Remote access for CMS@GRIF

A. Sartirana

Introduction.

- Since Run 1 **CMS@GRIF** relies on **remote files access** between GRIF subsites
  - i.e. LAL/LPNHE WN's access LLR/IRFU storage;
  - best matchmaking between CMS (original) policies and the GRIF distributed nature;
- …this was long before we all started doing remote access in one way or another…
  - **CMS model** was strictly **"data driven"**;
  - possible thanks to the **robust GRIF NW backbone;**
- …and today can be seen as a pioneering experience
  - massive remote access through a >~10Gb/s WAN NW is what we (want to) do to EU/WW scale today.

The historical view.

Why and how we ended up doing this?

- **4 subsites** ready to **contribute to CMS** computing…
  - IRFU, LAL, LPNHE, LLR;
- …**LLR/IRFU** hosting active **CMS groups**
  - ready to provide storage (beyond best effort);
  - none of the two wanted to bear alone the whole CMS GRIF storage;
- **CMS** was **not able to embed** sites with **more than 1 SE** and without a simple direct many-to-one relation between CE's and SE
  - the best solution was to **group sites by 2 and couple CE's.**

The setup.

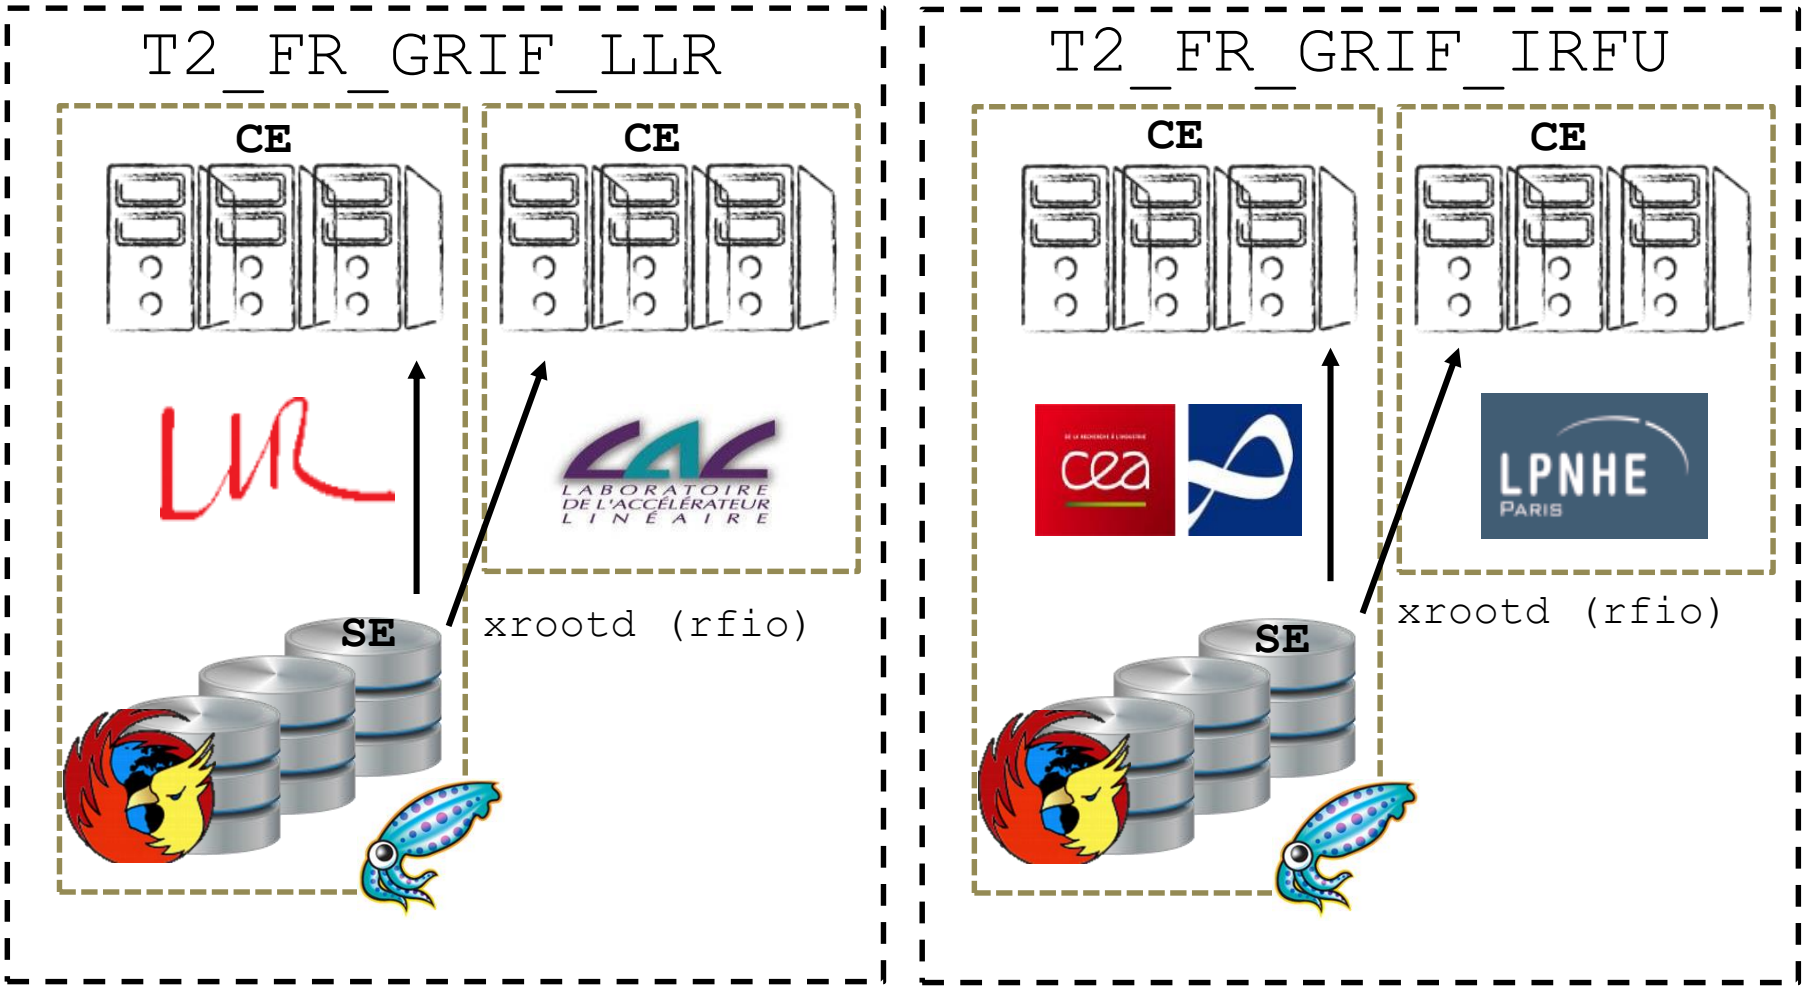
T2_FR_GRIF_LLR
CE
CE
SE
xrootd (rfio)
T2_FR_GRIF_IRFU
CE
CE
SE
xrootd (rfio)

Does it work?

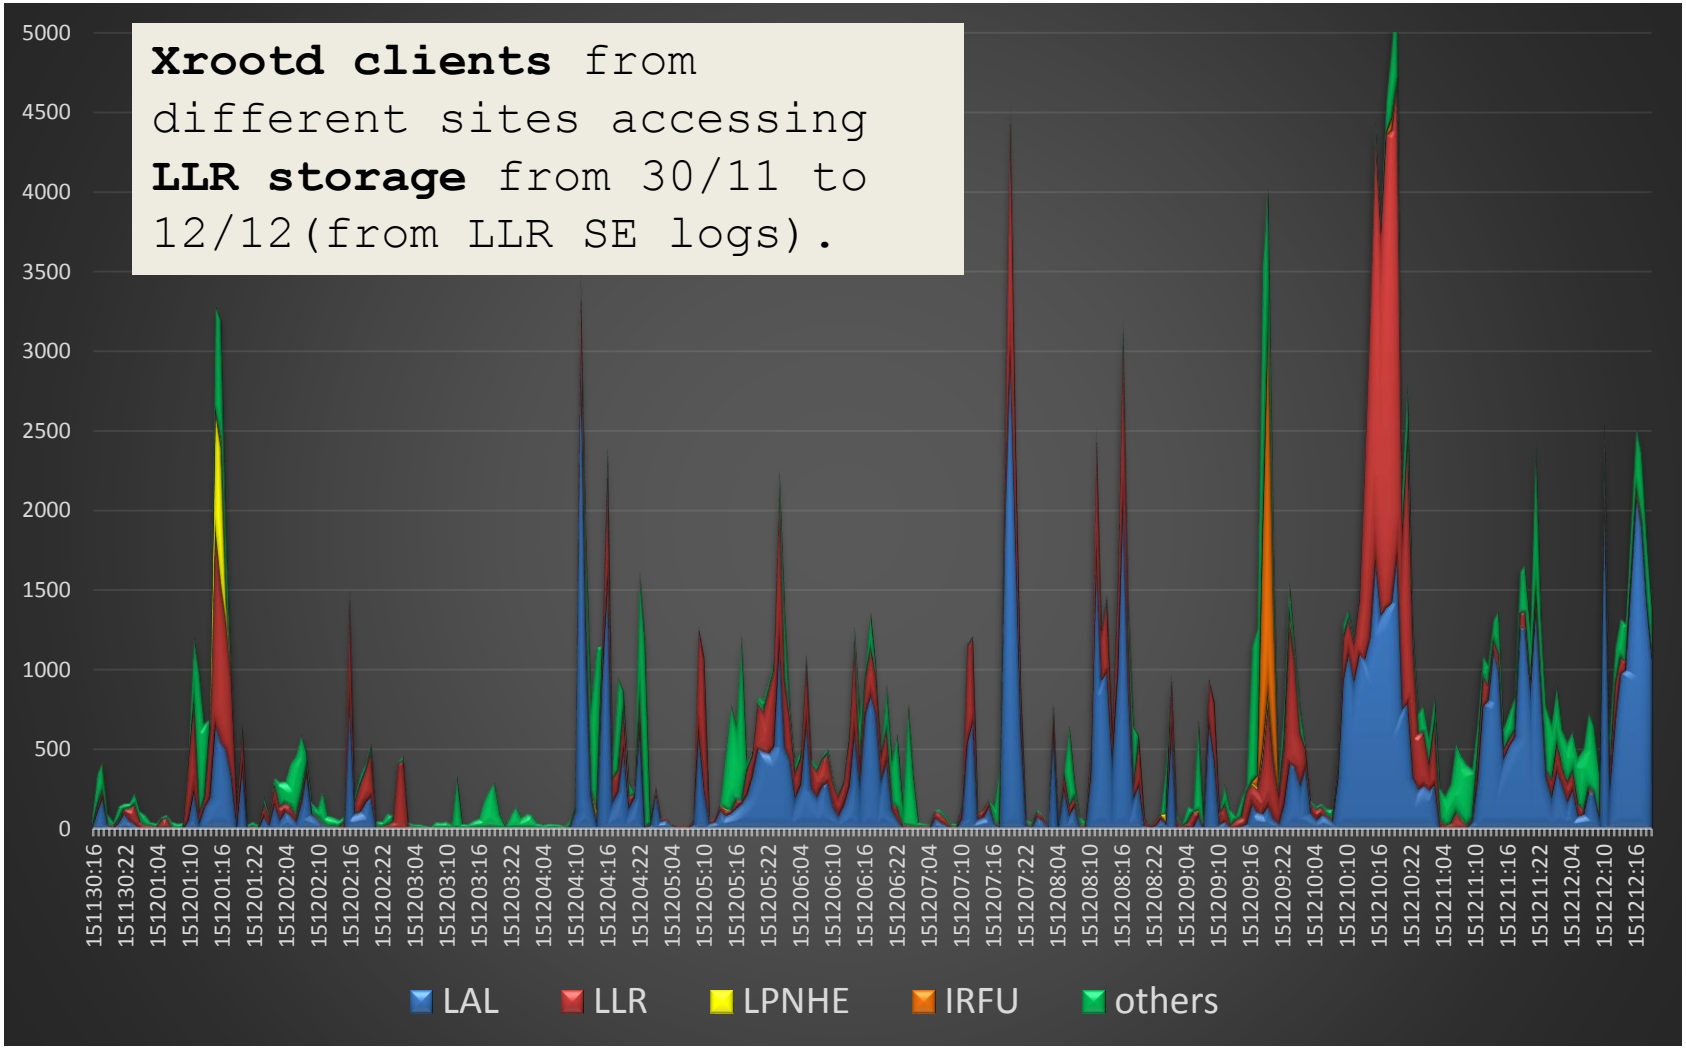
Xrootd clients from different sites accessing LLR storage from 30/11 to 12/12(from LLR SE logs).
5000
4500
4000
3500
3000
2500
2000
1500
1000
500
0
151130:16
151130:22
151201:04
151201:10
151201:16
151201:22
151202:04
151202:10
151202:16
151202:22
151203:04
151203:10
151203:16
151203:22
151204:04
151204:10
151204:16
151204:22
151205:04
151205:10
151205:16
151205:22
151206:04
151206:10
151206:16
151206:22
151207:04
151207:10
151207:16
151207:22
151208:04
151208:10
151208:16
151208:22
151209:04
151209:10
151209:16
151209:22
151210:04
151210:10
151210:16
151210:22
151211:04
151211:10
151211:16
151211:22
151212:04
151212:10
151212:16
LAL
LLR
LPNHE
IRFU
others

Does it work?

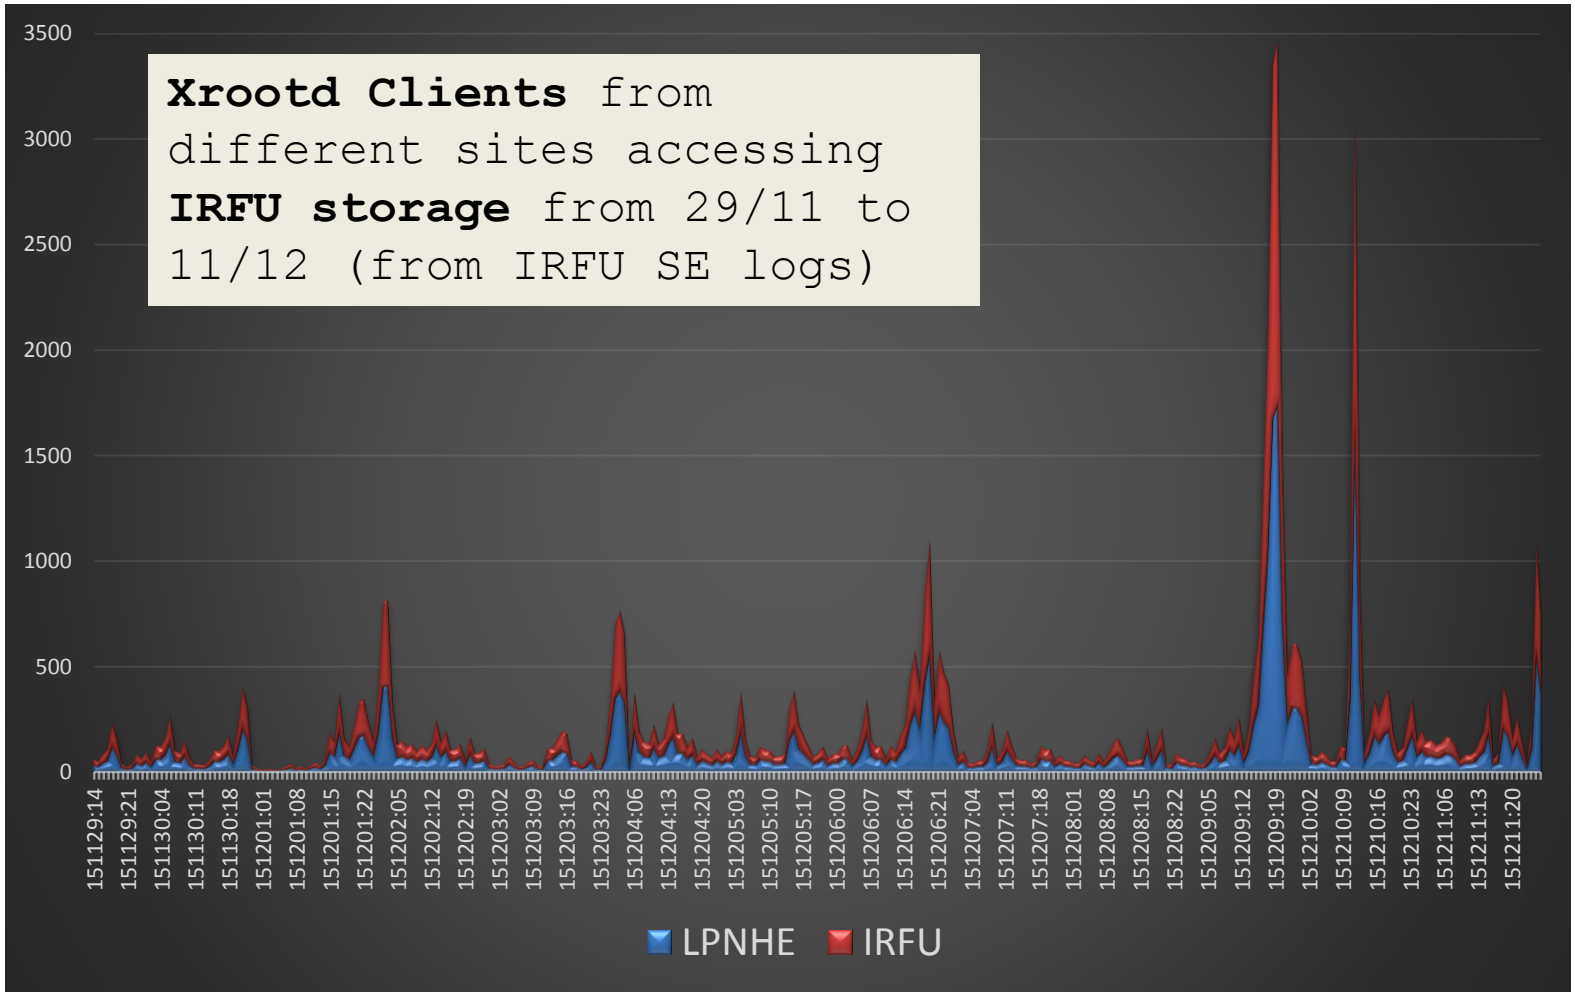
Xrootd Clients from different sites accessing IRFU storage from 29/11 to 11/12 (from IRFU SE logs)
3500
3000
2500
2000
1500
1000
500
0
151129:14
151129:21
151130:04
151130:11
151130:18
151201:01
151201:08
151201:15
151201:22
151202:05
151202:12
151202:19
151203:02
151203:09
151203:16
151203:23
151204:06
151204:13
151204:20
151205:03
151205:10
151205:17
151206:00
151206:07
151206:14
151206:21
151207:04
151207:11
151207:18
151208:01
151208:08
151208:15
151208:22
151209:05
151209:12
151209:19
151210:02
151210:09
151210:16
151210:23
151211:06
151211:13
151211:20
LPNHE
IRFU

Does it work?

This setup based on remote access has **worked quite fine** so far, despite **GRIF** is mostly **focused on** supporting CMS **analysis** and thus has, in general, a **data intensive** type of activity.

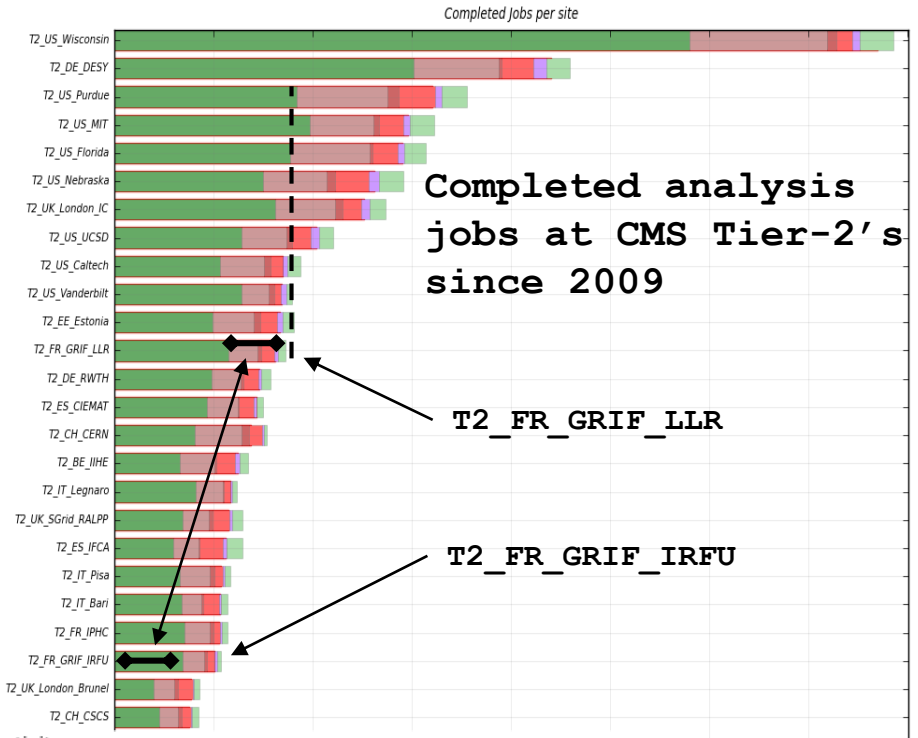

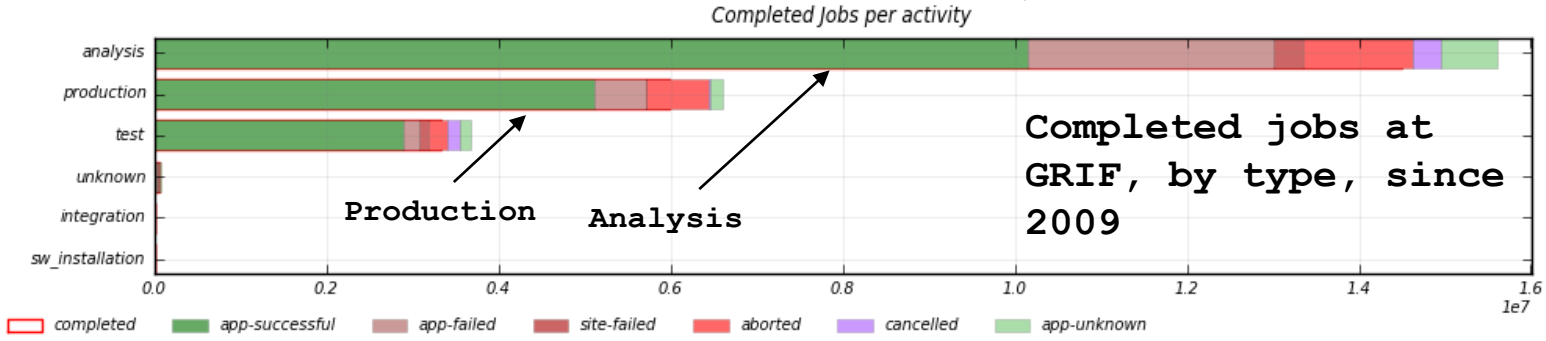

Does it work?

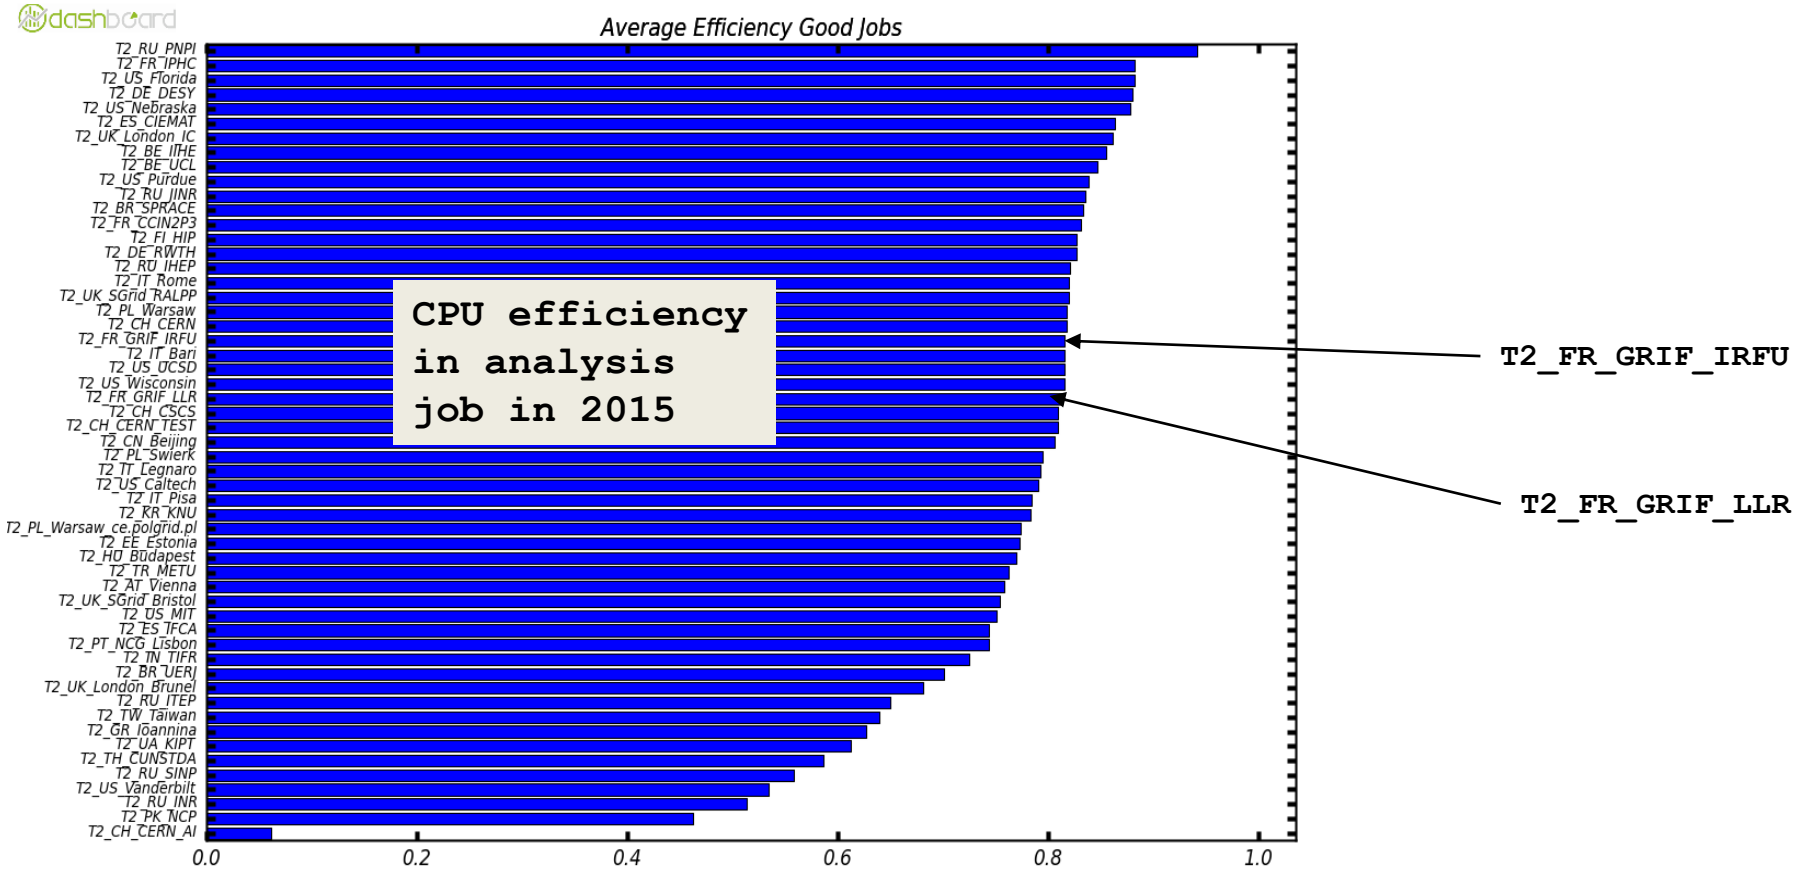

Indeed one can see that remote i/o **does not** seem to **affect GRIF efficiency.**

Impact on network.

So, what kind of **nw traffic** does this type of setup generates?

- **correlating** VO activities/type of access/network traffic, is really **not an easy tas**k (even in the easy LLR case)
  - VO **jobs monitoring** does **not** contain **all the information** and **federation** (AAA) **monitoring** is still **incomplete and unreliable;**
  - **network charts** are a bit confusing (at least to me)and have no real historical view;
  - even in the **local storage** logs we do **not** have **all the info;**
- …ok, anyways, let's try to make an example…

Impact on network.

LLR-GRIF NW gateway traffic (past year)

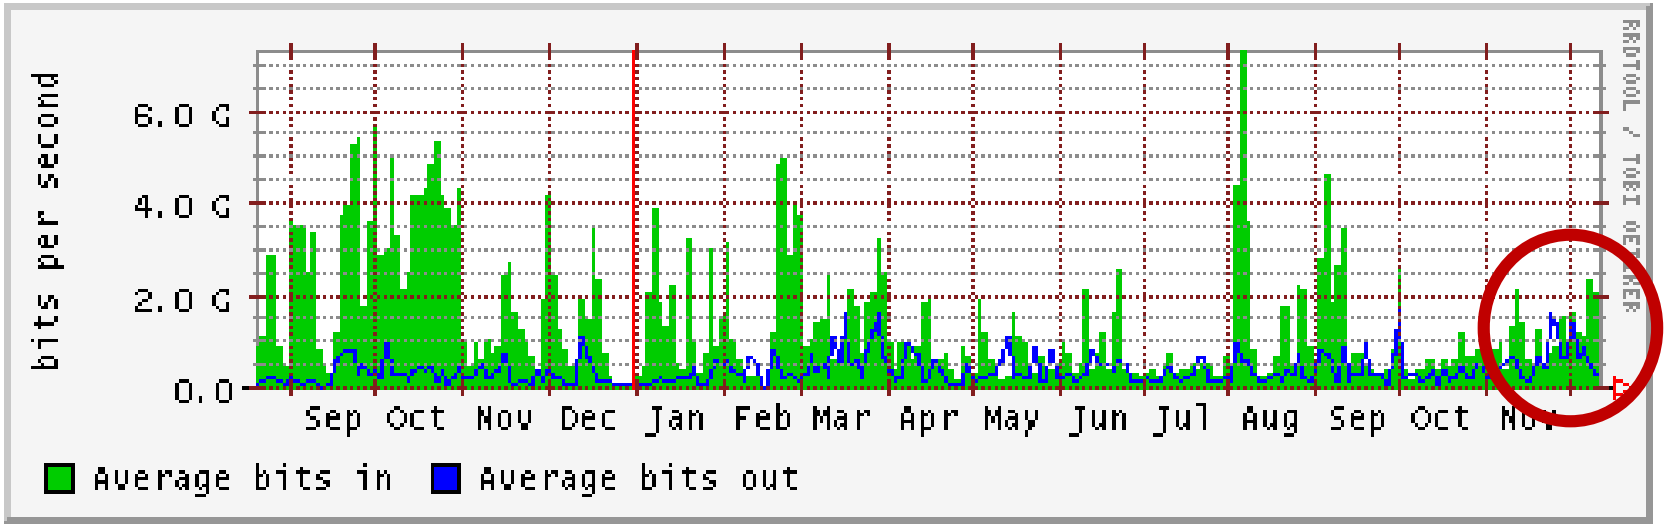

Let's **zoom on the last 2 weeks**…

Impact on network.

LLR-GRIF traffic (1/12 – 11/12)

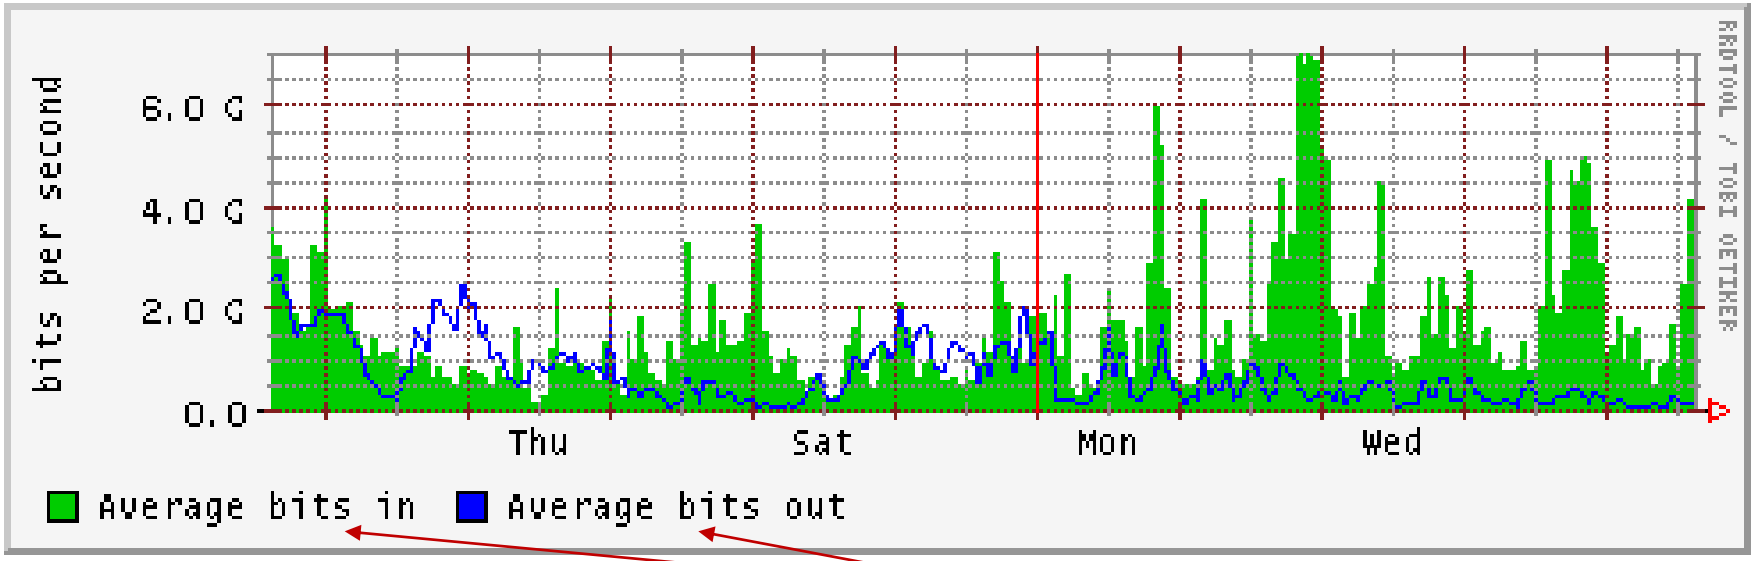

Relatively quiet times but we can see that there are several **traffic peaks**. Let's have a look to what they are.

Impact on network.

LLR-GRIF traffic (1/12 – 11/12)

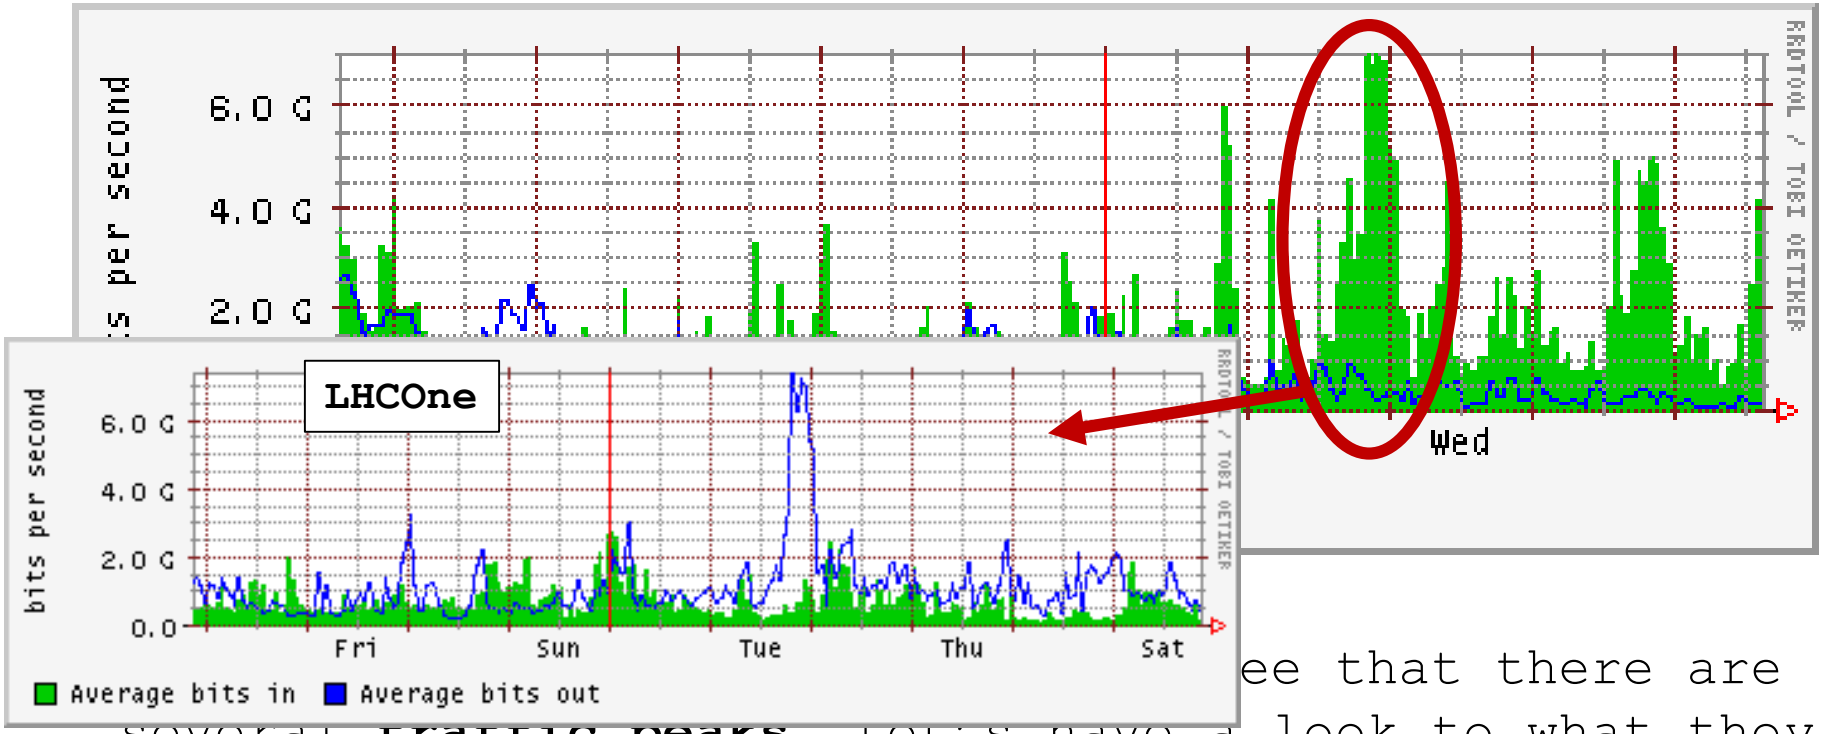

ee that there are several **traffic peaks**. Let's have a look to what they are.

Impact on network.

LLR-GRIF NW gateway traffic (1/12 – 11/12)

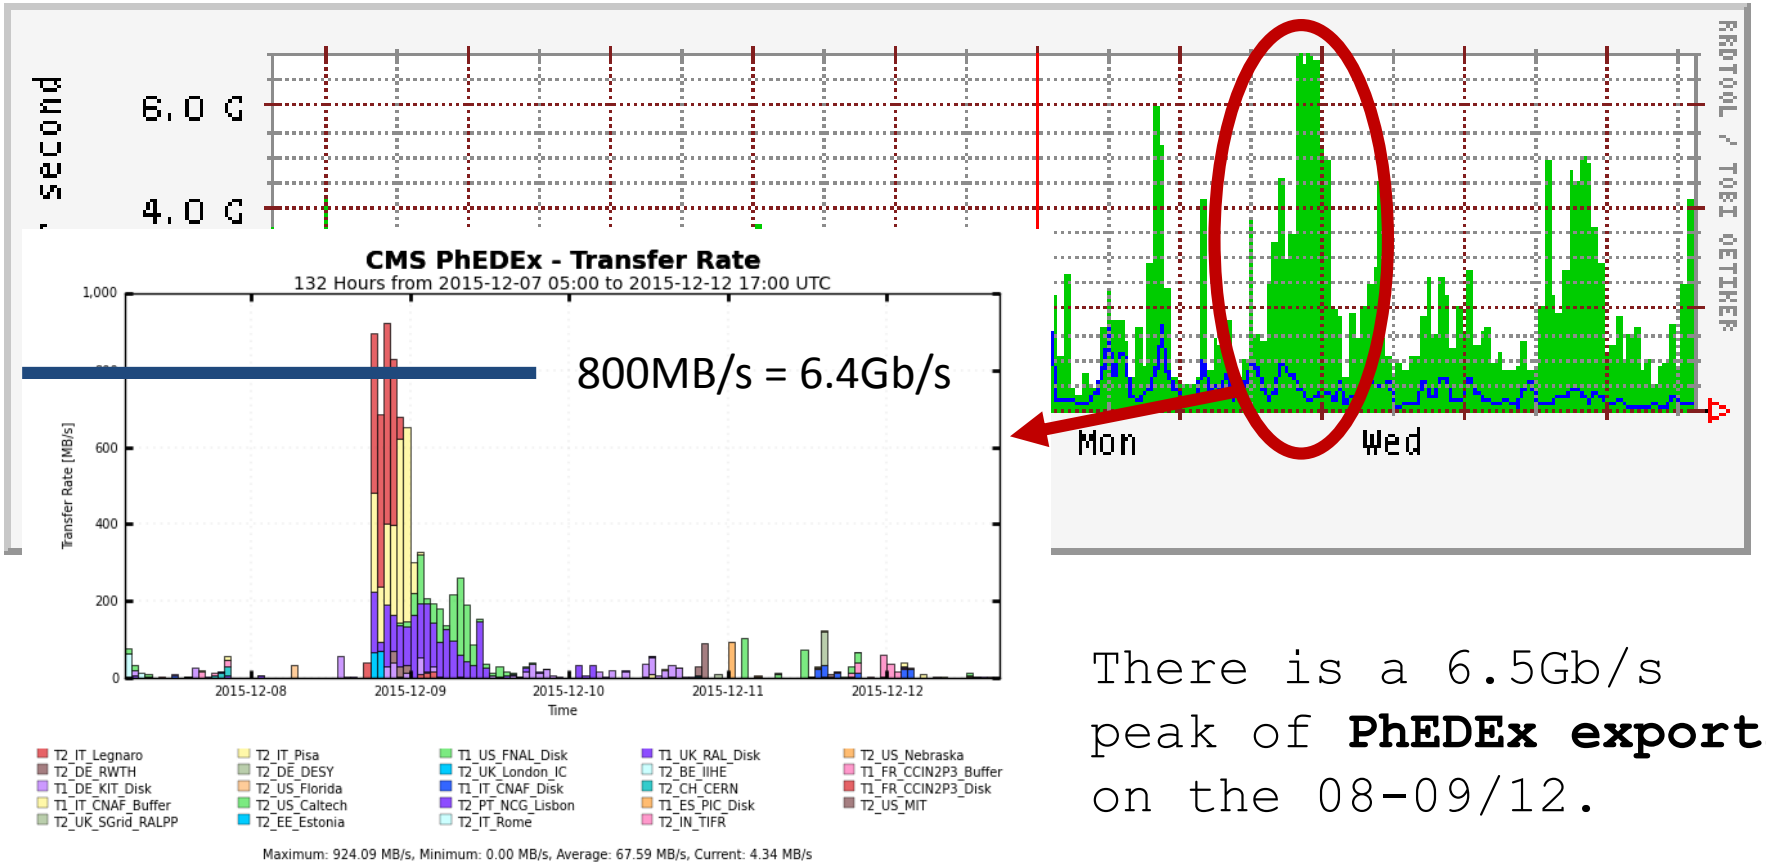

There is a 6.5Gb/s peak of **PhEDEx exports** on the 08-09/12.

Impact on network.

LLR-GRIF NW gateway traffic (1/12 – 11/12)

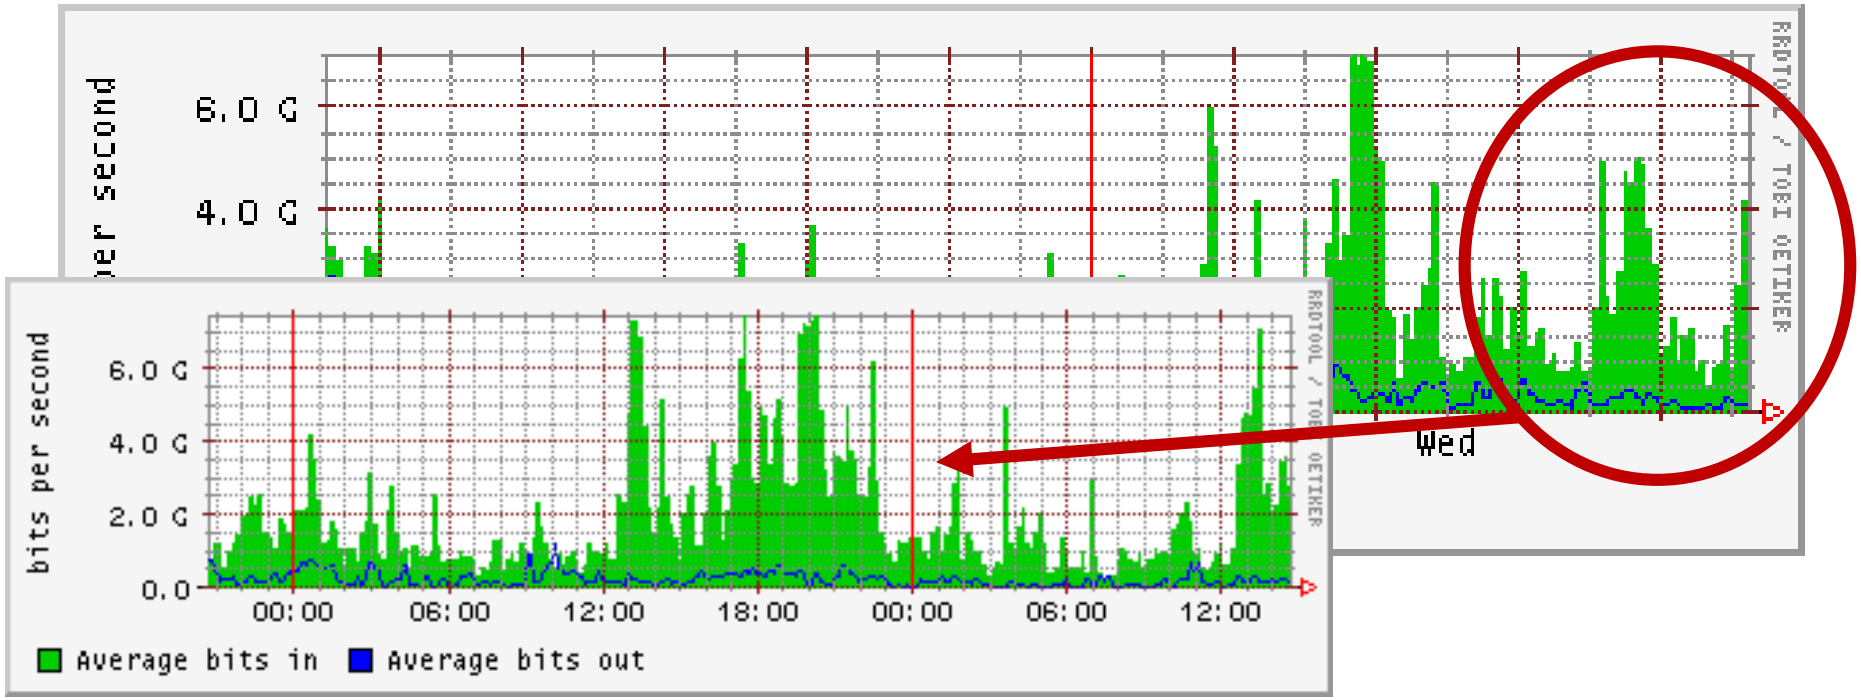

Zooming on the other peak we can try to correlate it with CMS activity.

Impact on network.

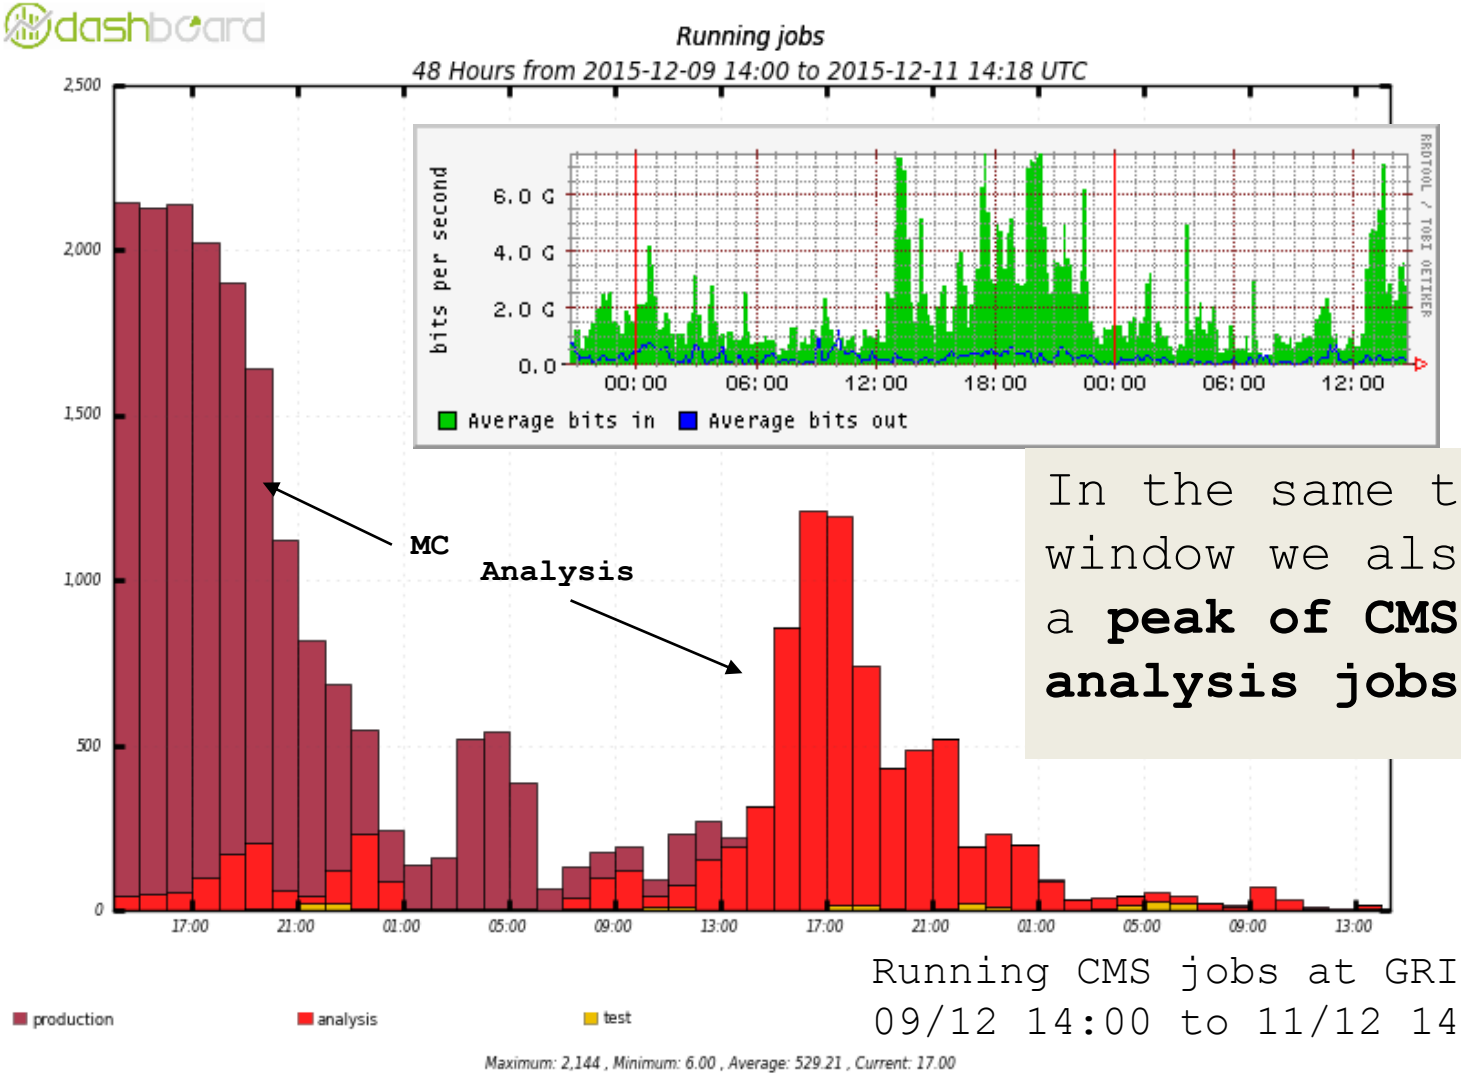

In the same time window we also see a **peak of CMS analysis jobs.**

Running CMS jobs at GRIF from 09/12 14:00 to 11/12 14:00

Impact on network.

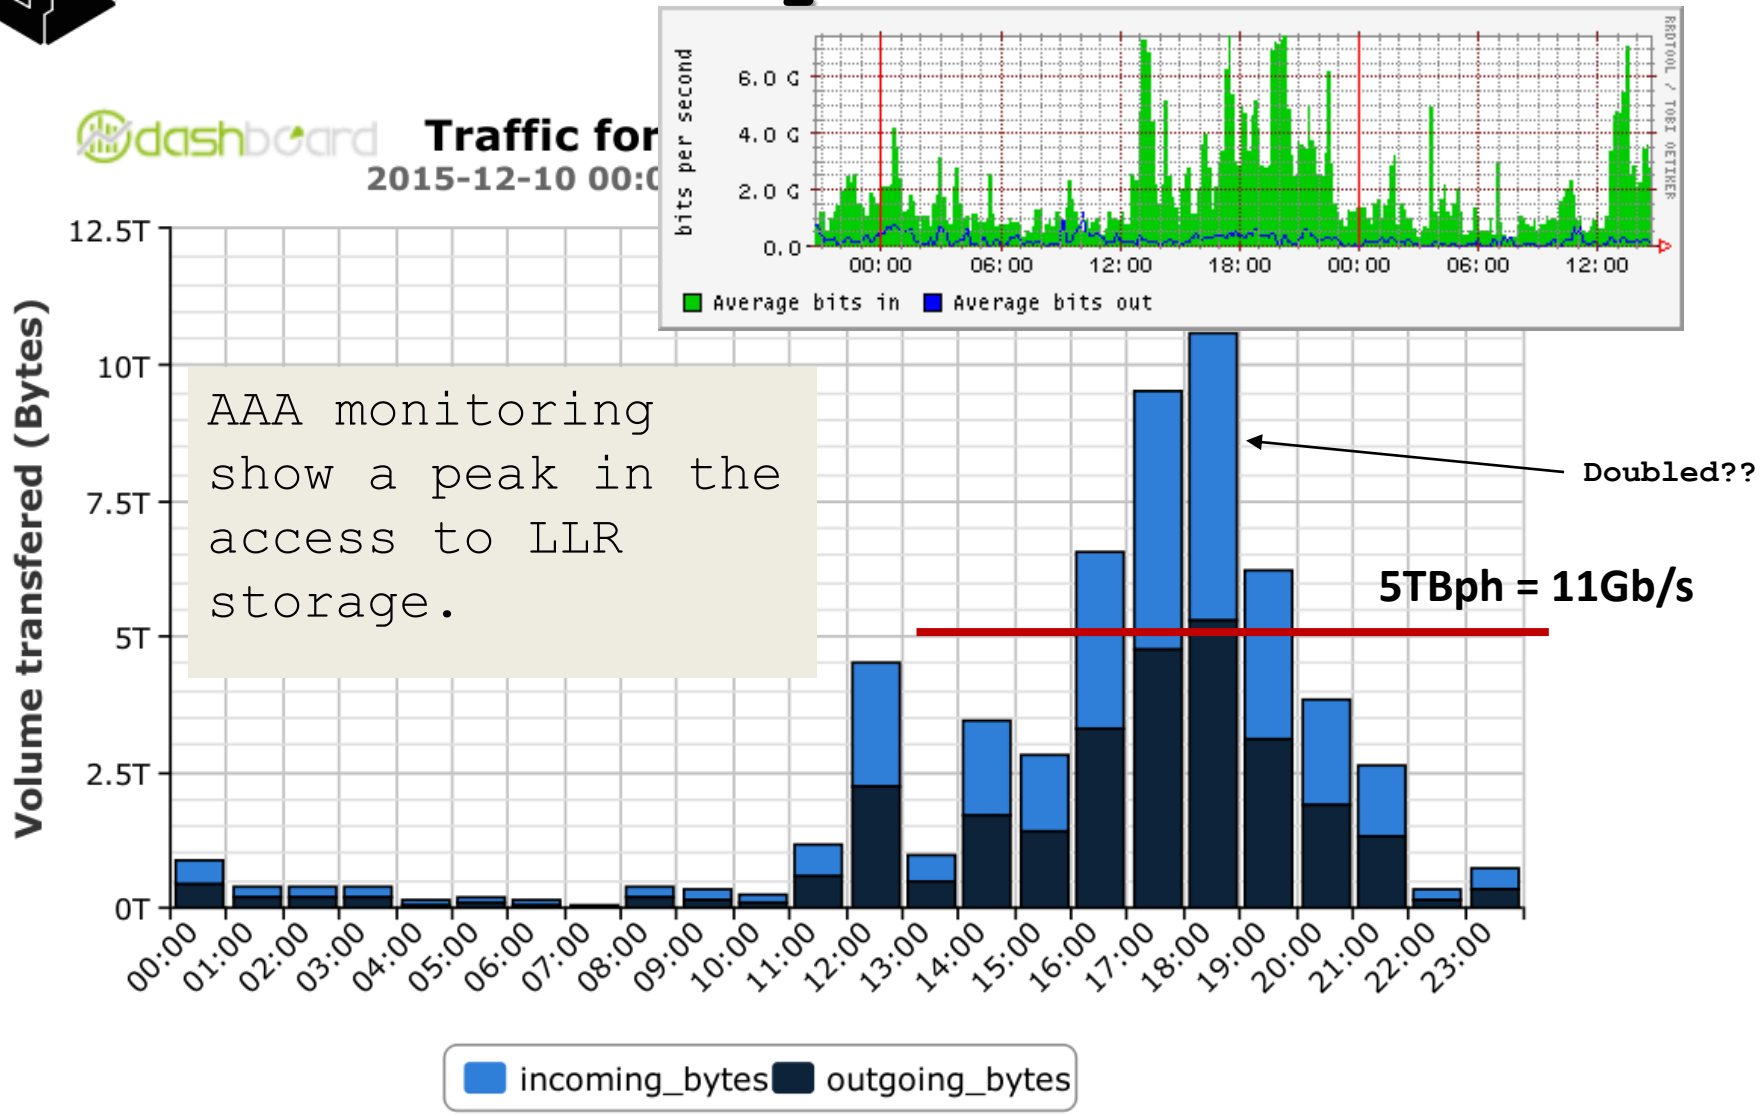

AAA monitoring show a peak in the access to LLR storage.

Impact on network.

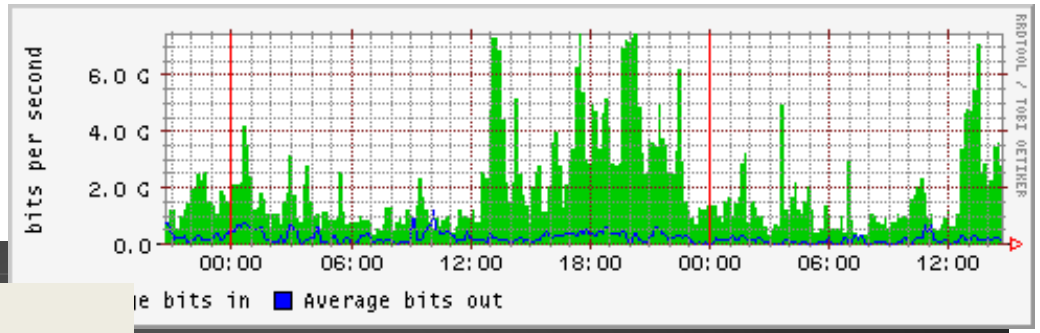

The access logs confirm the peak and tell us that ~40% of accesses come from LAL.

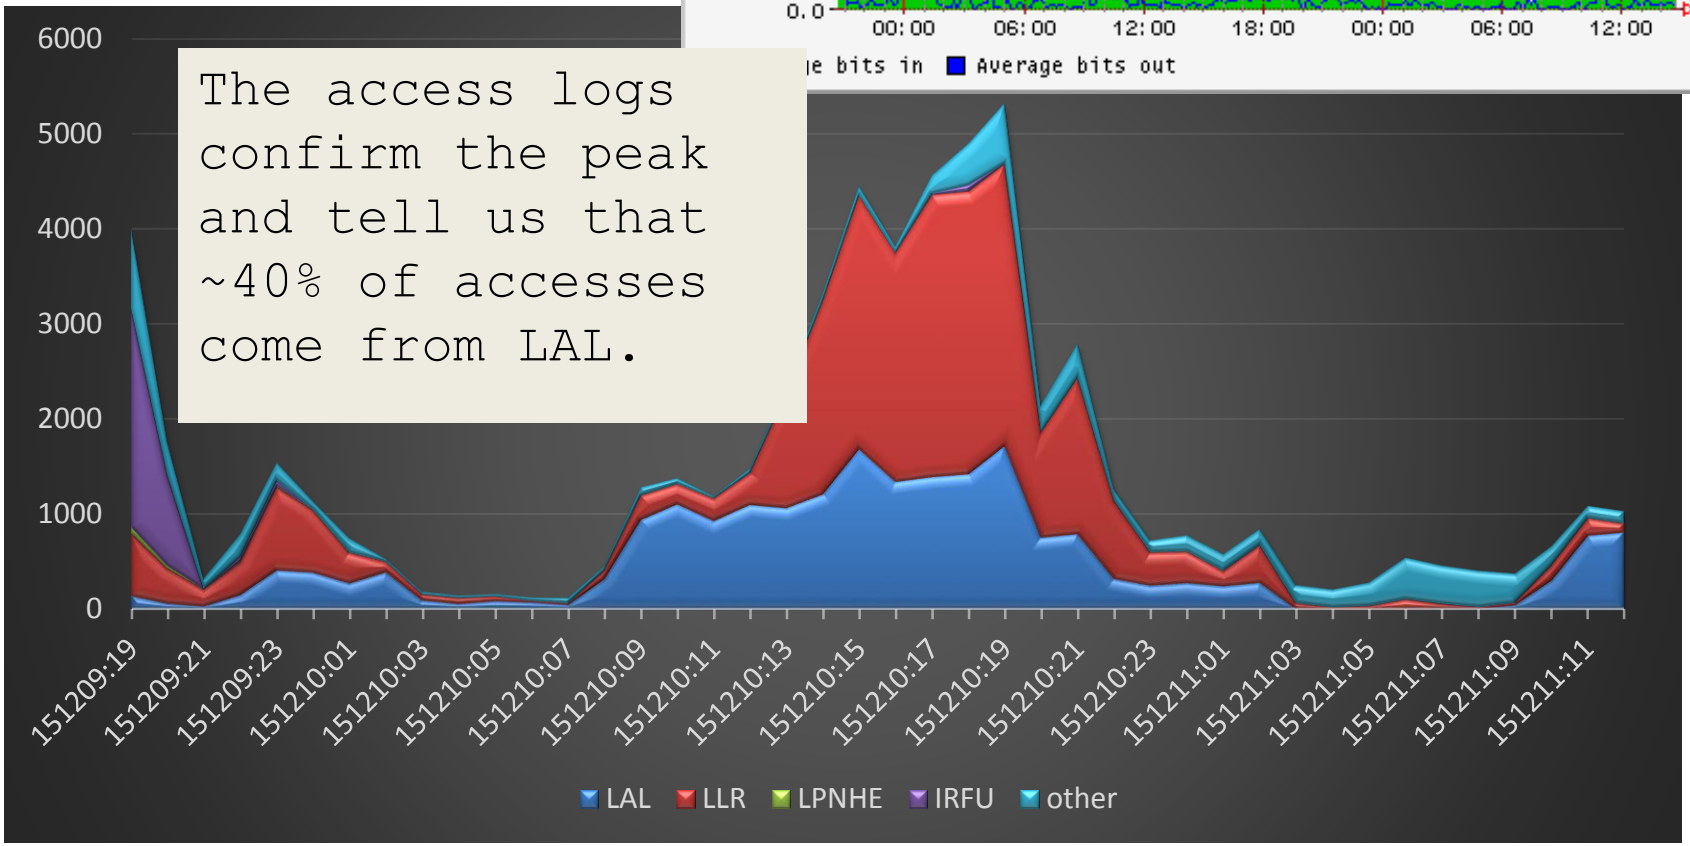

Impact on network.

The monitoring of LAL GW seems broken, but finally Guillaume gave us the smoking gun

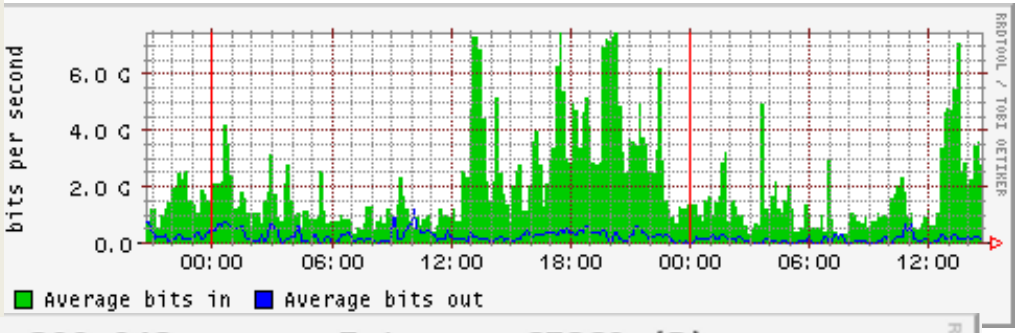

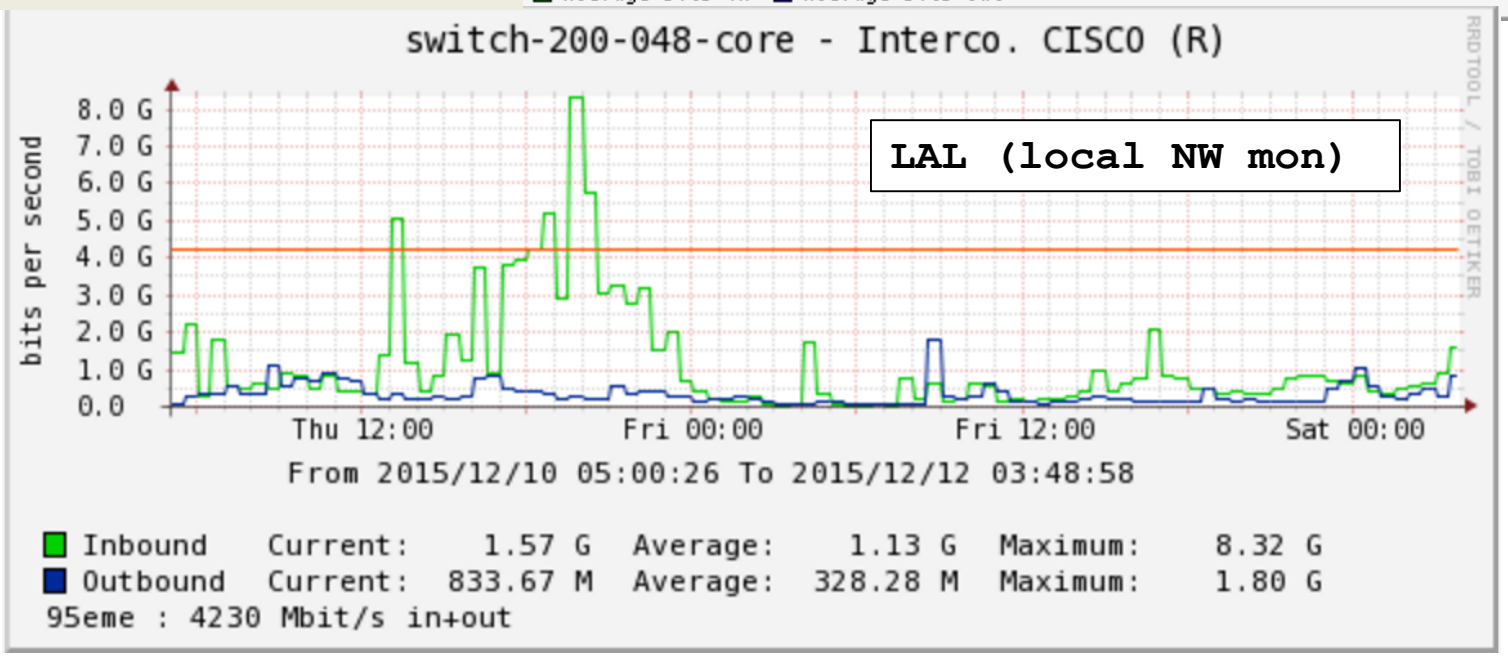

Impact on network.

…we managed to identify a nice ~8Gbps peak generated by remote access.

- this is representative of the type of traffic generated by CMS activity at LLR;
- one reflection:
  - this is a very simple case…
    - an almost-single-VO site: easy to identify activities;
    - few activity peaks: good time resolution;
  - …and we had to…
    - make (painful) searches into 2/3 different monitoring systems;
    - do some log mining and look at site specific nw monitoring (thanks Guillaume);
    - guess (a lot);
    - …and still we have some doubts on what we see;

Summing Up.

- remote access at for CMS at GRIF has been **working fine** since Run 1…
  - no impact on the quality of service;
- … can easily generate **O(10Gb/s)** traffic between the concerned sites…
  - need for at least N x 10Gbps if all (LHC) VO's start using it massively;
- … and can make **activity quite opaque**
  - even in the most simple cases it is difficult to understand what is going on relying on the current monitoring tools;
  - things can only grow worse if more VO's start to frequently use WAN access, unless our monitoring tools evolve accordingly.

Questions?

Impact on network.

An example of something we do not understand…

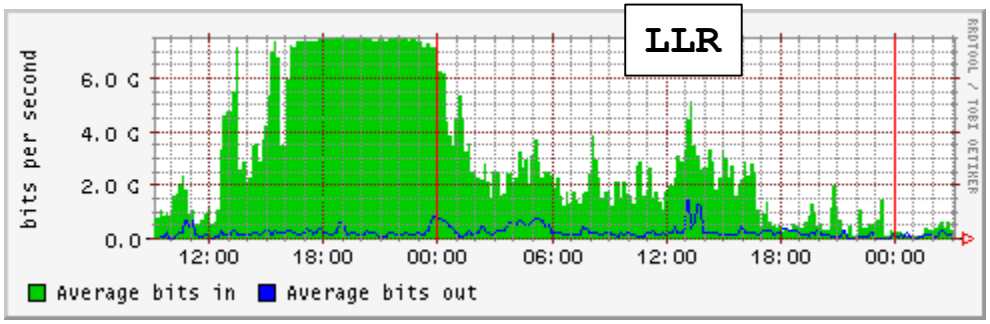

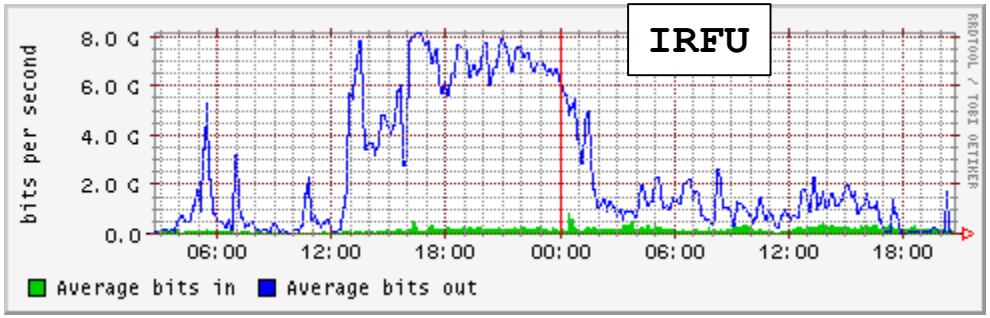

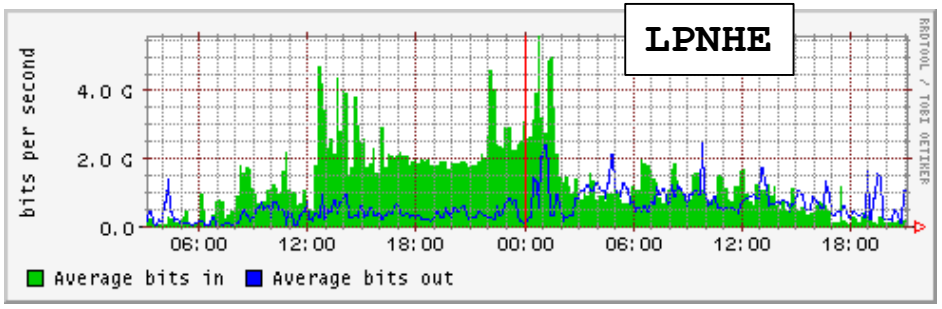

…something that seems an 8Gb/s sustained traffic out of LLR into IRFU on the 11/12…

… but I'm not able to correlate with some CMS (or others) activity.

… and I see a similar peak at LPNHE… but why?…