Present: Alexandre Beche, Barry Blumenfeld, Simone Campana, Alastair
    Dewhurst, Alessandro Di Girolamo, Dave Dykstra, Andrea Valassi

The current monitoring I'm talking about transferring to WLCG
responsibility is based on MRTG and is here:
    http://frontier.cern.ch/squidstats/indexcms.html
    http://frontier.cern.ch/squidstats/indexatlas.html
    http://frontier.cern.ch/squidstats/indexcvmfs.html

MRTG is used mainly for expert debugging after problems are detected
    because of either SUM test failures or detection of failover
    traffic at the Frontier or CVMFS servers
    - There is an urgency to get this moved, because since MRTG does
    polling via SNMP it requires specific IP address(es) allowed in
    the firewall at most sites and the squid access control lists at
    all sites, and the current monitor machines have to retire at
    least by the end of May 2013. We want to minimize the number of
    times the address(es) change.
ATLAS also sees a need for better automated notifications of problems to
    shifters
    - This can be a separate phase, but the top priority is to adapt the
      existing CMS notification system based on awstats statistics from
      the central servers that automatically notifies administrators (by
      email) of sites causing failover traffic. It also graphs recent
      failover traffic:
    http://frontier.cern.ch/squidstats/nonproxycms_summary.html
    - The awstats tool the above is based on monitors the reverse-proxy
      squids on the central "launchpad" servers:
    http://frontier.cern.ch/awstatscms.html
    http://frontier.cern.ch/awstatsatlas.html
    http://frontier.cern.ch/awstatscvmfs.html
      I wasn't thinking of transitioning that also to WLCG, because it
      is only for the central servers, but maybe it should be; they are
      squids, after all.

What does it mean to transition squid monitoring to WLCG?
    - do things the common WLCG way, report to the WLCG organization
    - may or may not be something run by CERN/IT. Might be the same
    people running it now, just in a slightly different way.

What about existing WLCG monitoring, could it be adapted instead?
    - WLCG Dashboard: doesn't have this type of real-time monitoring
    - SUM: can test direct connections (it doesn't need to use grid jobs),
    but even with that it cannot be as real time as MRTG; it is more
    like at most every half hour
    - A tool that gives MRTG-like frequently-updated performance
    information is required for debugging by the experts

The main development needed is to auto-generate the MRTG configuration
    from a common information database
    - ATLAS currently bases theirs on their own AGIS information system,
    and CMS configures theirs by hand but has an automated audit
    comparing it to another source of information about squids
    (local site configurations checked in to CERN CVS)
    - GOCDB & OIM are the WLCG ways to define information about sites
    - OSG/U.S. uses OIM, Europe uses GOCDB
    - we're not sure what to do with non-grid tier 3 sites
    - Alastair will check with GOCDB people, and I will check
        with Doug Benjamin about OIM and we will report what we
        found to the task force group
    - Alessandro says its no problem for AGIS to change to get its squid
        information from GOCDB/OIM instead of being the primary source

---