

# Network Incident Handling

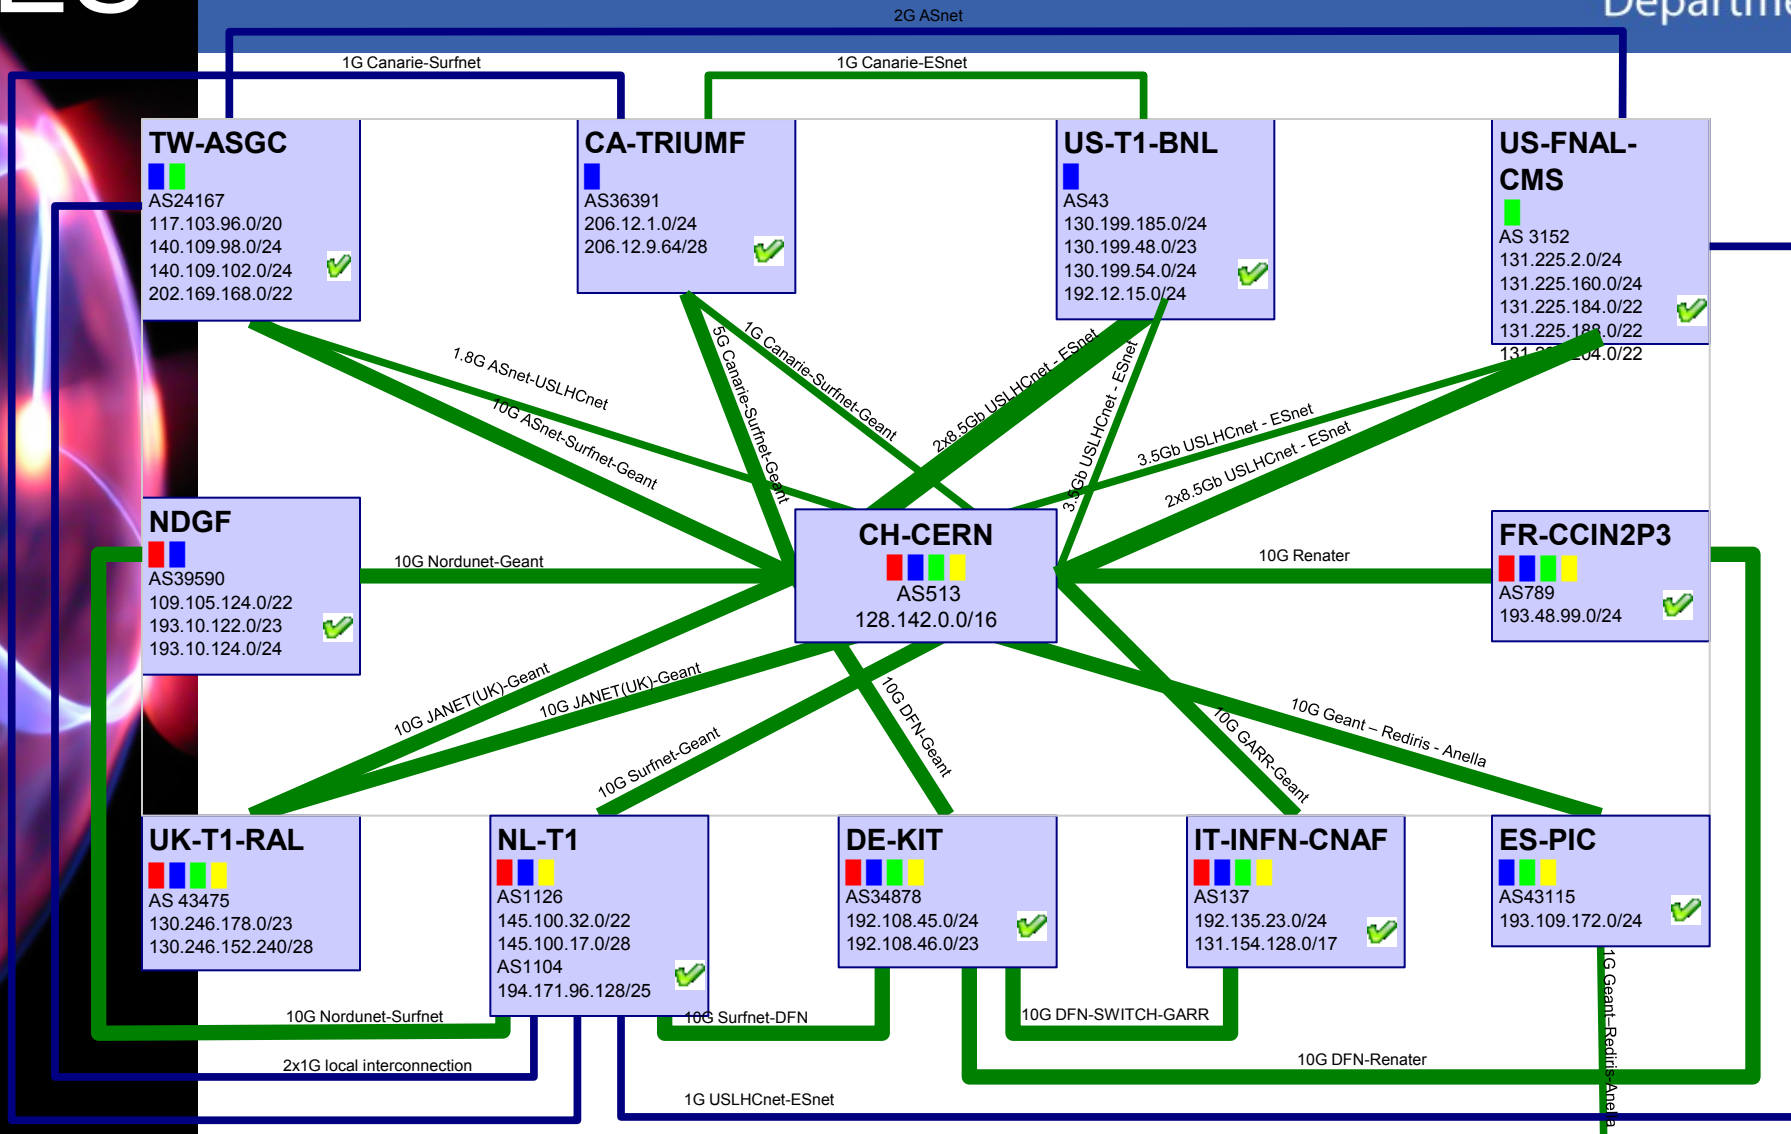
WLCG Grid Deployment Board  
November 2010

Alessandro Di Girolamo, John Shade,  
Jamie Shiers

(Revised 2 December 2010)

- The following is our understanding of how inter-site network incidents should be handled in WLCG
- The procedure is **simple**, involves a **well-defined** set of actors and is applicable to **both** LHCOPN and GPN incidents
- But first, a problem statement

1. Case 1: a “clean” cut of link between site A & B.
    - Traffic automatically rerouted, interruption transparent
    - Failover / back sometimes manual
    - No major service disruption; well understood
  
  2. Case 2: a degradation – lower than expected transfer rates and/or high failure rates
    - These are the cases that sometimes take a long time to diagnose and on which we will focus
- The procedure is applicable to both, but case 1 is not (really) a problem



### CFRN IT Department

- T0-T1 and T1-T1 traffic
- T1-T1 traffic only
- Not deployed yet
- (thick) >=10Gbps
- (thin) <10Gbps

= Alice = Atlas

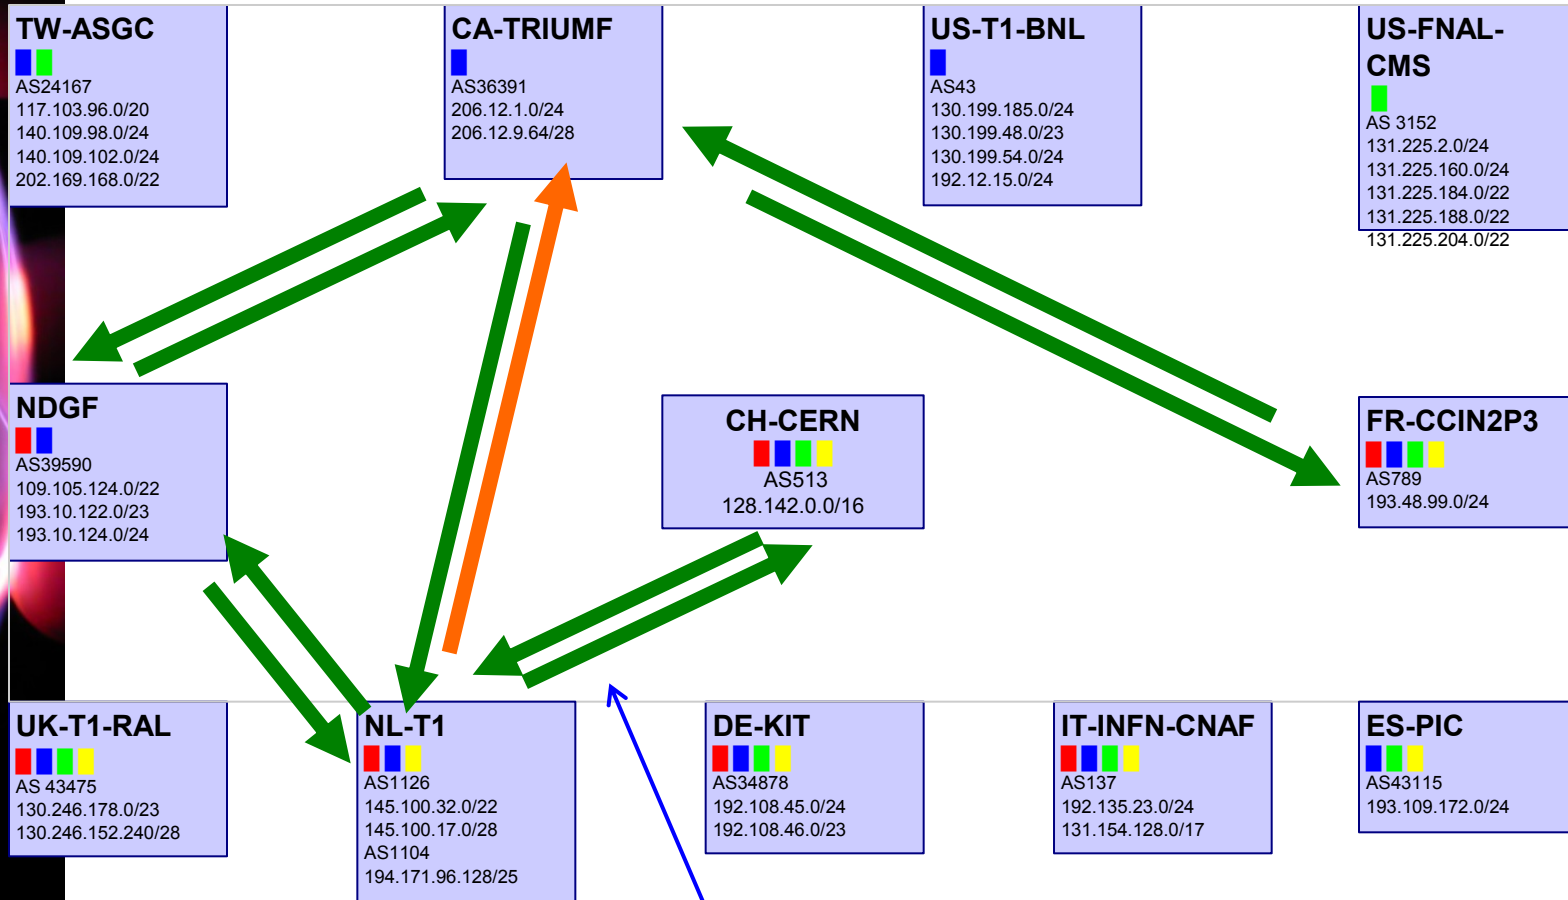
= CMS = LHCB

= internet backup available

p2p prefix: 192.16.166.0/24

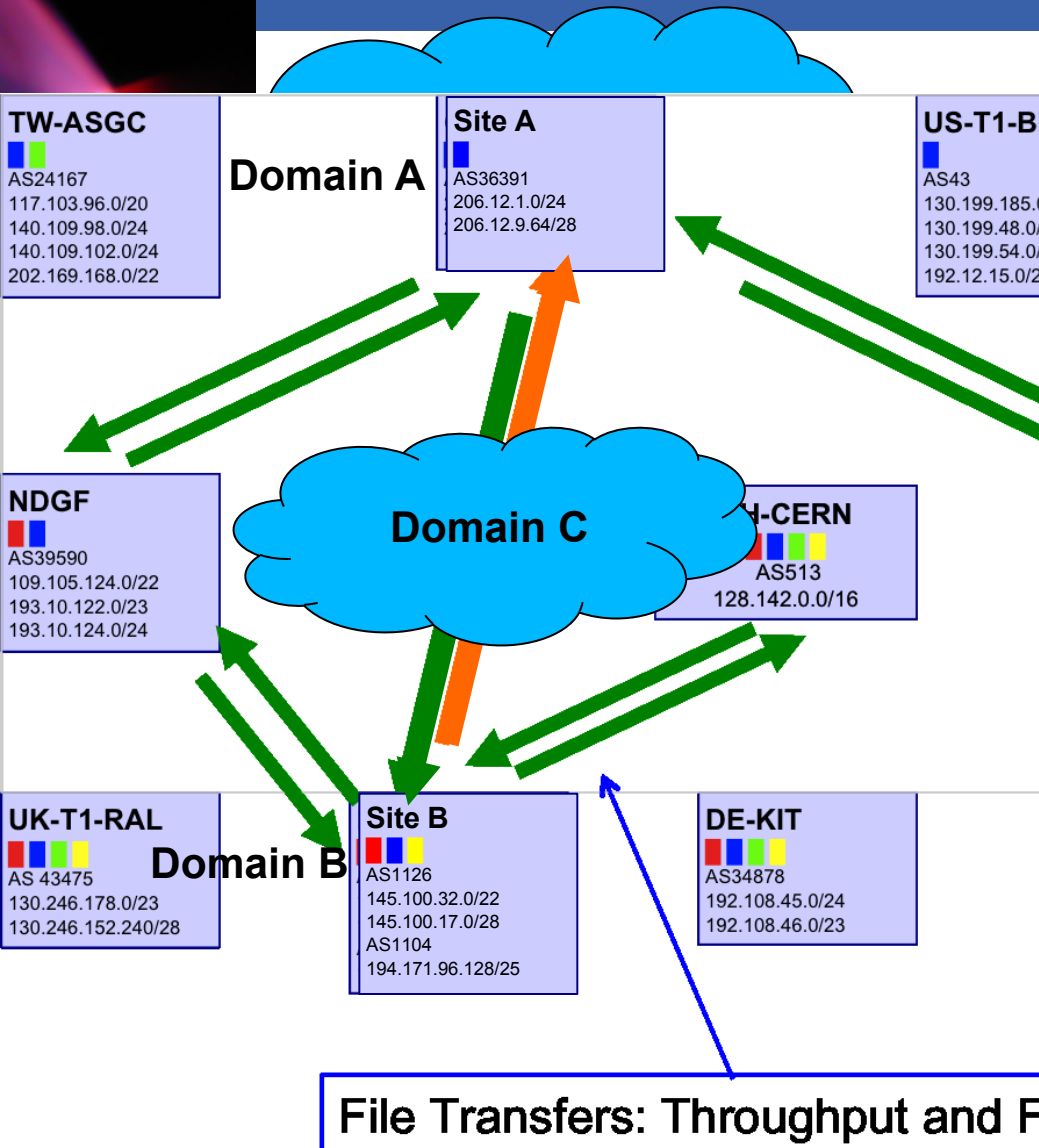
edoardo.martelli@cern.ch 20100916





File Transfers: Throughput and Failure Rates





- VO **X** observes high failure rates / low performance in transfers between sites **A** & **B**
- After basic debugging declared a “network issue”
- Site responsables at both site **A** & **B** informed (ticket)
- They are responsible for updating it and for interactions with network contacts at their respective sites
  - Ticket ownership follows FTS model – i.e. destination site
- All additional complexity – e.g. Domain C and possibly others – transparent to VO **X**
  - NRENs, GEANT, USLHCNET, etc.

- Are responsible for ensuring that the link between the sites is (or becomes) “clean”
- This includes interaction with all relevant partners “in the middle”
- Discussions between “network guys” should be transparent to VO & site representatives
- Site reps get info from their local network reps and ensure ticket is updated

- Assuming that this understanding is correct (which in itself would be a first step), we still see situations where:
  1. Tickets are not regularly updated
  2. Tickets are not updated on “change of state”.
    - E.g. experiment observes much better transfers but no record of what was done to achieve this

Q: is this understanding correct?

Q: can we improve on its implementation?



- A simple model for handling network problems has been discussed at the last LHC OPN meeting
- It applies not only to OPN but also non-OPN links – **highly desirable**
- If the model is agreed we need to formalize it and implement it
- IMHO regular ticket updates are also an important component of the model

- Network contacts A & B are **responsible** for ensuring that their connection to external network and end-to-end link between sites are “clean”
- Network contacts inform site contacts of change of state (preferably) or site contacts poll
- Experiments request at least **daily** update of ticket and on all change of state – at least for those tickets marked **high priority**
- **Problem followed until solved**
- Escalation, if required, based on common WLCG timelines

# ES

# backup



- VO **X** observes high failure rates / low performance in transfers between sites **A** & **B**
- After basic debugging (e.g. to see that its not a simple problem with e.g. source or destination SE and/or transfer s/w), declared a “network issue”
- Site responsables at both site **A** & **B** informed (ticket)
- They now “own” ticket and are responsible for updating it and for interactions with network contacts at their respective sites