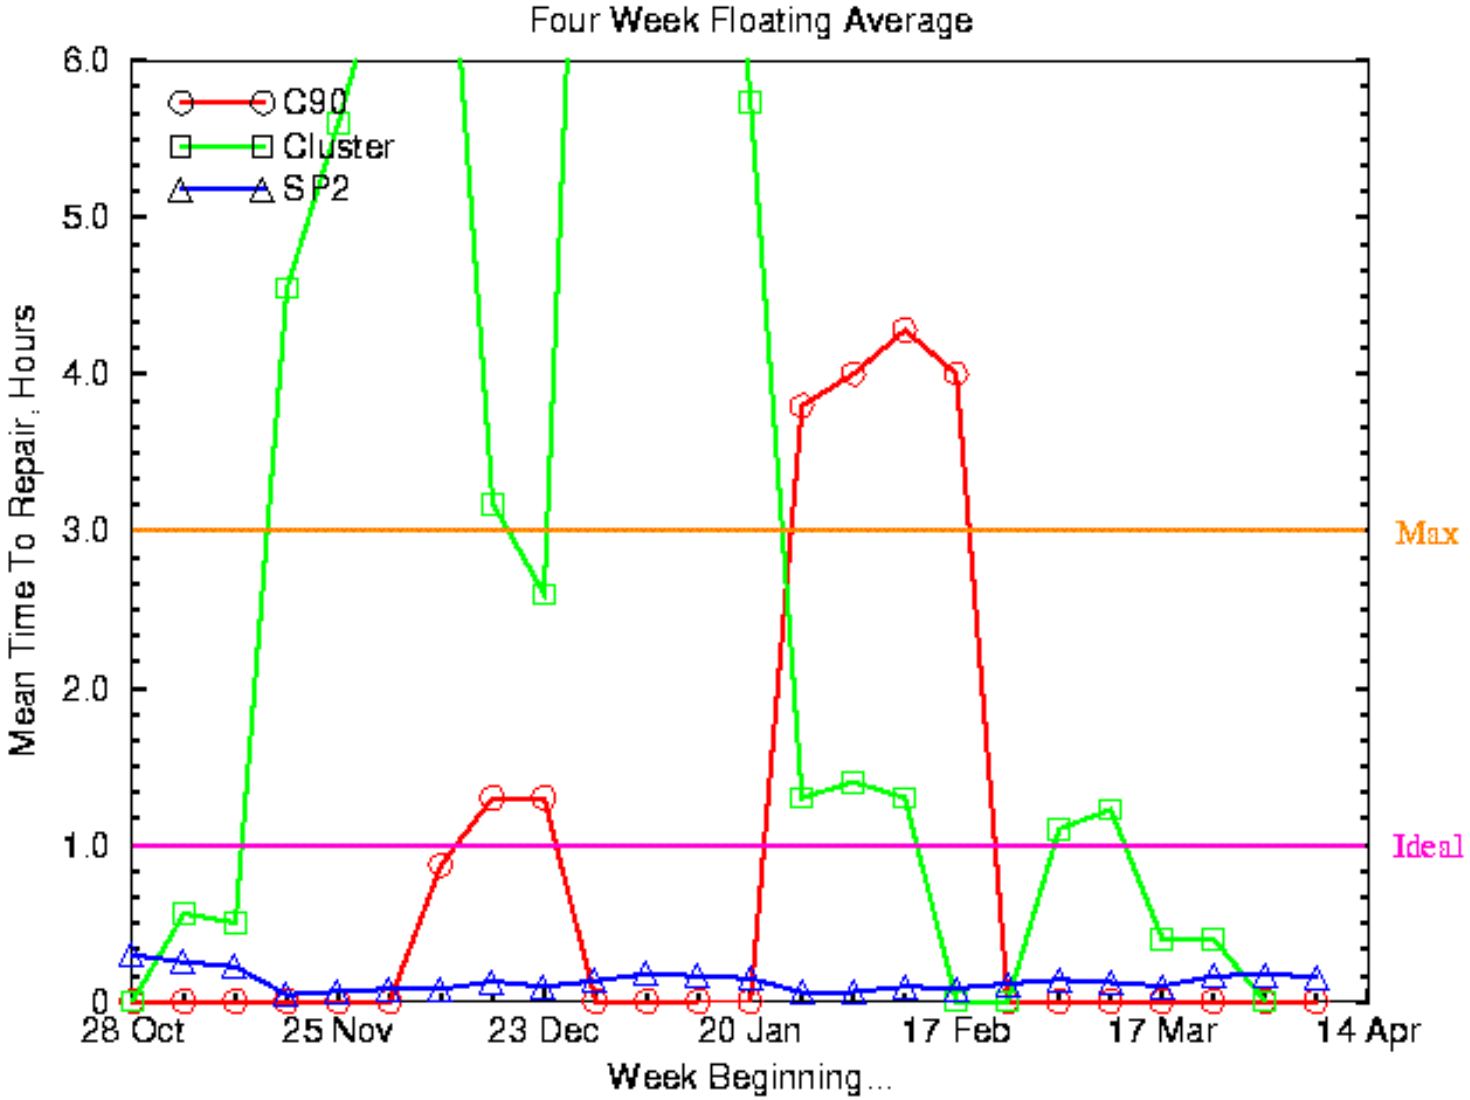

Net Availability

Shows the time that the machine was available as a percentage of the time that the machine was scheduled to be available. The minimum acceptable value is 95%, and the goal is 97%.

Net Availability

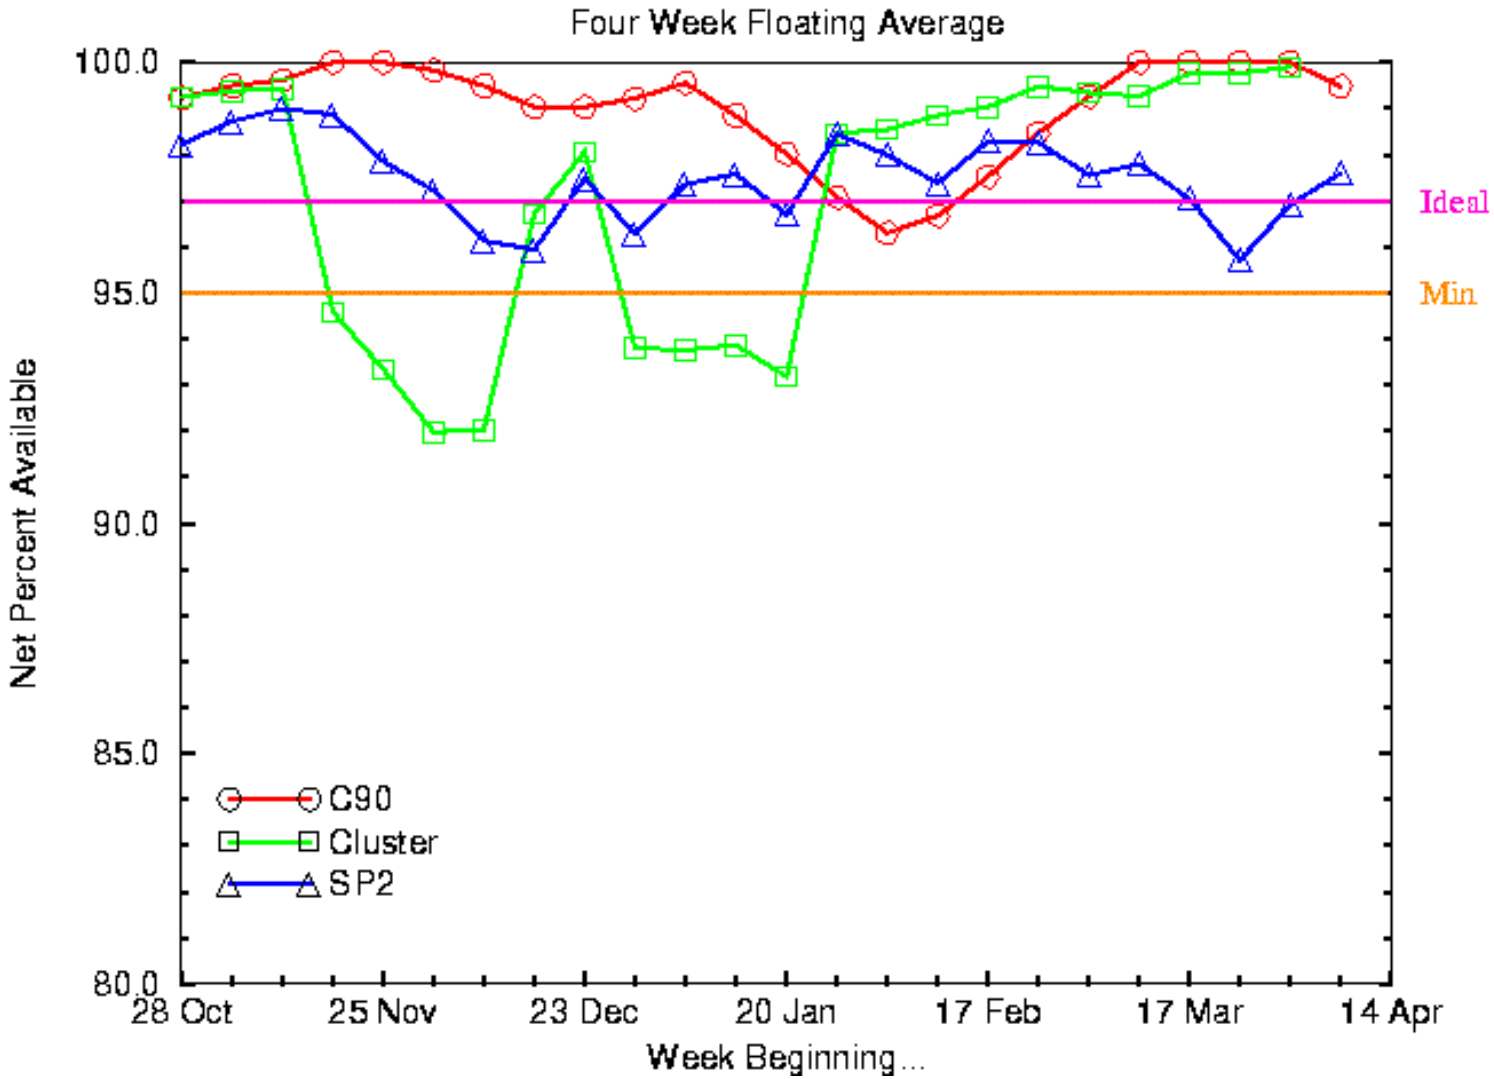

Number of Batch Jobs

Number of jobs run through the batch subsystem during each week. Not all jobs, even long jobs, are run though a batch scheduler on all systems. There are no defined goals for this metric.

Number of Batch Jobs

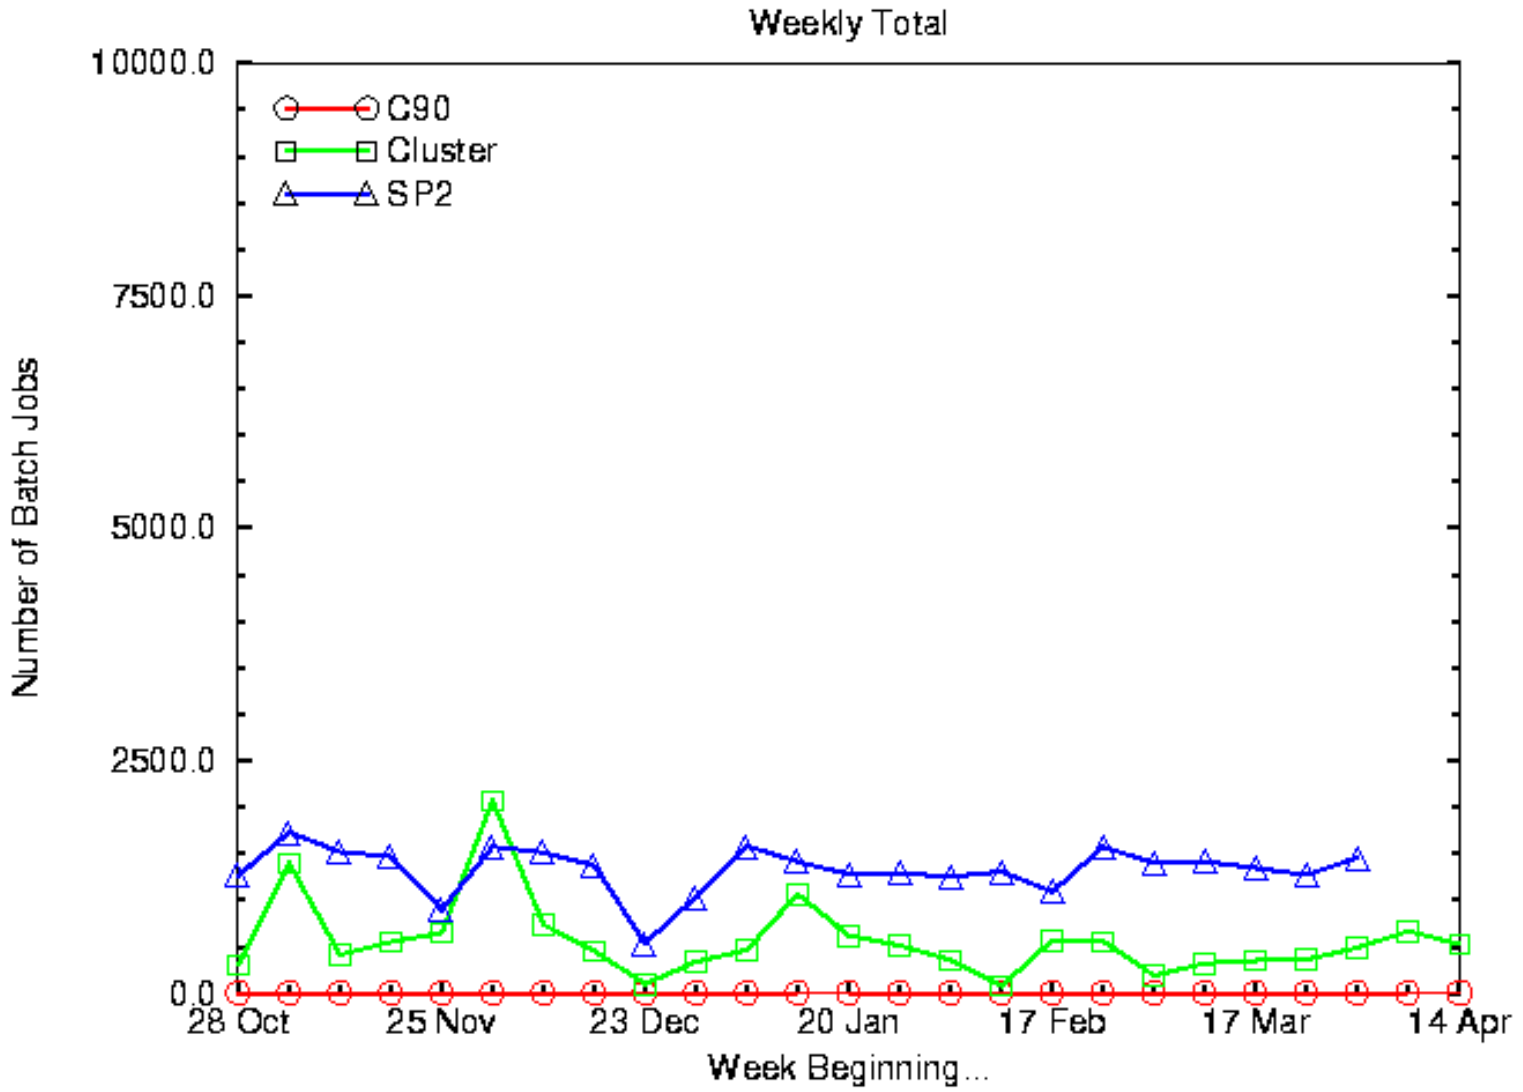

Number of Logins

Number of people logged in to the machine in question. Sampled every fifteen minutes throughout the week and averaged. This metric has no goals specified (for purely batch systems, a heavily used system may have few, or no, logins).

Number of Logins

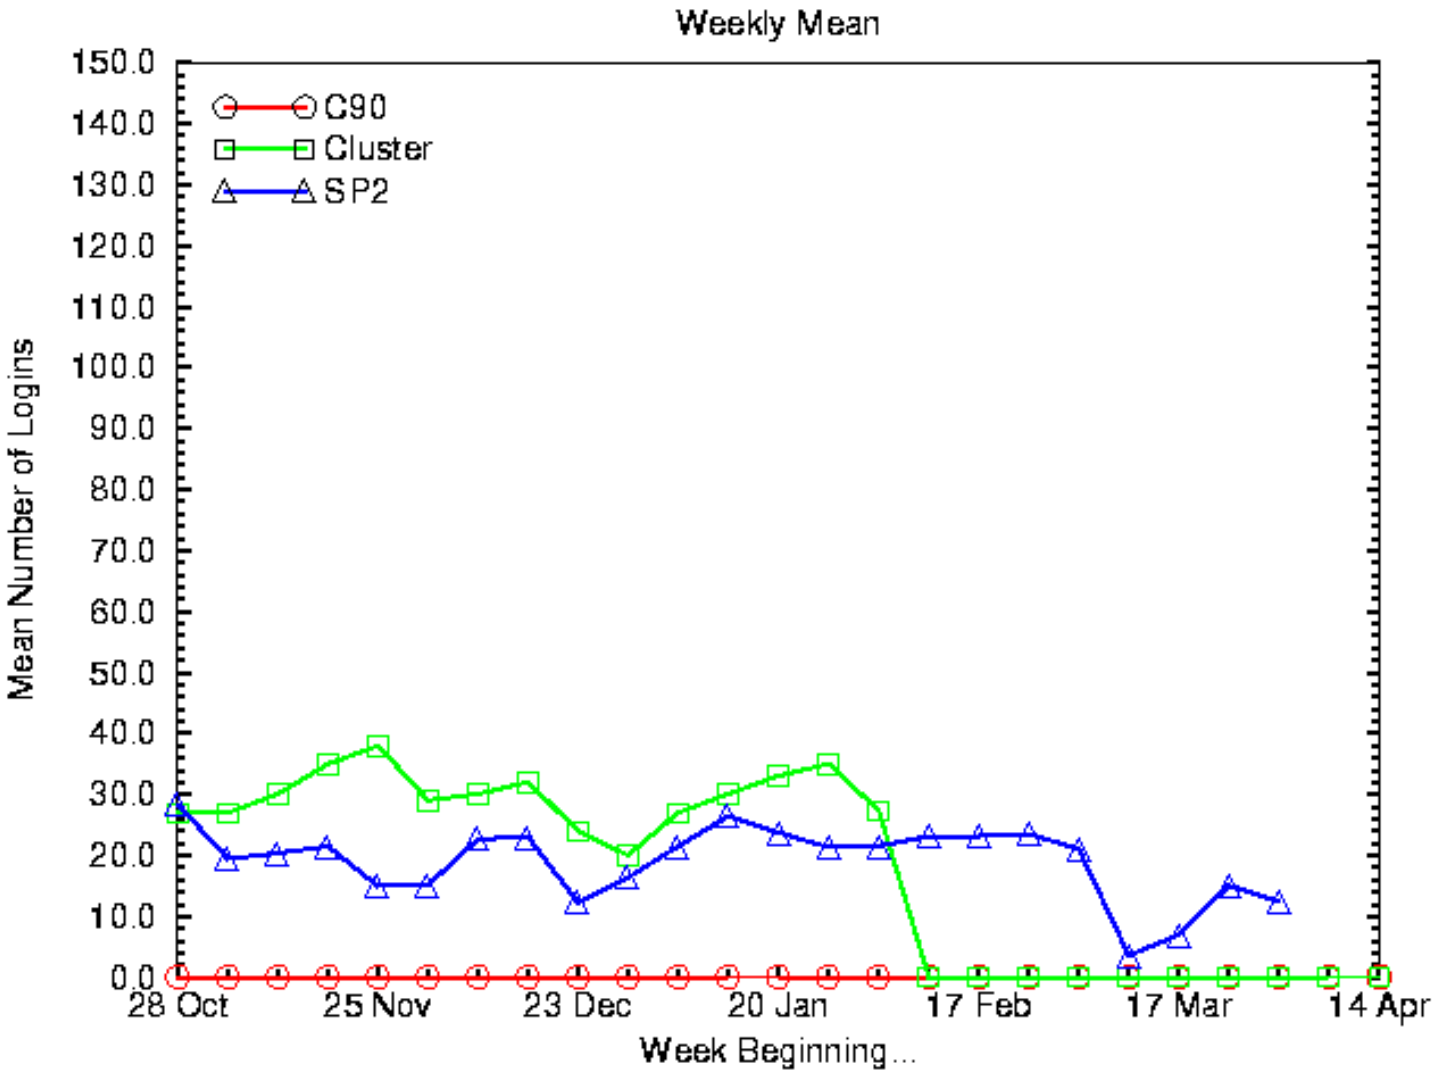

CPU Utilization

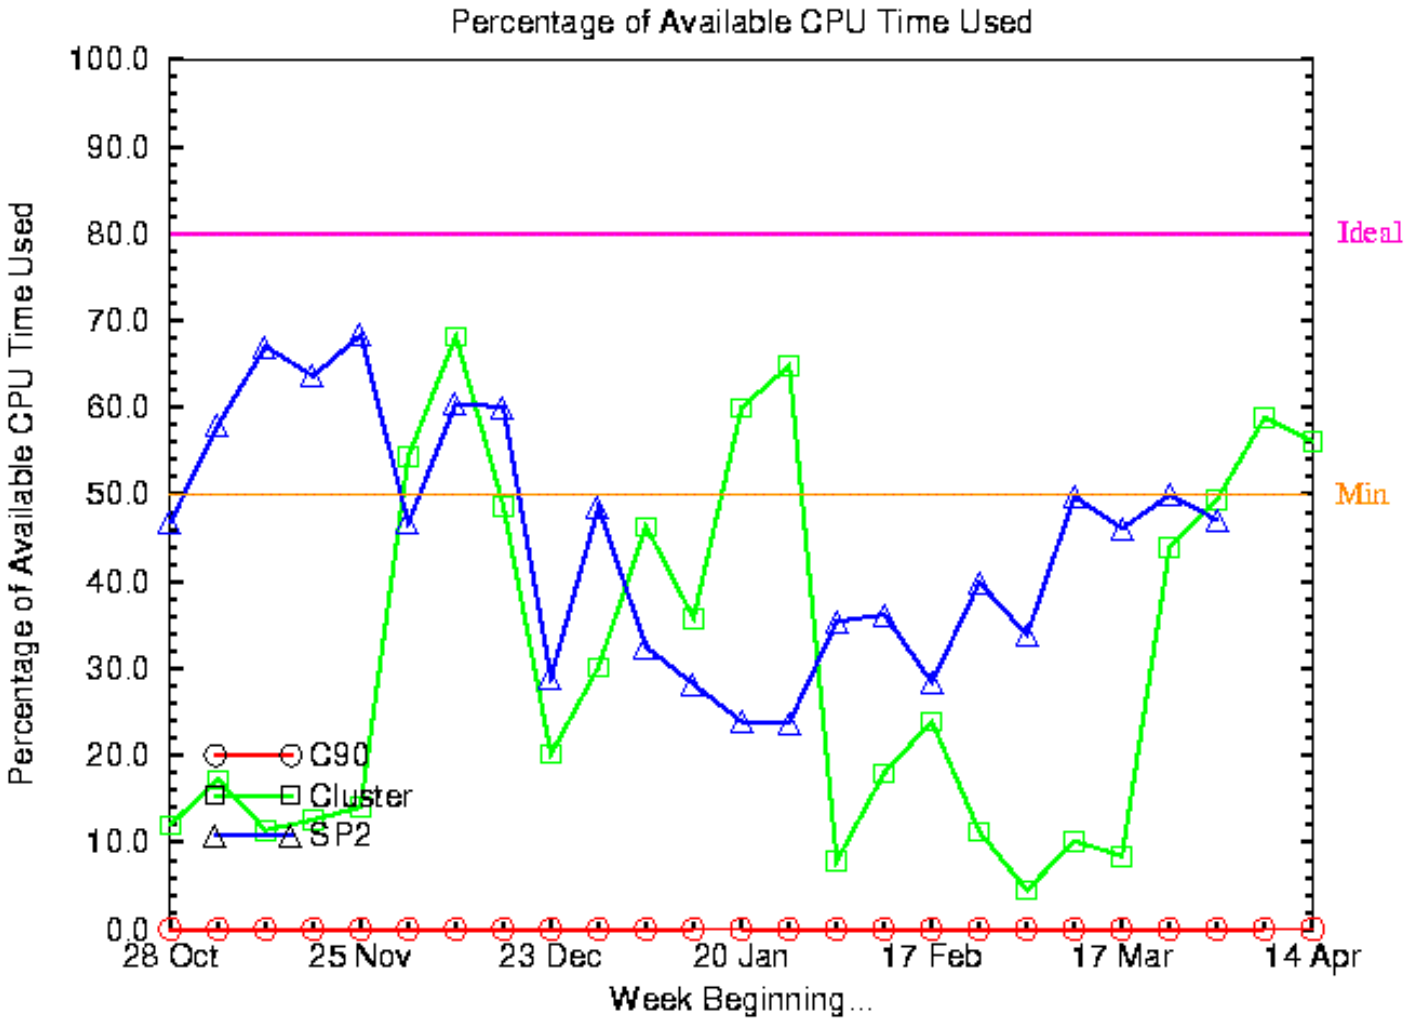

Memory Utilization

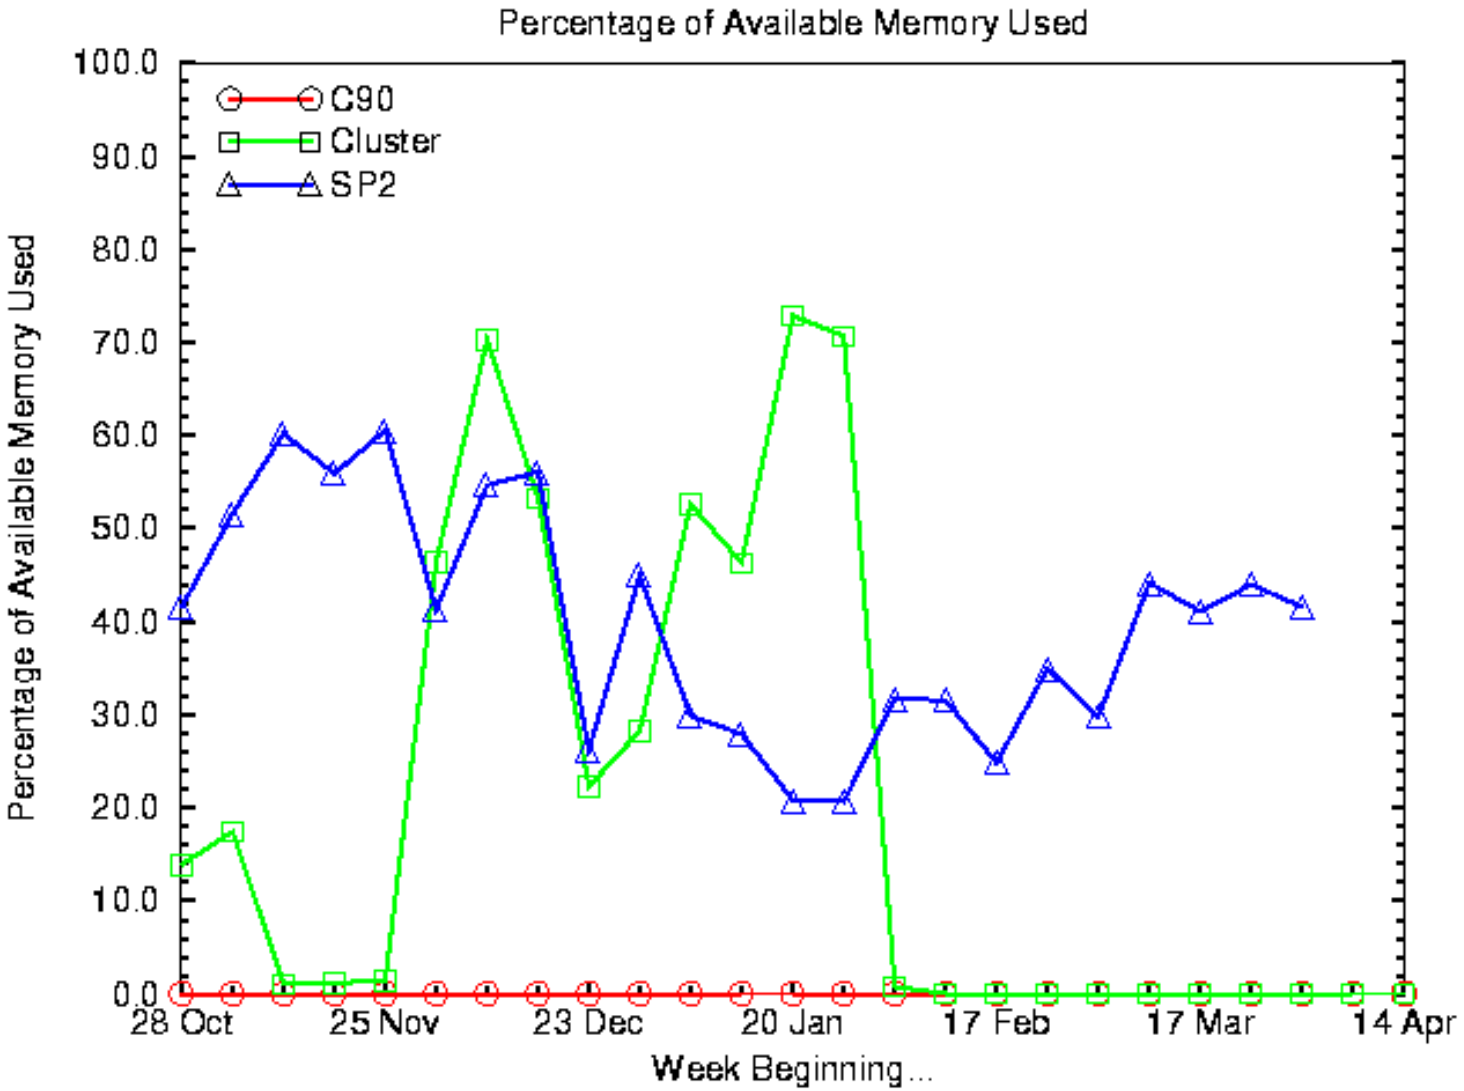

NAS HPCCP SP2 (Babbage) Utilization

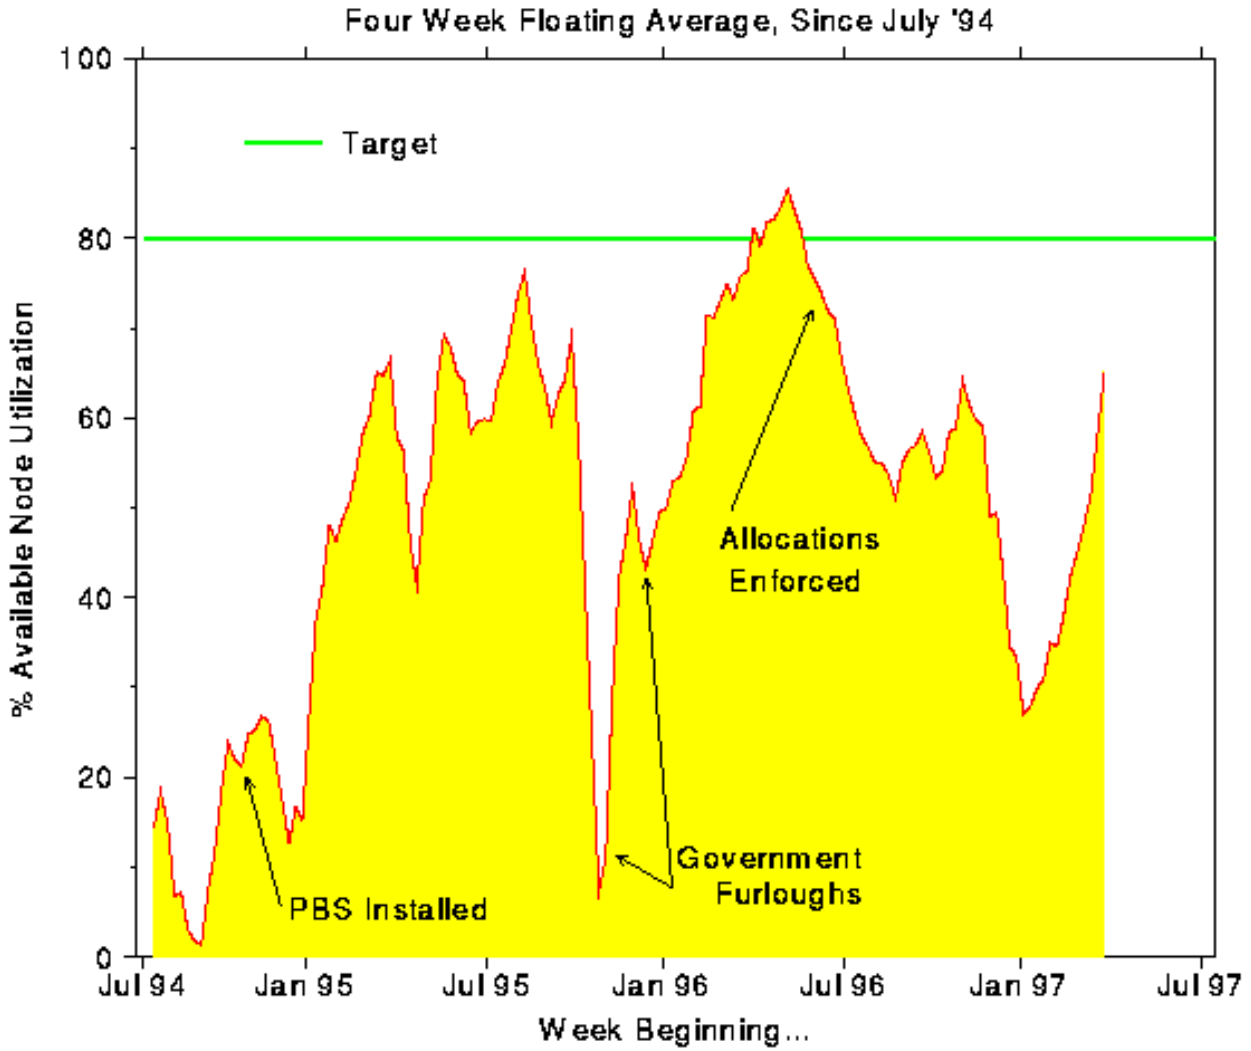

NAS HPCCP Cluster (Davinci) Utilization

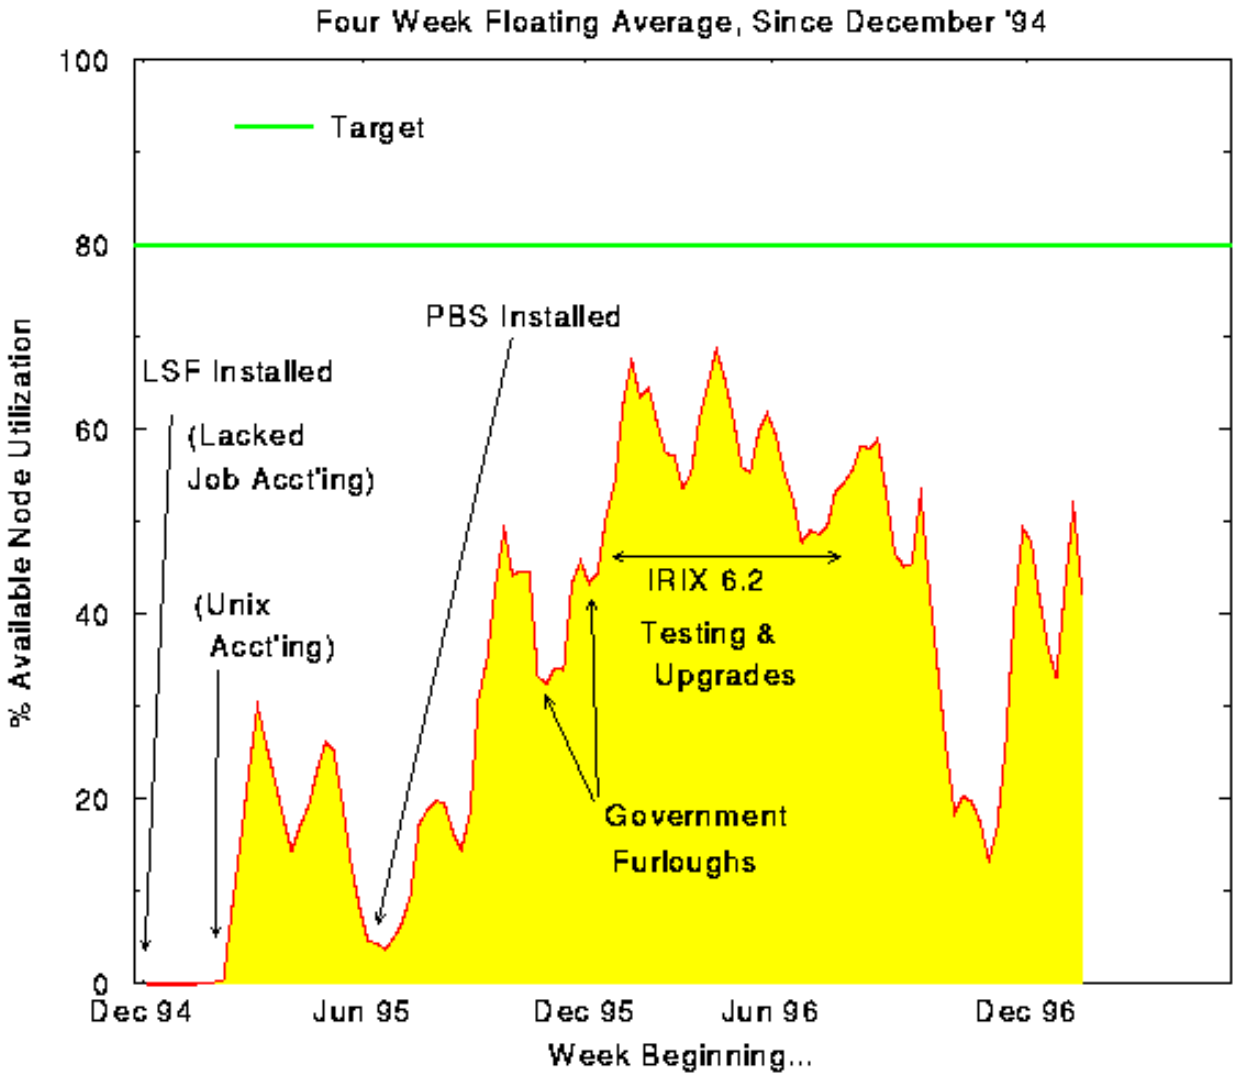

Meta Center SP2 Utilization

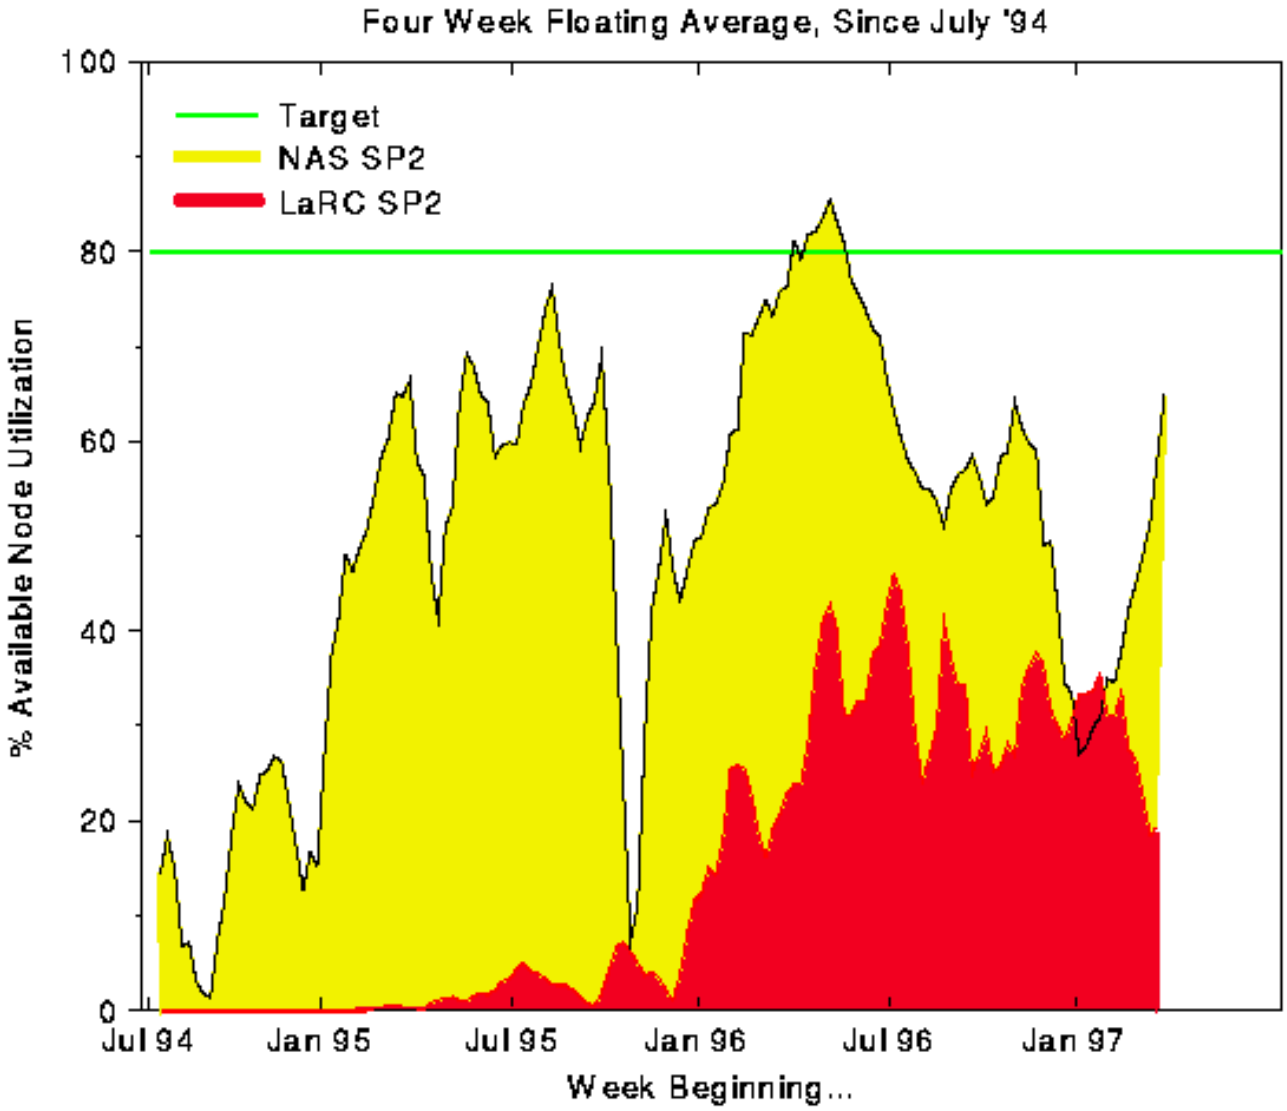

Meta Center SP2 Job Migration

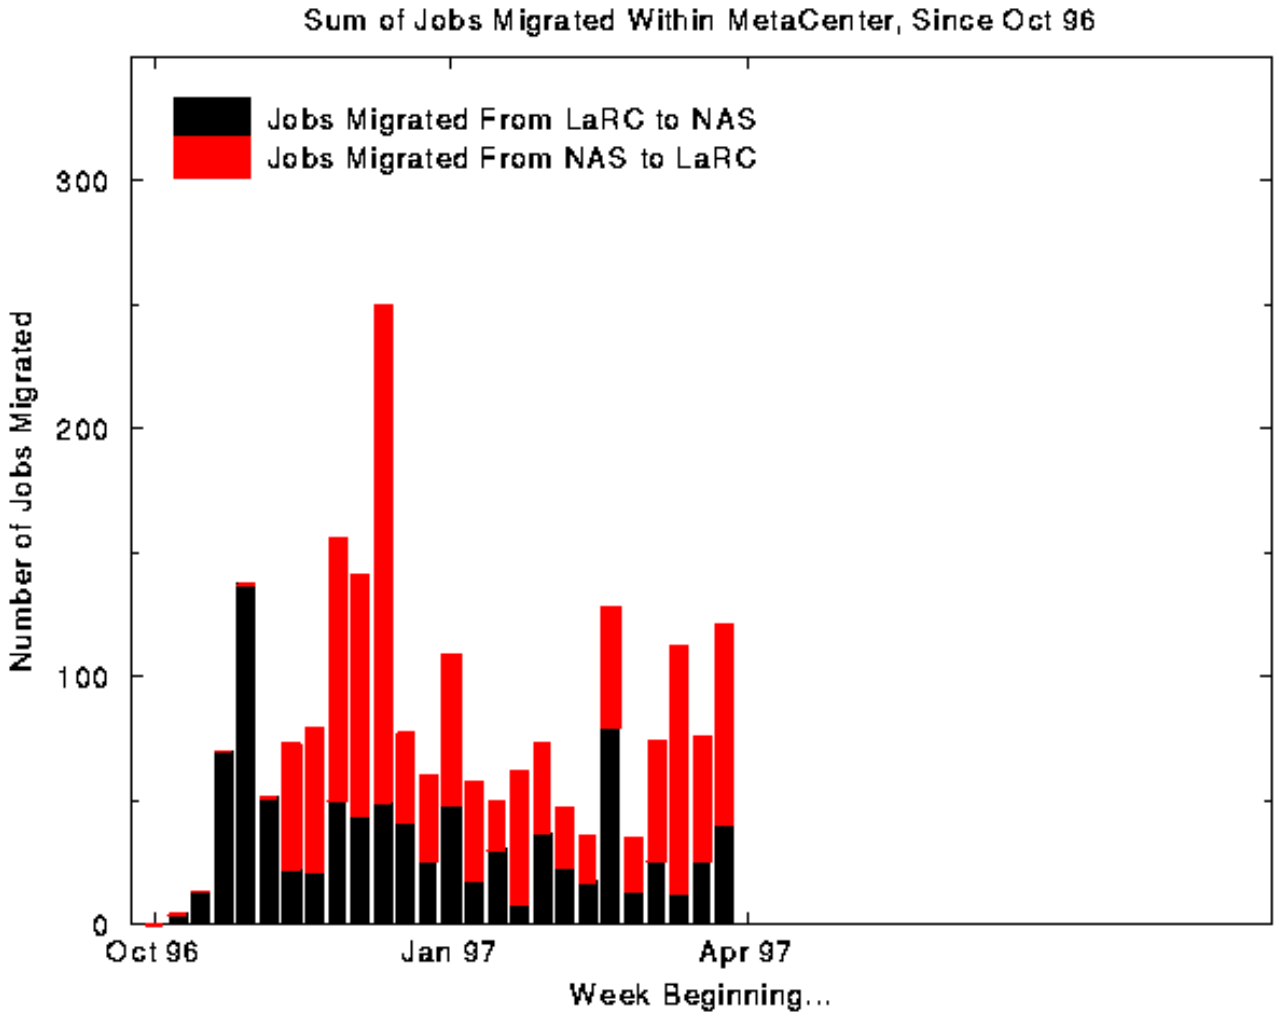
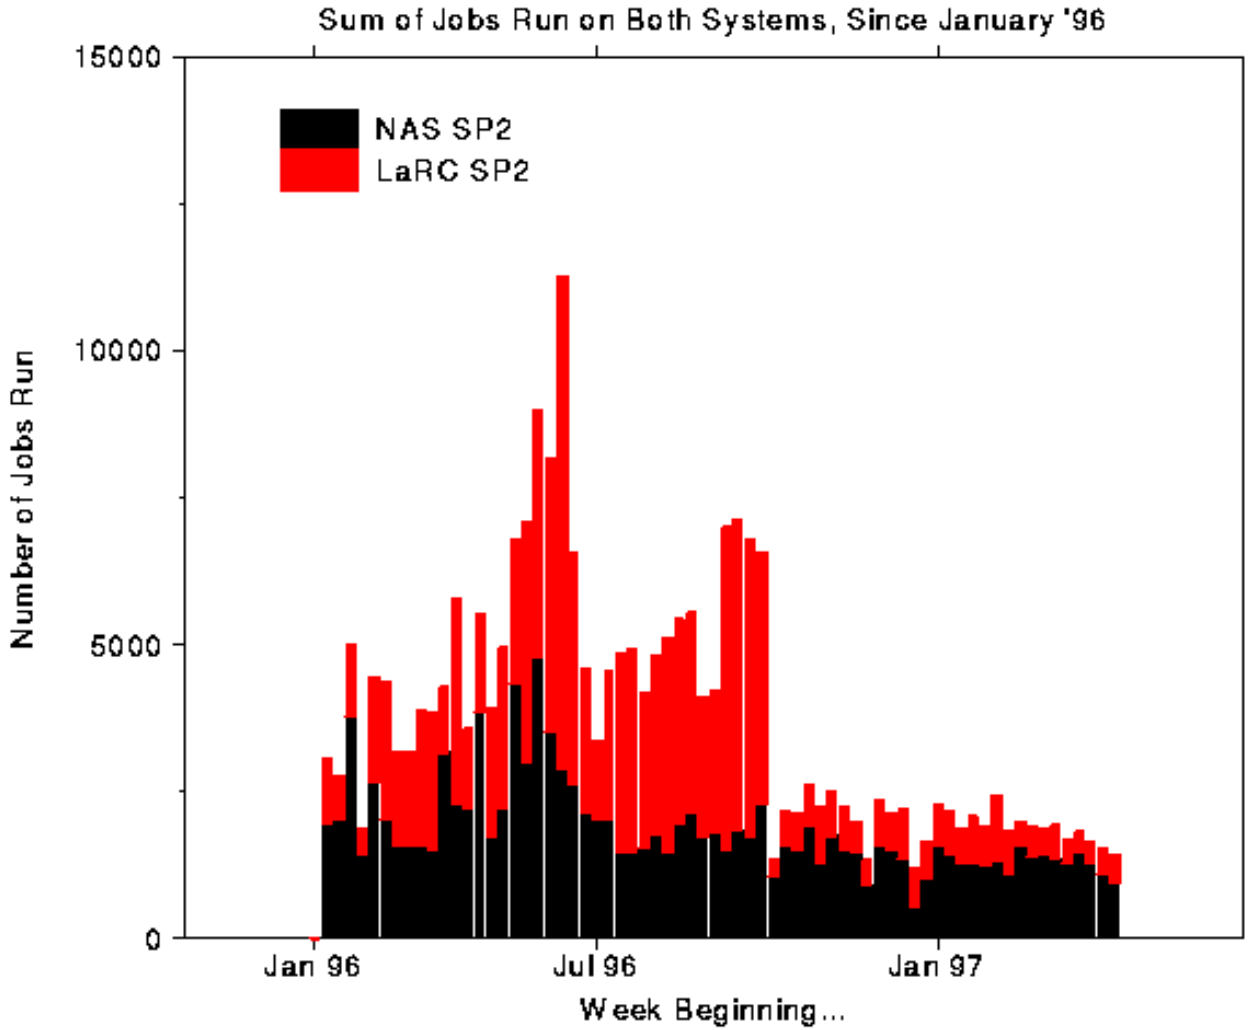
Meta Center SP2 Batch Jobs