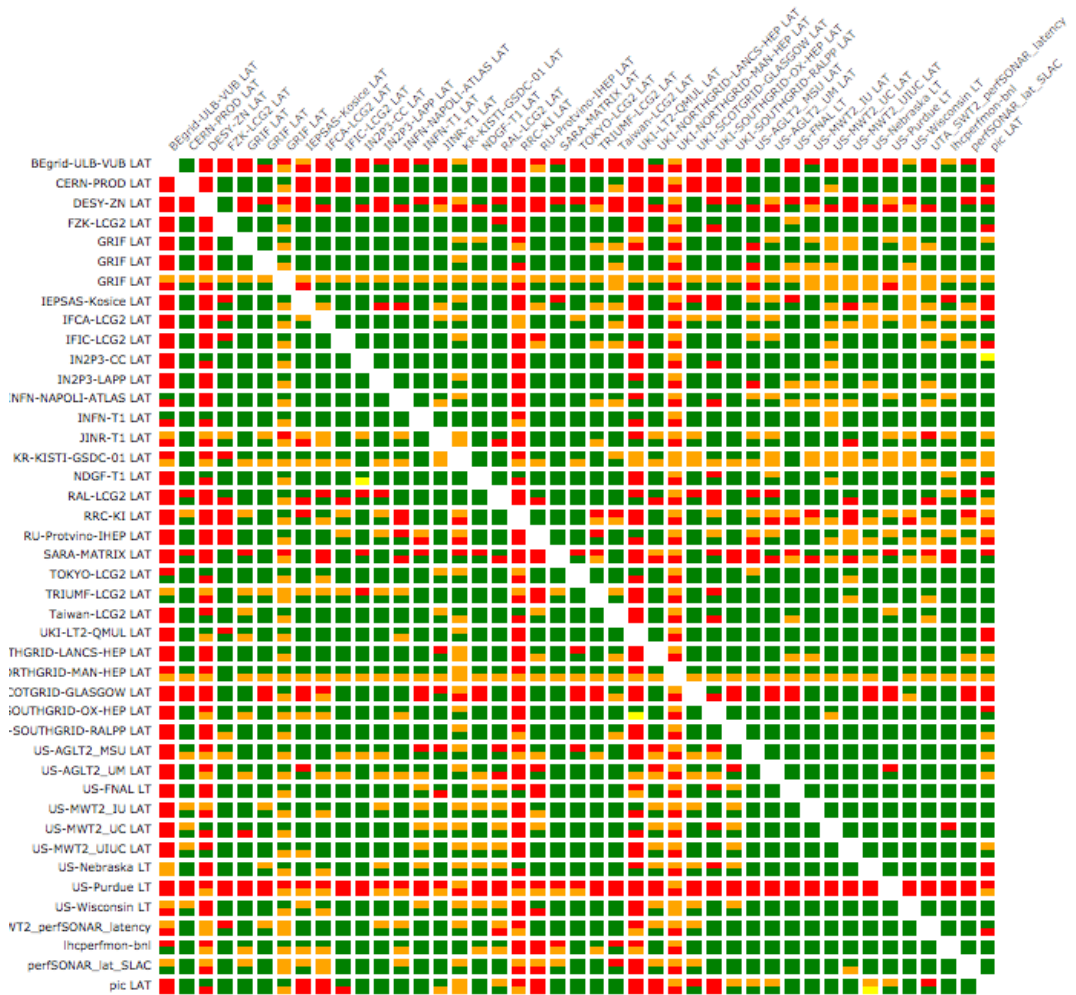
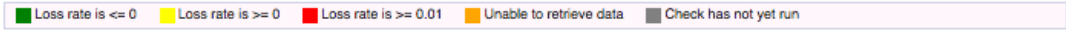


# Latency tests between all WLCG hosts - Latency Tests Between WLCG Latency Hosts



# WLCG Latency Mesh Comments

- +
  - It can be done, works consistently and already provides useful data
  - Latency mesh stable, once configured sonars are stable
  - Delivering ~ 98% of theoretical full mesh (45 sonars, ~1940 of 1980 links monitored, continuous 10Hz, each link is tested twice)
  - Very good response from sites, most of the issues were trivial to fix (ntp, regular testing not working, mesh configuration incorrect/missing, DNS aliases used)
  - Many new sonars appeared recently - only discovered during ramp up
- -
  - Many different hardware profiles (CPU, memory) – some sonars have 2GB memory and 32bit CPU ? perfSONAR 3.4 needs at least 4GB
  - Ramp up required follow up with sites, even if all tests OK in OMD, issues only pop up (visible) once sonar is added to the mesh
  - Some issues still remain, KISTI, JINR-T1 don't monitor certain links, US-EU links show some degradation during certain period of the day (under investigation)

# WLCG-wide meshes

- Proposal to create full mesh latency, tracepath and bandwidth tests
  - Traceroutes (hourly), Latency (10Hz), Throughput (weekly)
  - Needed to accomplish the use cases as described by the experiments in the use case document
  - It can significantly help in debugging network issues as you can see the situation across many sites
  - Establish sonar core - sonars that work fine and can test to each other
  - Determine the limits of the perfSONAR testing
  - Simplify monitoring and debugging of the sonars – one mesh for given event type is a lot easier to follow up than six

# WLCG-wide meshes

- Incremental approach to determine list of core sonars
  - We start with sonars/sites already in the WLCG latency mesh
  - Add only few sonars at a time (3-5), check their performance and performance of the meshes, contact sites to fix issues
  - Initial target would be to reach coverage of top 100 sites (based on contributed storage and location)
  - So we end up with 3 meshes, 100 sonars each, testing latency, traces and bandwidth
- For the duration of the ramp up, we disable all regional and project meshes – this is needed to ensure sonars are not overloaded
- Establish global mesh – with all inactive sonars (eventually those with hardware or configuration issues) – sonars still useful for debugging purposes
- We would use global mesh in the transition period to ensure every sonar in some mesh – this is needed as we're unable to run sonars with empty meshes
- Once we have WLCG-wide meshes, we restore regional and project meshes
  - For project meshes, recommendation will be to prefer using sonars already in the WLCG-wide mesh
- Manually configured tests unaffected