



ALICE internal and external network requirements @T2s

WLCG Workshop 2017
University of Manchester

20 June 2017
Latchezar Betev

The ALICE Grid today



Average CPU count

~130 CPU cores - April 2017

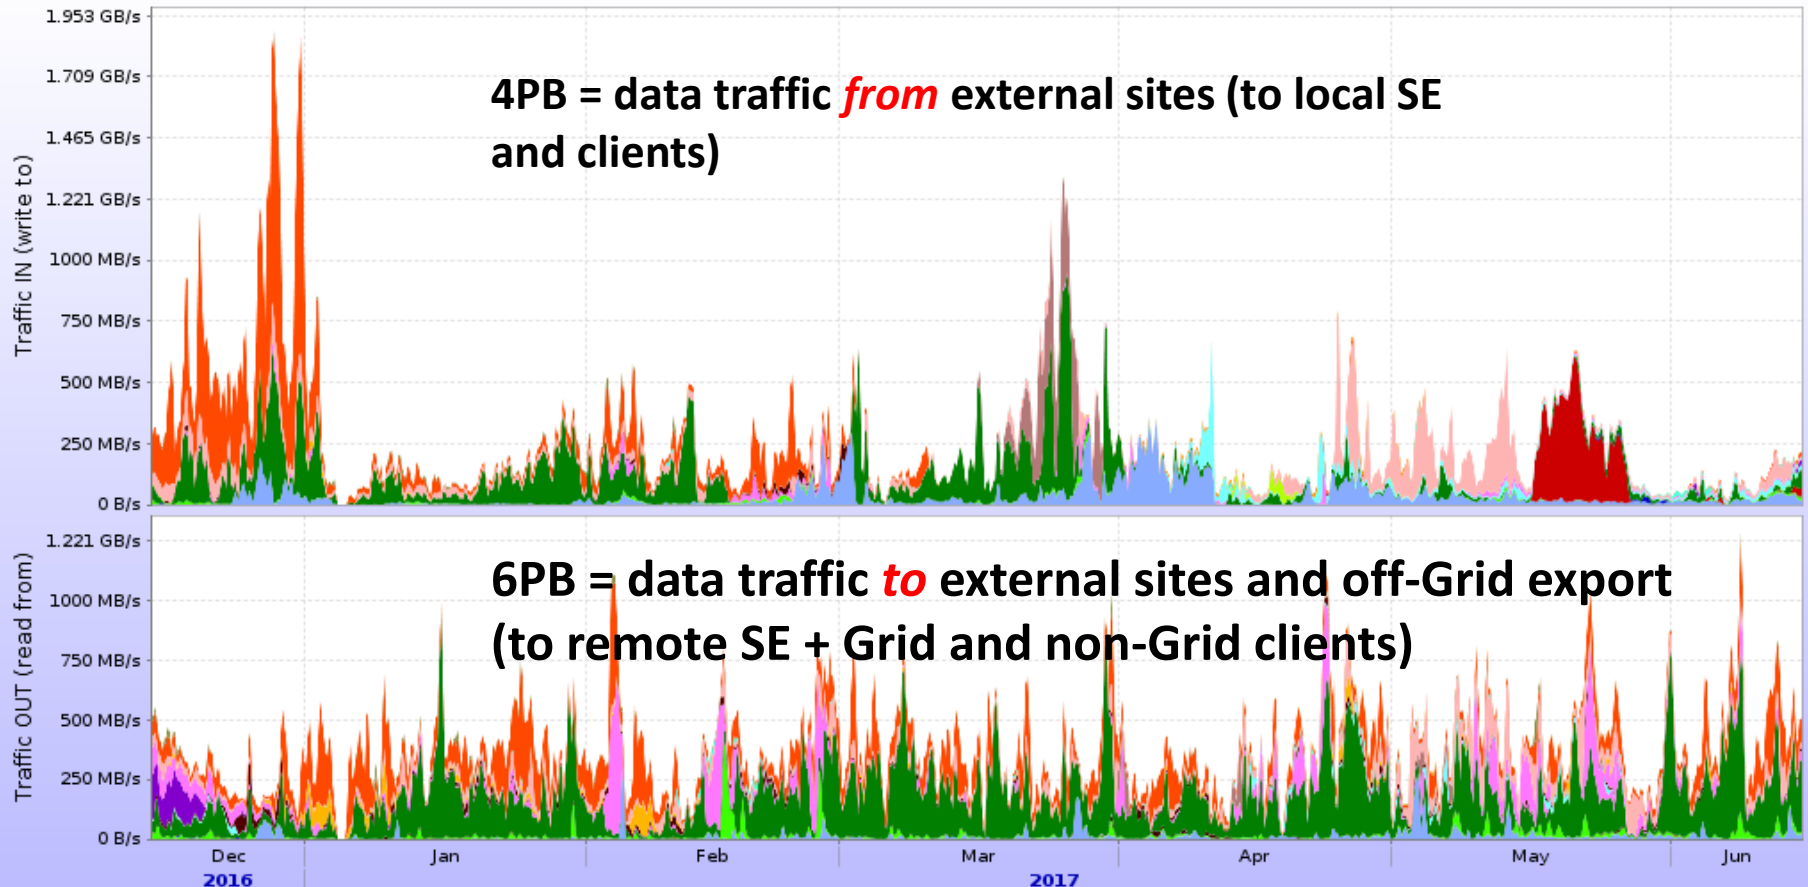


60% @T0+T1s, 40% @T2s

WAN @T2s

WAN server traffic

4PB = data traffic *from* external sites (to local SE and clients)

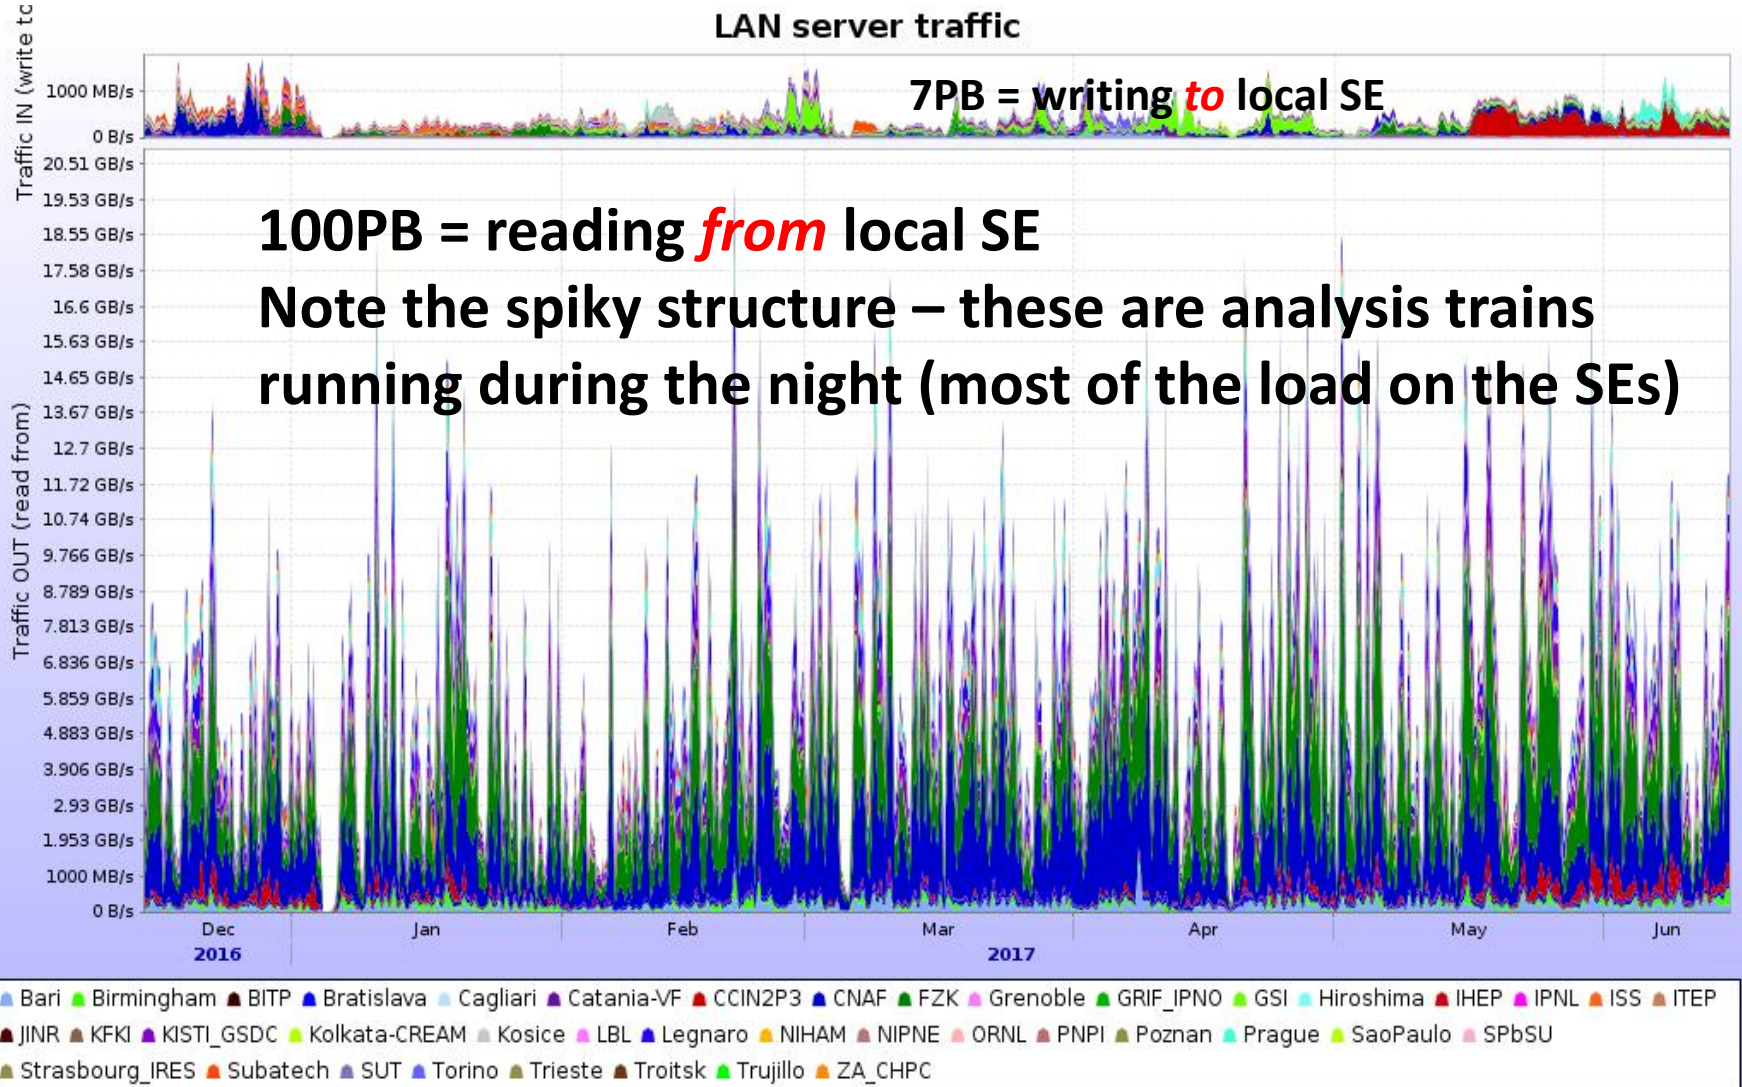


WAN Numbers @T2s

- Total WAN traffic are 1/11 of LAN traffic
 - WAN total T2s = 650MB/sec => **12.5KB/sec/core** (52K cores)
- Our ‘advisory rule’ for T2s is “**100Mb WAN network per 1K cores**”
 - Usually the available bandwidth is superior to the above
 - ~Independent of T2 role – a diskless T2 has to export ~2x the produced data; T2 with an SE exports one copy and accepts copies from other centres.

LAN @T2s

LAN server traffic



LAN Numbers @T2s

- LAN sustains by far the highest traffic (jobs go to data)
 - About 2% of data is read remotely (for example, local SE issue)
 - 90% of the LAN traffic is analysis, 10% is local + remote clients write (for a T2 with SE)
 - Diskless T2s - LAN=WAN traffic
- Figure of merit (for 85%+ CPU efficiency):
2.5MB/sec/core from local SE
- Remote access penalty
 - RTT is introducing efficiency penalty of ~15% per 20ms
 - Remote read is feasible for small data samples and as failover

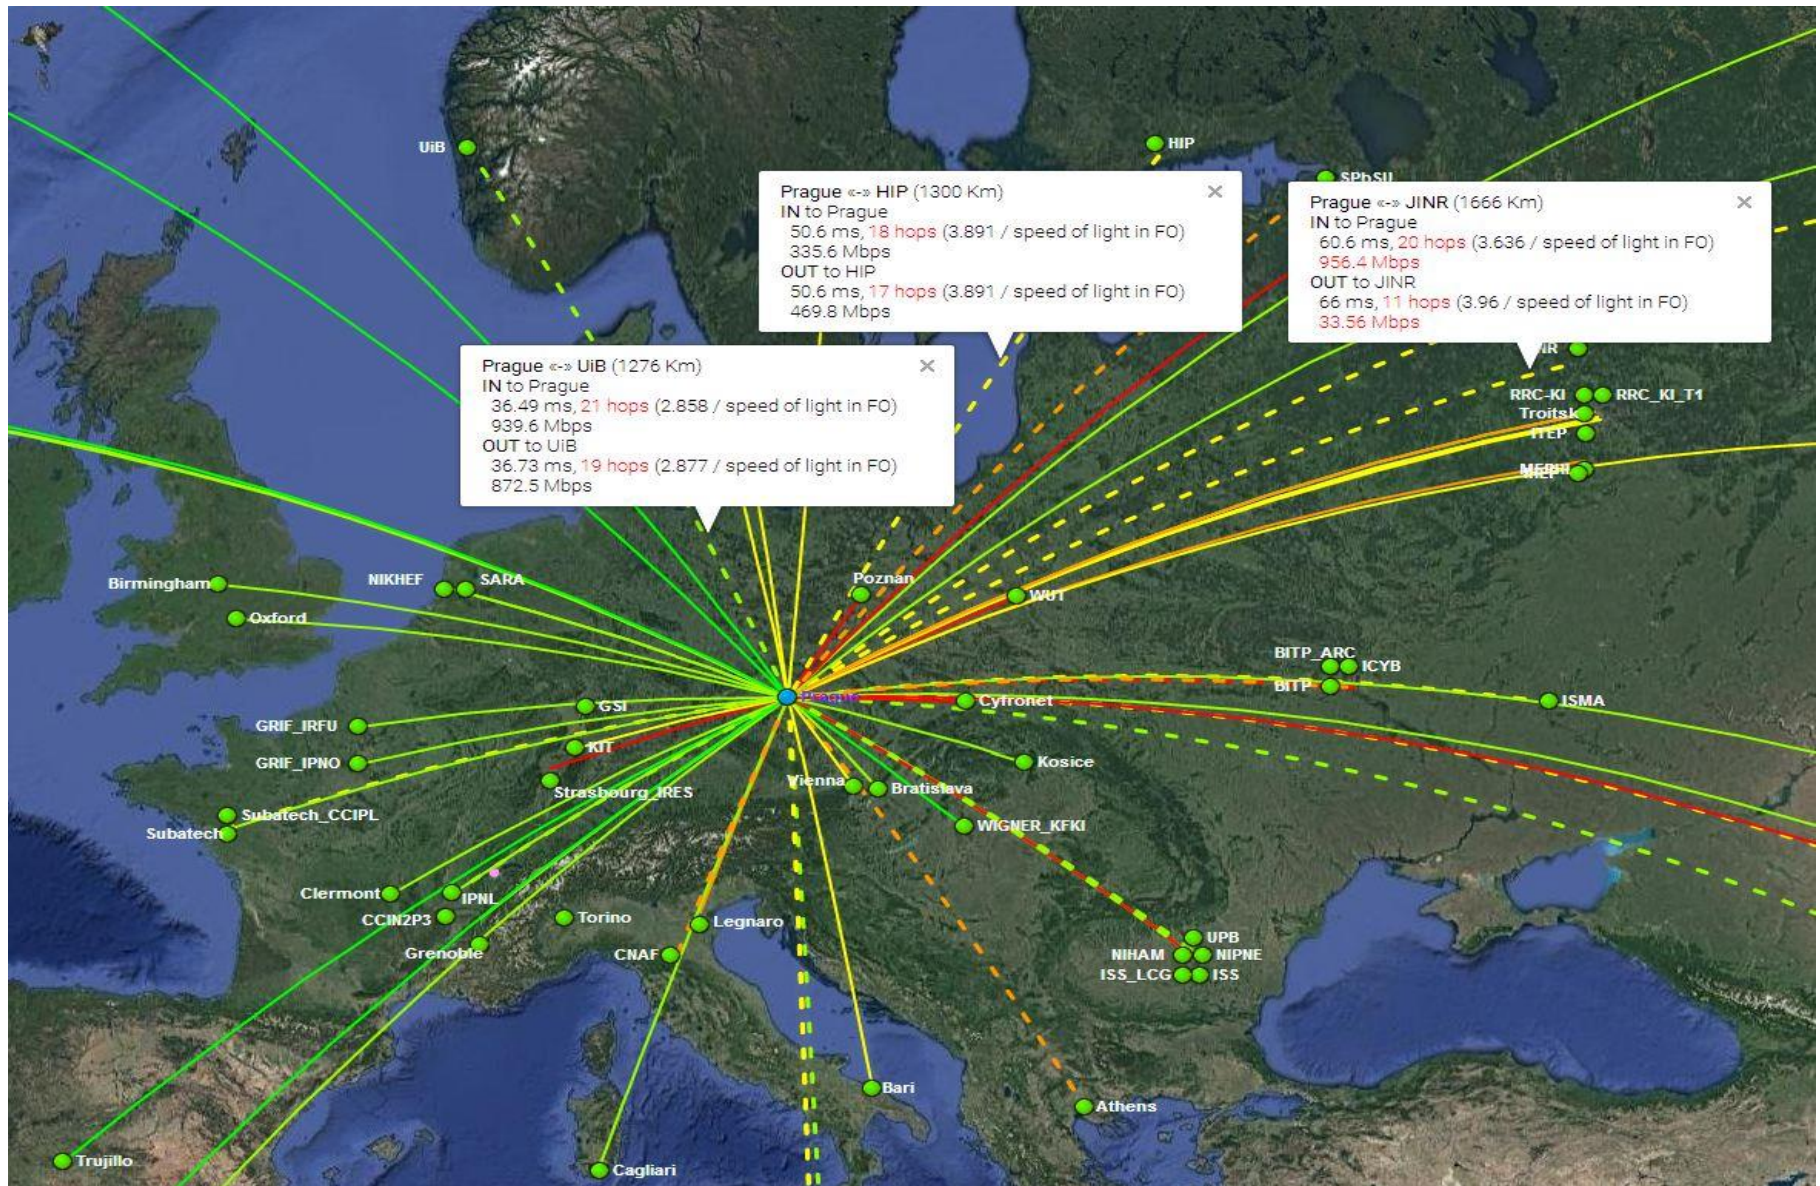
Role of tiers

- T0/T1/T2s – only differentiation is RAW data storage and processing (T0/T1s)
- All tiers with storage do MC and analysis
- Diskless T2s can do only MC, all data is exported to close T0/T1s/T2s

Role of network monitoring

- ALICE is using in-house measured network map
 - Through ***all-to-all*** VO-box periodic [bandwidth tests](#)
- Used to determine the best placement for data from every client (Storage Auto-discovery) in real time
 - Based on network topology
 - Traceroute/tracepath for RTT (bidirectional)
 - 1 TCP stream bandwidth tests (like jobs do)
 - => Distance metric function of the above measurements plus storage availability and occupancy levels

Network paths



Network tuning

- Comprehensive WAN network picture is essential for efficient data transfers and placement
- ALICE depends on and strongly encourages the full implementation of LHCONE at all sites
 - Plus the associated tool (like PerfSONAR), properly instrumented to be used in the individual Grid frameworks
- Huge progress (Thank you network experts!)
 - Number of sites in well-connected parts of the world still have to join
 - Entire regions with high local growth of computing resources, *in particular Asia*, require a lot of further work