NAF Status Report

Untertitel der Präsentation

Yves Kemp
NUC, 11.07.2018

HELMHOLTZ SPITZENFORSCHUNG FÜR GROSSE HERAUSFORDERUNGEN

DESY.

Decomissioning of SGE/BIRD // Status of HTCondor WGS

Summary: SGE/BIRD is history!

- No user complaints
- Still need to migrate (=reinstall) some WGS from SGE/BIRD to HTCondor/BIRD
  - Very few ones

HTCondor WGS

- Every experiment should now have as many HTCondor WGS as before SGE/BIRD WGS
- We streamlined the naming scheme
  - No more nafhh-... , only naf-...
  - We had to rename the WGS
  - ATLAS and CMS have load-balanced aliases. This is not rocket science. If others also want that, please ask

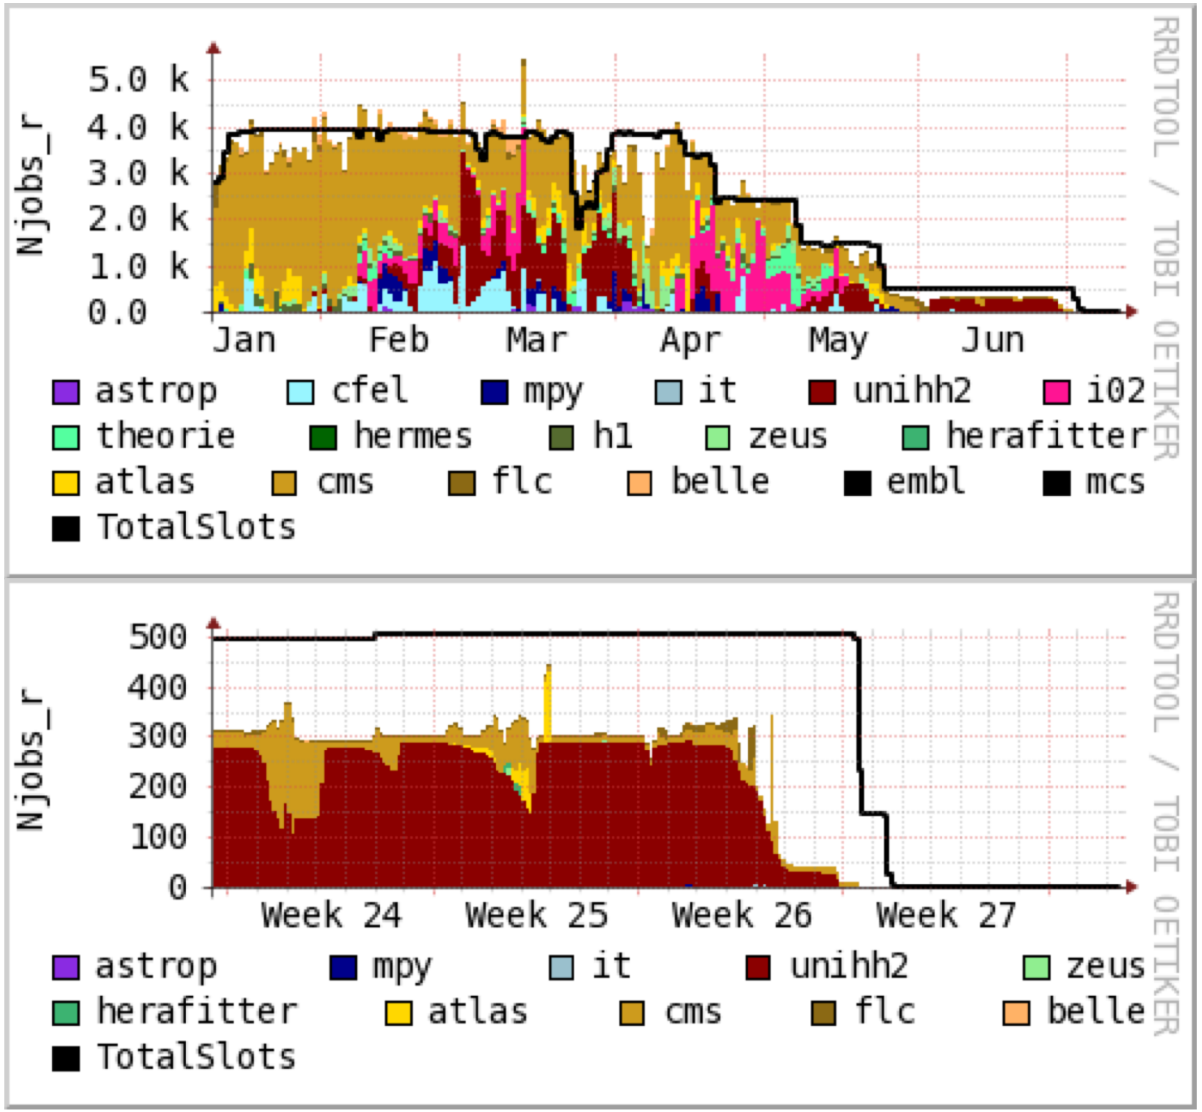

Status of HTCondor Batch

- Now 9320 cores !
- Generally stable running.
  - AFS & Token issues are gone
  - DUST NFS is stable
- Issues observed last month:
  - Small NFS group server are causing troubles. In contact with one group. Generally migrate to NetApp (or DUST)
  - Migration of Kerberos Servers induced short downtime
  - Short troubles during CMS dCache migration
  - Observing badly behaving CMVFS clients – neet to investigate

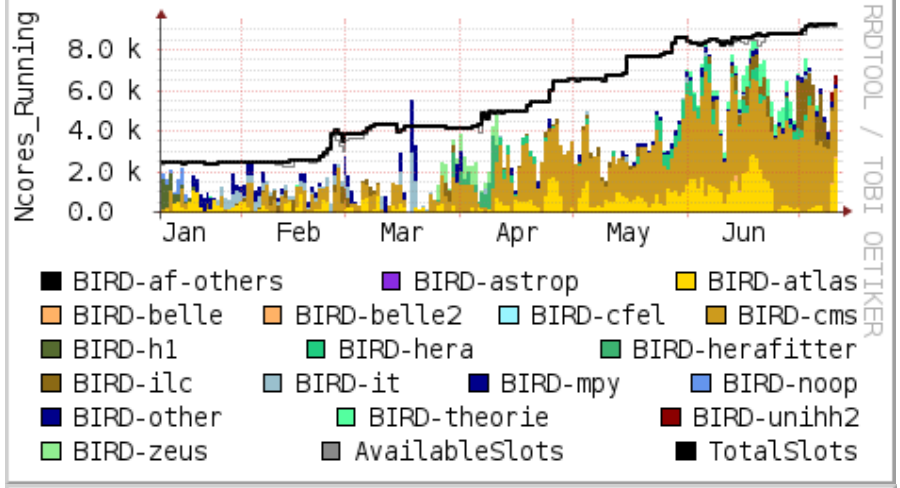

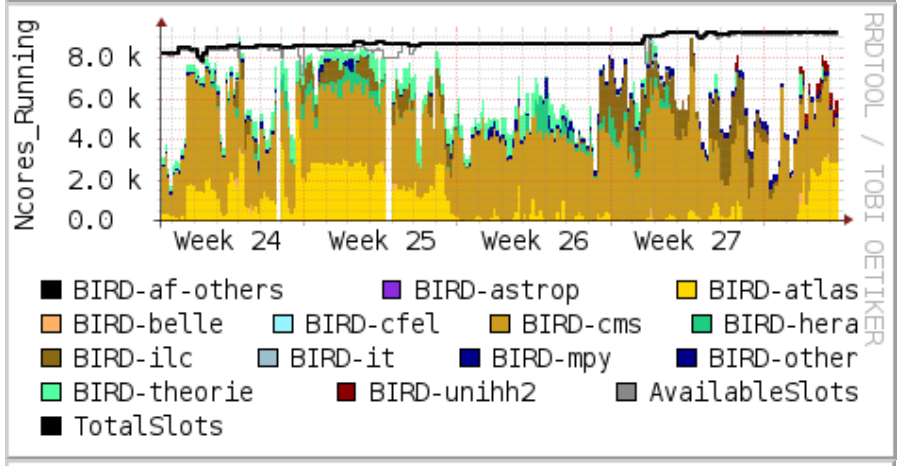

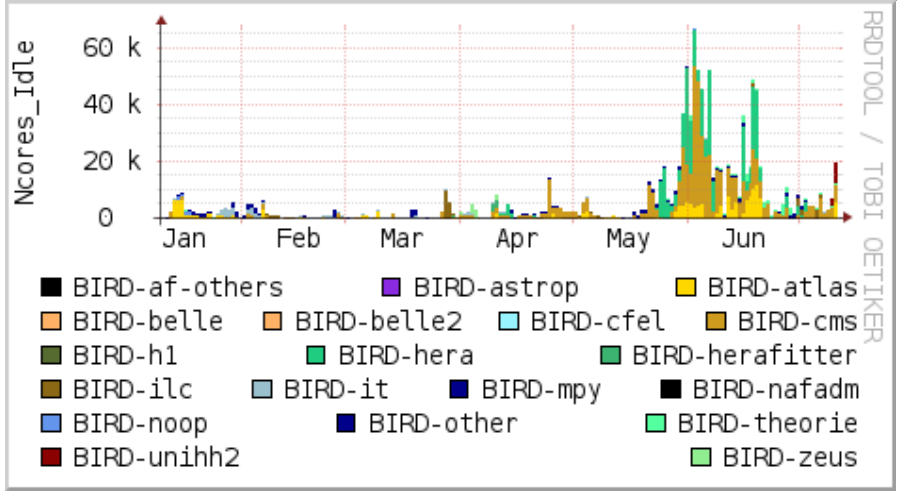

Status of HTCondor Schedulers

- Over factor 10 more jobs/day than Grid ... With Grid having 2x more cores!
- Abundance of very short jobs
- Not a problem per se, but
  - Need to scale schedulers horizontaly (more schedulers, put them on hardware)
  - Users need to help HTCondor:
    - Using „queue n“ for submitting n jobs in one submit already helps
    - Using late materialization if n>1000 (or so) helps in addition

DUST status

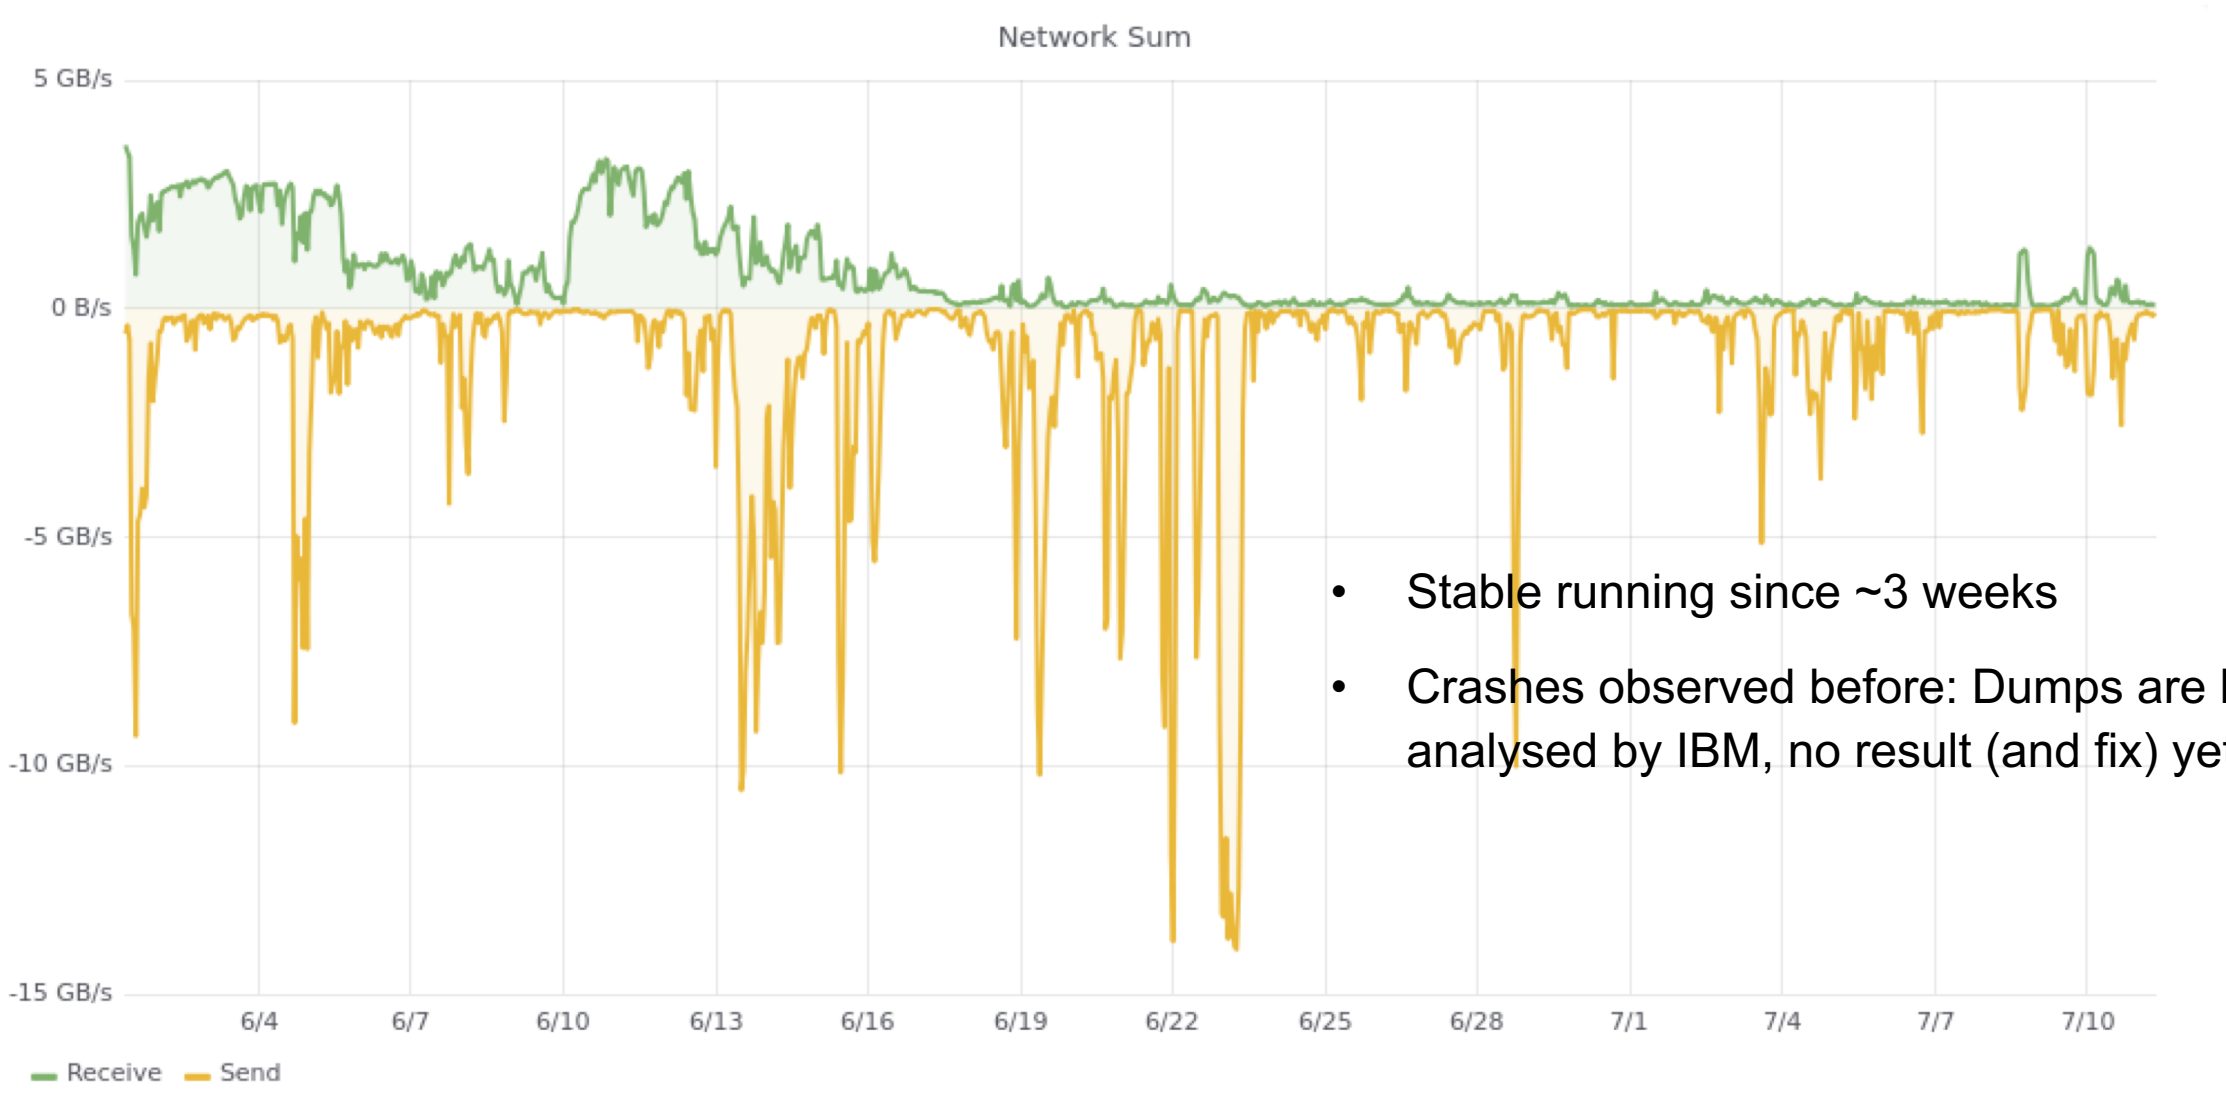

- Stable running since ~3 weeks
- Crashes observed before: Dumps are being analysed by IBM, no result (and fix) yet

Documentation

- We had https://naf-wiki.desy.de/Main_Page in the past. Although WIKI, not much used.
  - This kind of Wiki no longer supported
- Migrating to confluence:
  - https://confluence.desy.de/display/IS/Resources
  - Already lots of documentation, most naf-wiki stuff is already migrated
- Communicated entry point for documentation is
  - **http://bird.desy.de/**
  - This is a redirect to confluence

Network

- External connectivity via DFN
- Previous capacity:
  - General: 2x5 Gbit Downlink, 2x10 Gbit Uplink
  - 2x 10 Gbit LHCONE (symmetrisch)
  - 2x 10 Gbit to Zeuthen (symmetrisch)
- Since, 6.7.2018, Change of general access (XWIN):
  - 2x 15 Gbit Downlink, 2x20 Gbit Uplink
  - Some LHC sites (RAL) go via general XWIN network
- LHCONE is still saturated
- Mid-Term-Plan: Go to 2x50 Gbit/s for all external connections
  - Shared between XWIN/LHCONE/Zeuthen

DFN WAN

Last 2 days

RT-WAN-01 - LHCONE

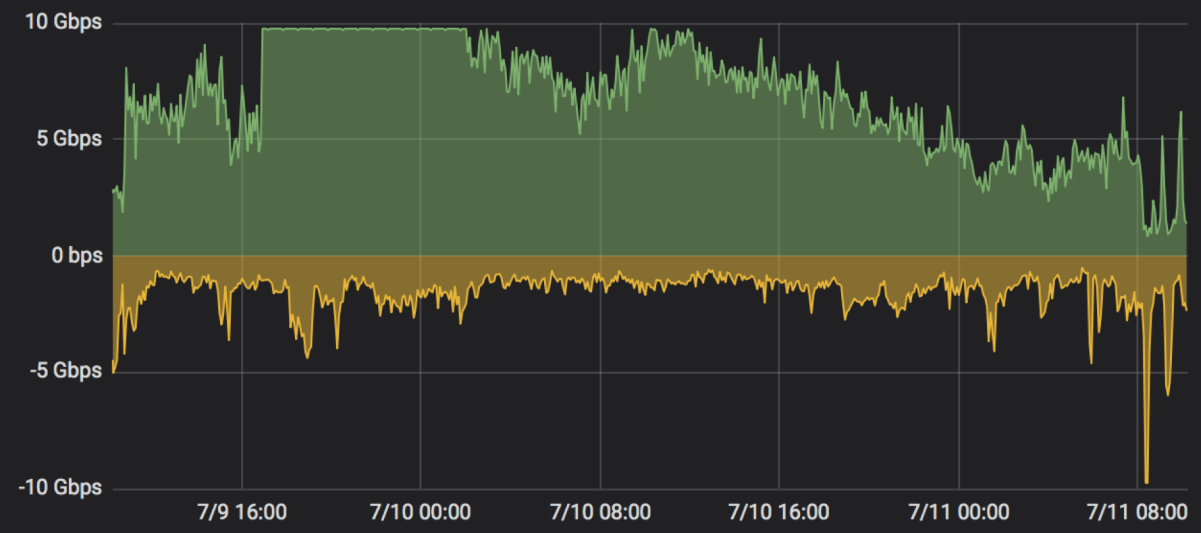

LHCOne → DESY (via rt-wan-01) Max: 9.77 Gbps Avg: 6.83 Gbps
DESY → LHCOne (via rt-wan-01) Max: 9.76 Gbps Avg: 1.57 Gbps

RT-WAN-02 - LHCONE

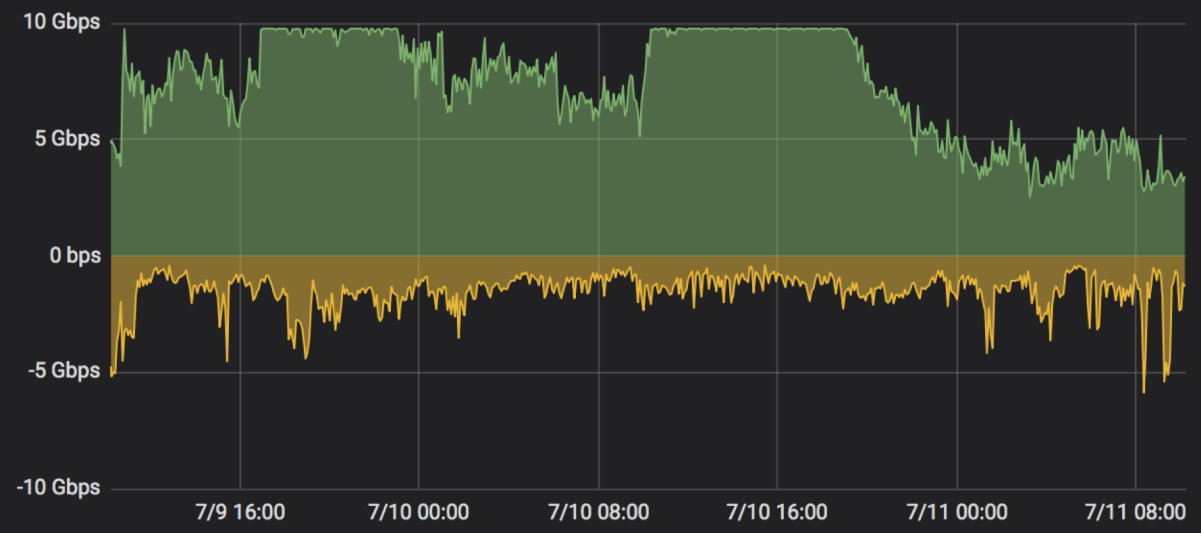

LHCOne → DESY (via rt-wan-02) Max: 9.76 Gbps Avg: 7.33 Gbps
DESY → LHCOne (via rt-wan-02) Max: 5.88 Gbps Avg: 1.52 Gbps

RT-WAN-01 - X-Win

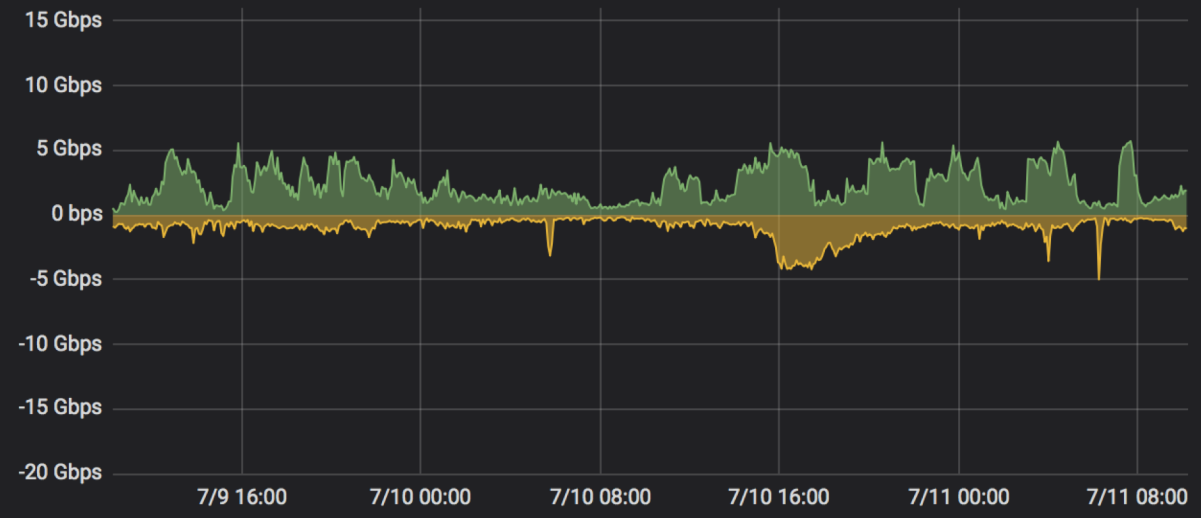

WAN → DESY (via rt-wan-01) Max: 5.7 Gbps Avg: 2.3 Gbps
DESY → WAN (via rt-wan-01) Max: 5.0 Gbps Avg: 968 Mbps

RT-WAN-02 - X-Win

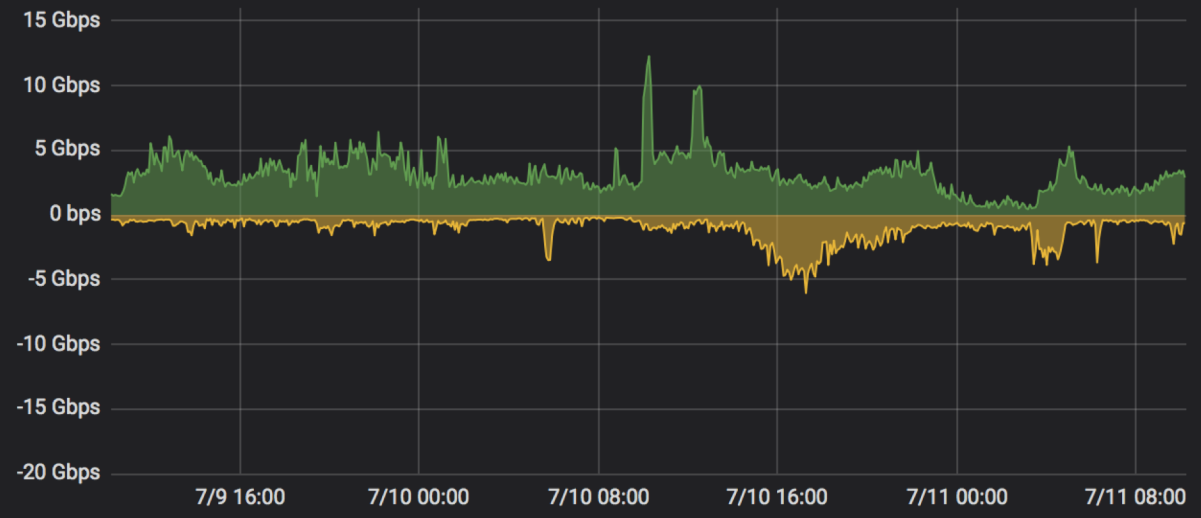

WAN → DESY (via rt-wan-02) Max: 12.3 Gbps Avg: 3.1 Gbps
DESY → WAN (via rt-wan-02) Max: 6.1 Gbps Avg: 1.0 Gbps