

LHCOPN Operation status

*Guillaume Cessieux
(CNRS/CCIN2P3, EGEE networking support)*

LHCOPN meeting, 2009-12-10, Bologna

- **Background & since last LHCOPN meeting**
- **Update on implementation status**
 - Tools and processes
- **Statistics**
 - From TTS and network monitoring
 - Correlation
- **Possibilities for improvements**
- **What's next?**

- **Last LHCOPN meeting: Vancouver, 2009-08-31**
 - <http://indico.cern.ch/conferenceDisplay.py?confId=59842>
- **Outcomes around operations**
 - Pure T1-T1 links (not allowed to carry T0-T1 traffic) to be included in LHCOPN to benefit of existing process and tools
 - Events of interest for WLCG and interactions around to be clarified
 - Existing BGP monitoring no longer suitable
 - Model needs some extra practice before being assessed
- **Pending from Utrecht's LHCOPN meeting, 2009-04-21**
 - <http://indico.cern.ch/conferenceDisplay.py?confId=50345>
 - Information on Twiki need to be reviewed
 - Access to information repositories to be regularly checked
 - KPIs for Ops to be defined and provided

- **LHCOPN Ops phoneconf, 2009-10-08**
 - <https://twiki.cern.ch/twiki/bin/view/LHCOPN/8thOctober2009>
 - 8 sites, no major outcomes
 - Details around GGUS, BGP monitoring tbd
- **LHCOPN Ops training 4, 2009-12-10**
 - <http://indico.cern.ch/conferenceDisplay.py?confId=69191>
 - Attendances NDGF, IT-INFN-CNAF

	Trained	R/W Access to the twiki verified	Access to the TTS verified	Started Ops production mode	Review of twiki
CA-TRIUMF	2009-04-08	2009-04-30	2009-04-30	2009-04-30	Partial 2009-06-19
CH-CERN	2009-04-02	2009-02-04	2009-02-04	2009-02-04	
DE-KIT	2009-04-02	2009-02-23	2009-02-04	2009-02-23	2009-09-14
ES-PIC	2009-04-02	2009-08-12	2009-02-04	2009-02-04	2009-08-12 (Twiki access issue)
FR-CCIN2P3	2009-04-02	2009-02-04	2009-02-04	2009-02-04	2009-12-07
IT-INFN-CNAF	2009-12-10		2009-03-25		
NDGF	2009-06-16	2009-07-14	2009-07-06	2009-07-14	
NL-T1	2009-06-16	2009-06-19	2009-03-20	2009-06-19	2009-10-07
TW-ASGC	2009-04-08	2009-06-03	2009-04-13	2009-06-03	2009-10-30
UK-T1-RAL	2009-06-16	2009-06-23	2009-06-19	2009-06-23	2009-11-04 Postponed 2009-11-23
US-FNAL-CMS	2009-04-08	2009-06-22	2009-05-04	2009-06-22	Started 2009-11-23
US-T1-BNL	2009-04-08	2009-05-27	2009-05-08	2009-05-27	2009-10-20 (twiki access issue)

- **Several “pure” T1-T1 links added**
 - = T1-T1 links not able to carry T0-T1 traffic
 - <https://twiki.cern.ch/twiki/bin/view/LHCOPN/NamingConventionAndLinksIDs>

T1-T1 links					
T1	T1	Link ID	Bandwidth	Default Responsible site	Able to carry T0-T1 traffic
US-T1-BNL	CA-TRIUMF	BNL-TRIUMF-LHCOPN-001	1G	CA-TRIUMF	Y
IT-INFN-CNAF	DE-KIT	CNAF-GRIDKA-LHCOPN-001	10G	IT-INFN-CNAF	Y
NL-T1	CA-TRIUMF	SARA-TRIUMF-LHCOPN-001	1G	NL-T1	N
DE-KIT	NL-T1	GRIDKA-SARA-LHCOPN-001	10G	DE-KIT	Y
DE-KIT	FR-CCIN2P3	GRIDKA-IN2P3-LHCOPN-001	10G	FR-CCIN2P3	Y
NDGF	NL-T1	NDGF-SARA-LHCOPN-001	10G	NDGF	Y
TW-ASGC	NL-T1	SARA-ASGC-LHCOPN-001	1G	TW-ASGC	N
US-FNAL-CMS	NL-T1	SARA-FNAL-LHCOPN-001	1G	US-FNAL-CMS	N
TW-ASGC	US-FNAL-CMS	ASGC-FERMI-LHCOPN-001	2G	TW-ASGC	N

- **Now 34 links IDs in production**
 - 18 T0-T1 + 9 T1-T1 (4 pure) + 7 ‘internal’

- **Process to put new links in production documented**
 - [https://twiki.cern.ch/twiki/bin/view/LHCOPN/OpsFAQ#What is the process to release a](https://twiki.cern.ch/twiki/bin/view/LHCOPN/OpsFAQ#What_is_the_process_to_release_a)
 - Lot of steps usually forgotten
 - Is link monitored
 - Are technical details documented on twiki
 - Is new link advertised

- **Feedbacks from third training session**
 - NDGF:
 - A lot of things are internal site process
 - IT-INFN-CNAF
 - Lack of central operational coordination
 - Monitoring is a key tool
 - Centralising tools is excellent
 - Grid interactions are a complex topic

From Wayne's Vancouver follow up:

- **Relevant Incident for WLCG**
 - Any incident which results in a reduction in the committed bandwidth (in general 10Gbps except for CA-TRIUMF - 5Gbps) or an instability causing unreliable data transfers
- **Major Incident for WLCG**
 - A WLCG Relevant Incident lasting more than 1 working day for a single site and/or a WLCG Relevant Incident affecting multiple sites

- **Review of Twiki information**
 - 6 sites have still reviewed information into
 - Huge amount of wrong information corrected
 - We nearly discovered new links part of the LHCOPN...
 - Indispensable, but taking too much time (~2 weeks/site)

- **Sustainability of GGUS ensured**
 - <https://edms.cern.ch/document/1021528/>
 - Seems moving from EGEE to WLCG

- **IMHO GGUS in a very good shape not preventing good operations to happen**

- **Complete list of opened items maintained on twiki**
 - <https://twiki.cern.ch/twiki/bin/viewauth/LHCOPN/GGUSstatusList>

Major wish list ordered by priority:

1. Display organisation's name of people acting on tickets
2. Link GGUS LHCOPN tickets to standard GGUS LHCOPN tickets
3. API to remotely handle GGUS tickets from sites
4. Checkbox: "clone this ticket and assign it to all impacted sitename" . The main ticket assigned to the submitting site being master of all cloned tickets.
5. Regular tickets / to be clarified by us


Major bug solving ordered by priority:

1. Ticket should be "in progress" instead of "assigned" if self assigned
2. Last diary of tickets in e-mails sometimes disappearing


Rejected:

1. Entire tickets in notification e-mails


Statistics



- **232 'valid' tickets in the TTS**
- **57 privileged users**




Sum of tickets assigned to a site




They now have 6 LHCOPN links

Between 2009-08-01 and 2009-11-30

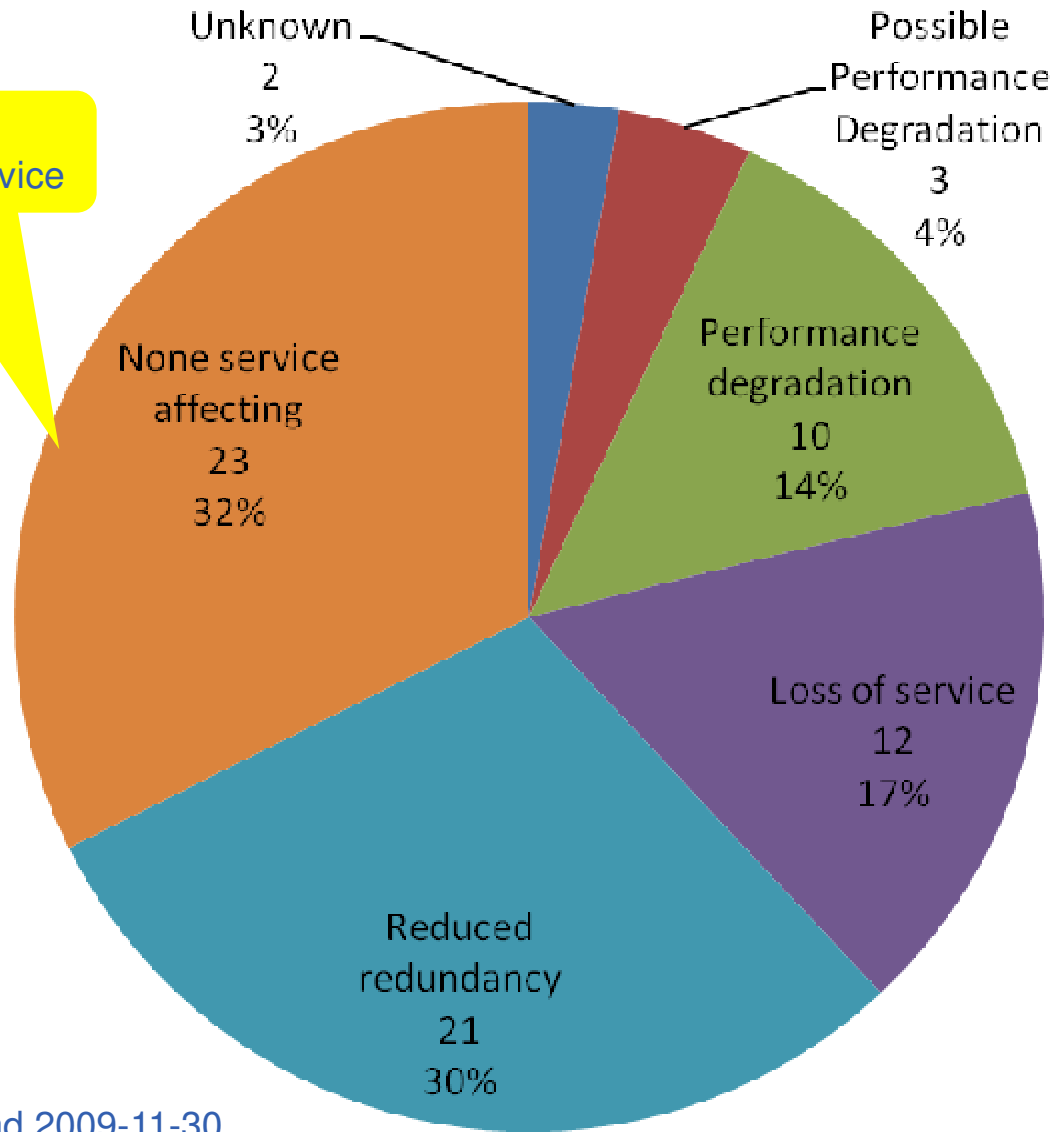


Overall kind of tickets and problems



Between 2009-08-01 and 2009-11-30

~2/3 of events are without impact on service



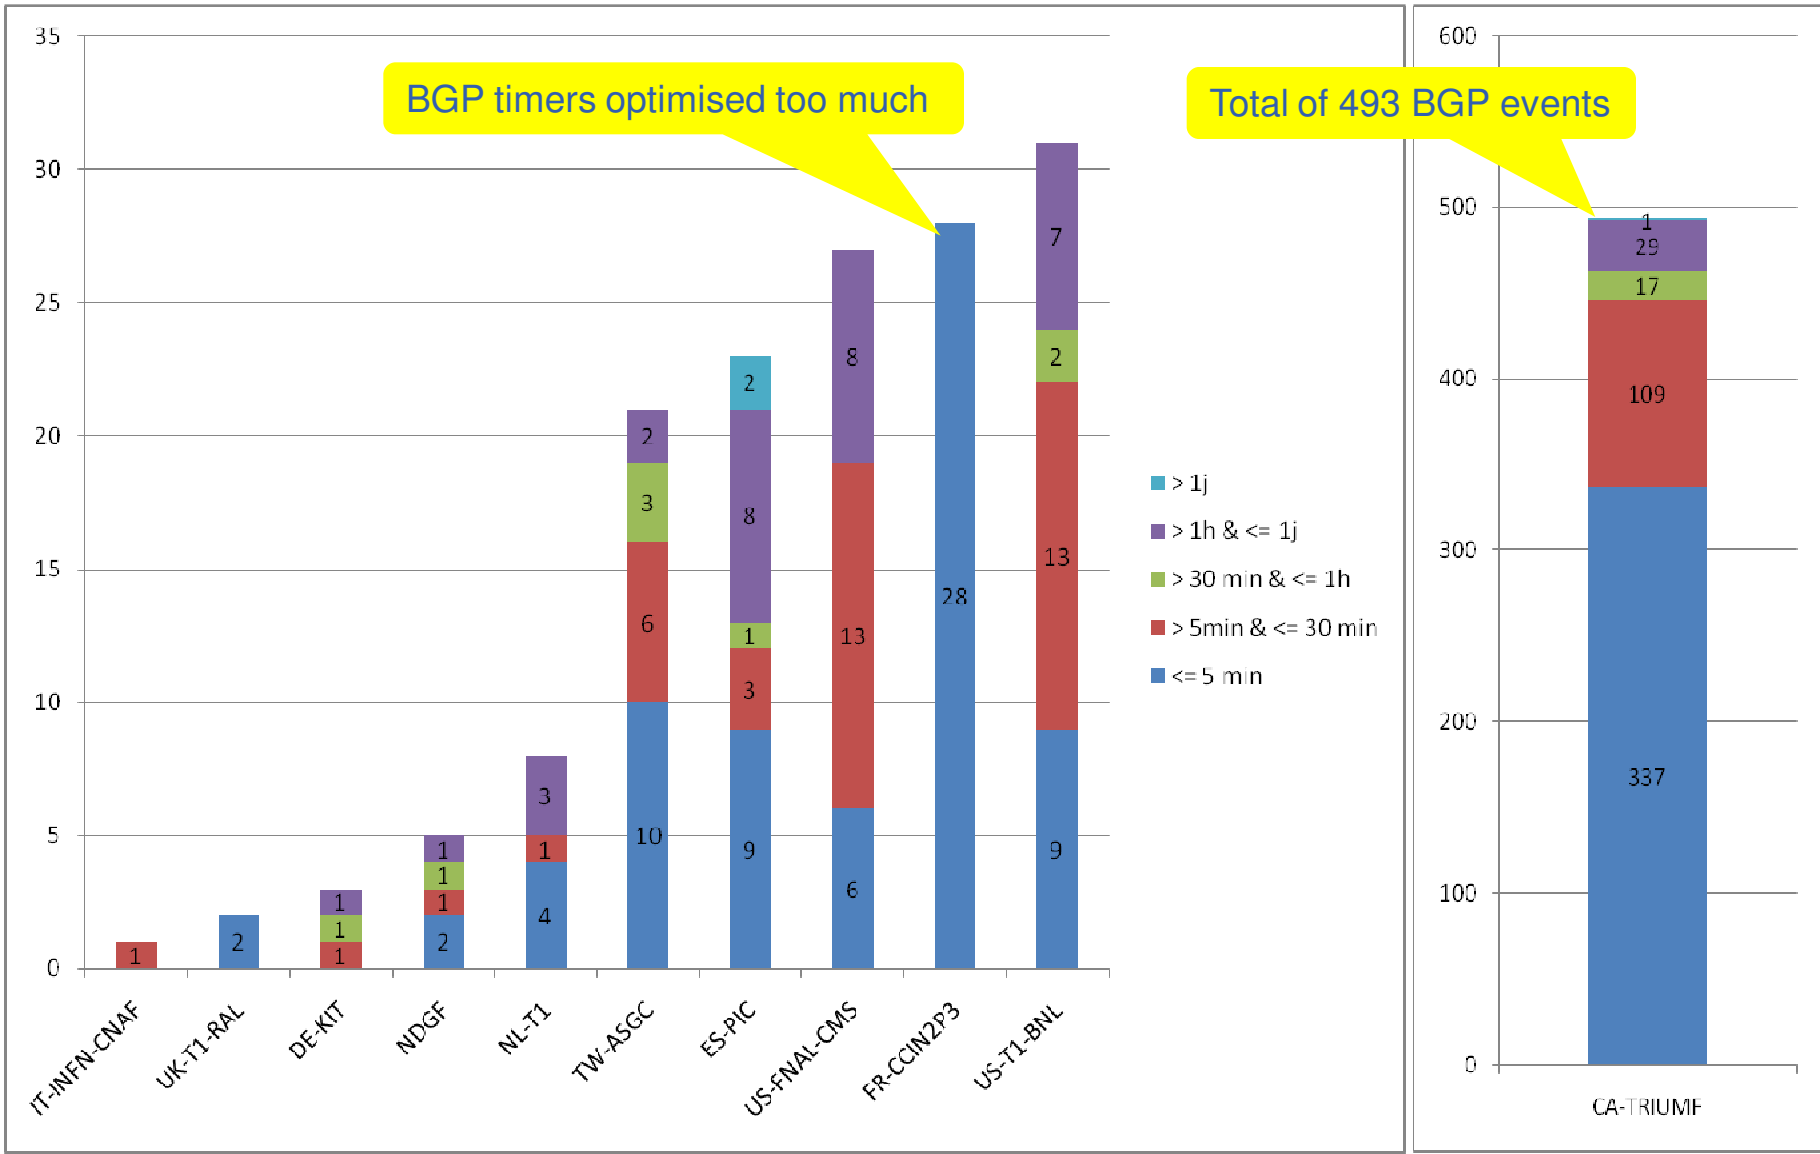
Between 2009-08-01 and 2009-11-30

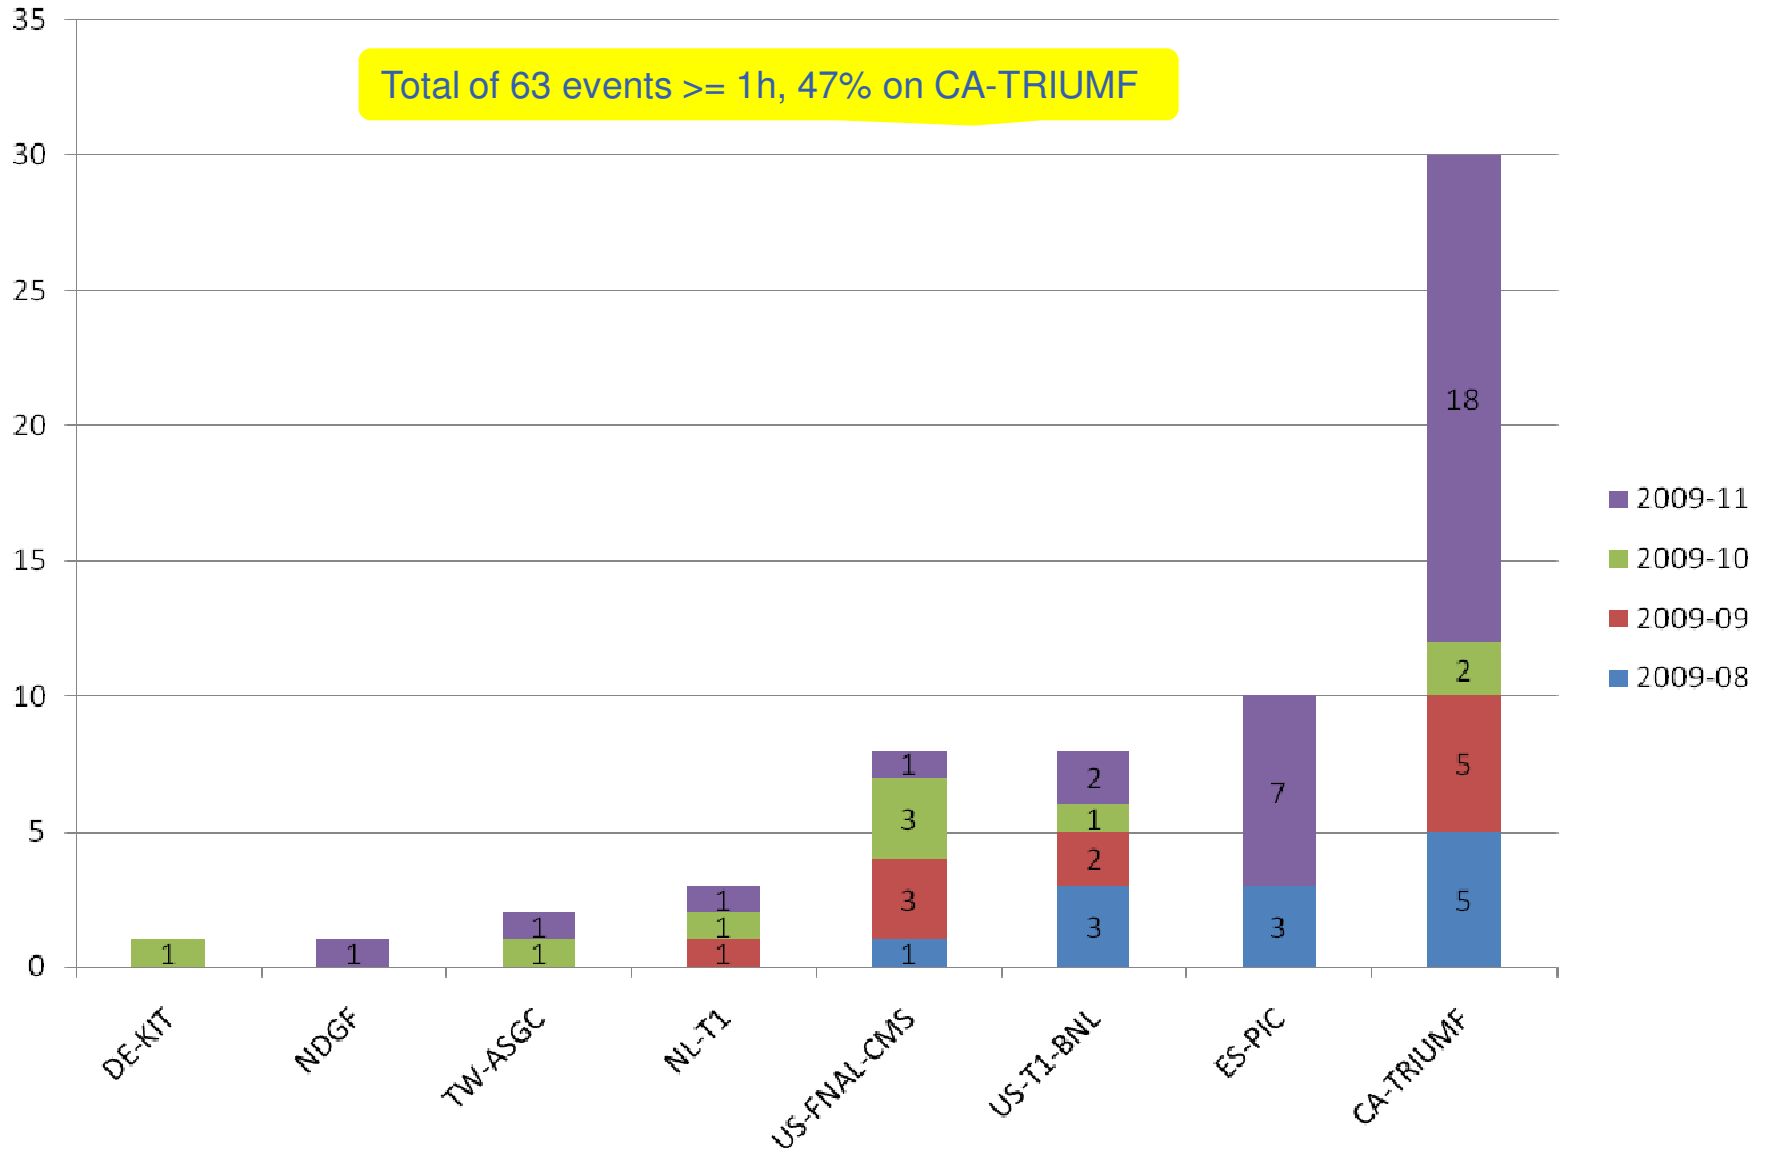
- **Previous trends (2009-02 → 2009-08) confirmed**
 - 2/3 of reported events without impact on service
 - 80% of reported events occurring at L2
- **Hard to assess if behaviour around tickets is normal**
 - Drop in number of tickets etc. strongly linked to infrastructure behaviour
- **GGUS is only reflecting what is registered into**
 - Does this match what occurs on the network?

- **What is occurring on the network?**
- **BGP monitoring from the ENOC stopped 2009-10**
 - You can close SNMP access previously requested
 - Backup through default route and new routers killed it
 - No monitoring is better than wrong
- **CERN offered to provide Spectrum's data**
 - In a database format!
 - Big thanks to Carles Kishimoto (CERN IT-CS) for the technical arrangements
 - BGP peer monitoring events
 - T0-T1 only, from CERN's viewpoint
 - Service impact of BGP peer down unclear

Dispatching per duration

Number of BGP events between 2009-08-01 and 2009-11-30





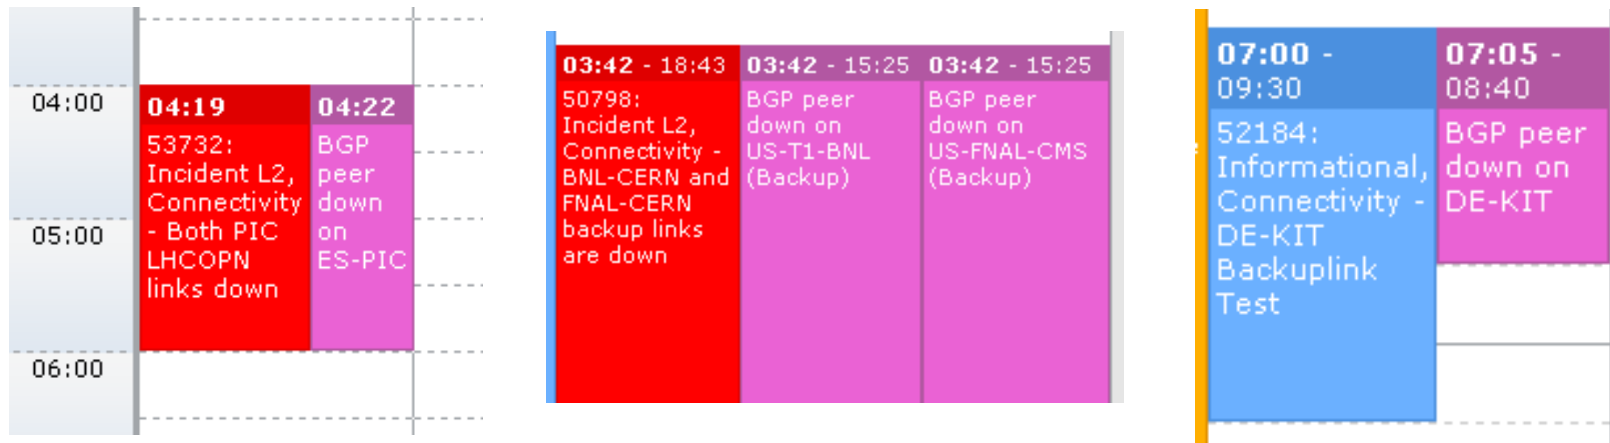
- **Some bias**
 - Only CERN viewpoint, no T1-T1 links
- **Real service impact of BGP events unknown**
 - Dedicated backup path and also through generic IP
 - We miss the service view
- **Not so many events exceeding threshold**
 - Reasonable amount of GGUS tickets expected
 - < 5 tickets per month – But also T1-T1 issues
 - Maybe a little more for CA-TRIUMF

- **Objective: Check if behaviours of operations and infrastructure match**

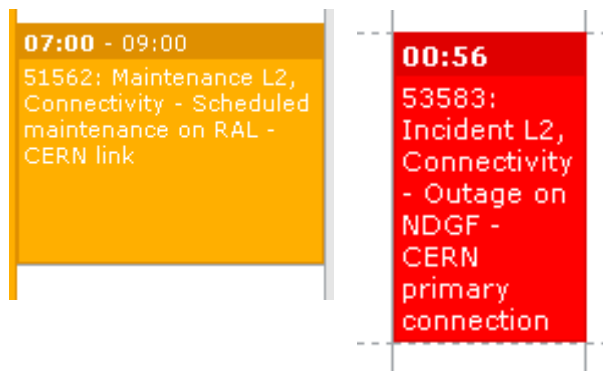
- **Process: See for each monitoring event exceeding threshold if we have a corresponding GGUS ticket**
 - Are significant faults reported as they should?

- **The previously agreed network thresholds were:**
 - <https://twiki.cern.ch/twiki/bin/view/LHCOPN/OperationalModel#Processes>
 - Any event with an impact on the service must be reported with at least 1 ticket per issue not per event
 - For non service impacting events, those lasting more than 1 hour or occurring more than 5 times an hour should be reported in the TTS

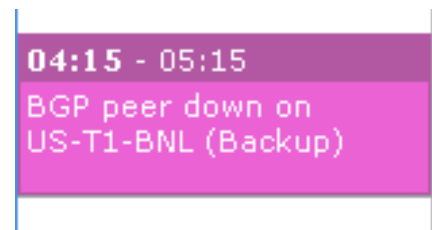
Monitoring 1 – Ticket 1



Monitoring 0 – Ticket 1



Monitoring 1 – Ticket 0



Legend:


- LHCOPN Incident
 - LHCOPN Maintenance
 - LHCOPN Other
 - LHCOPN Spectrum
- events from CH-CERN

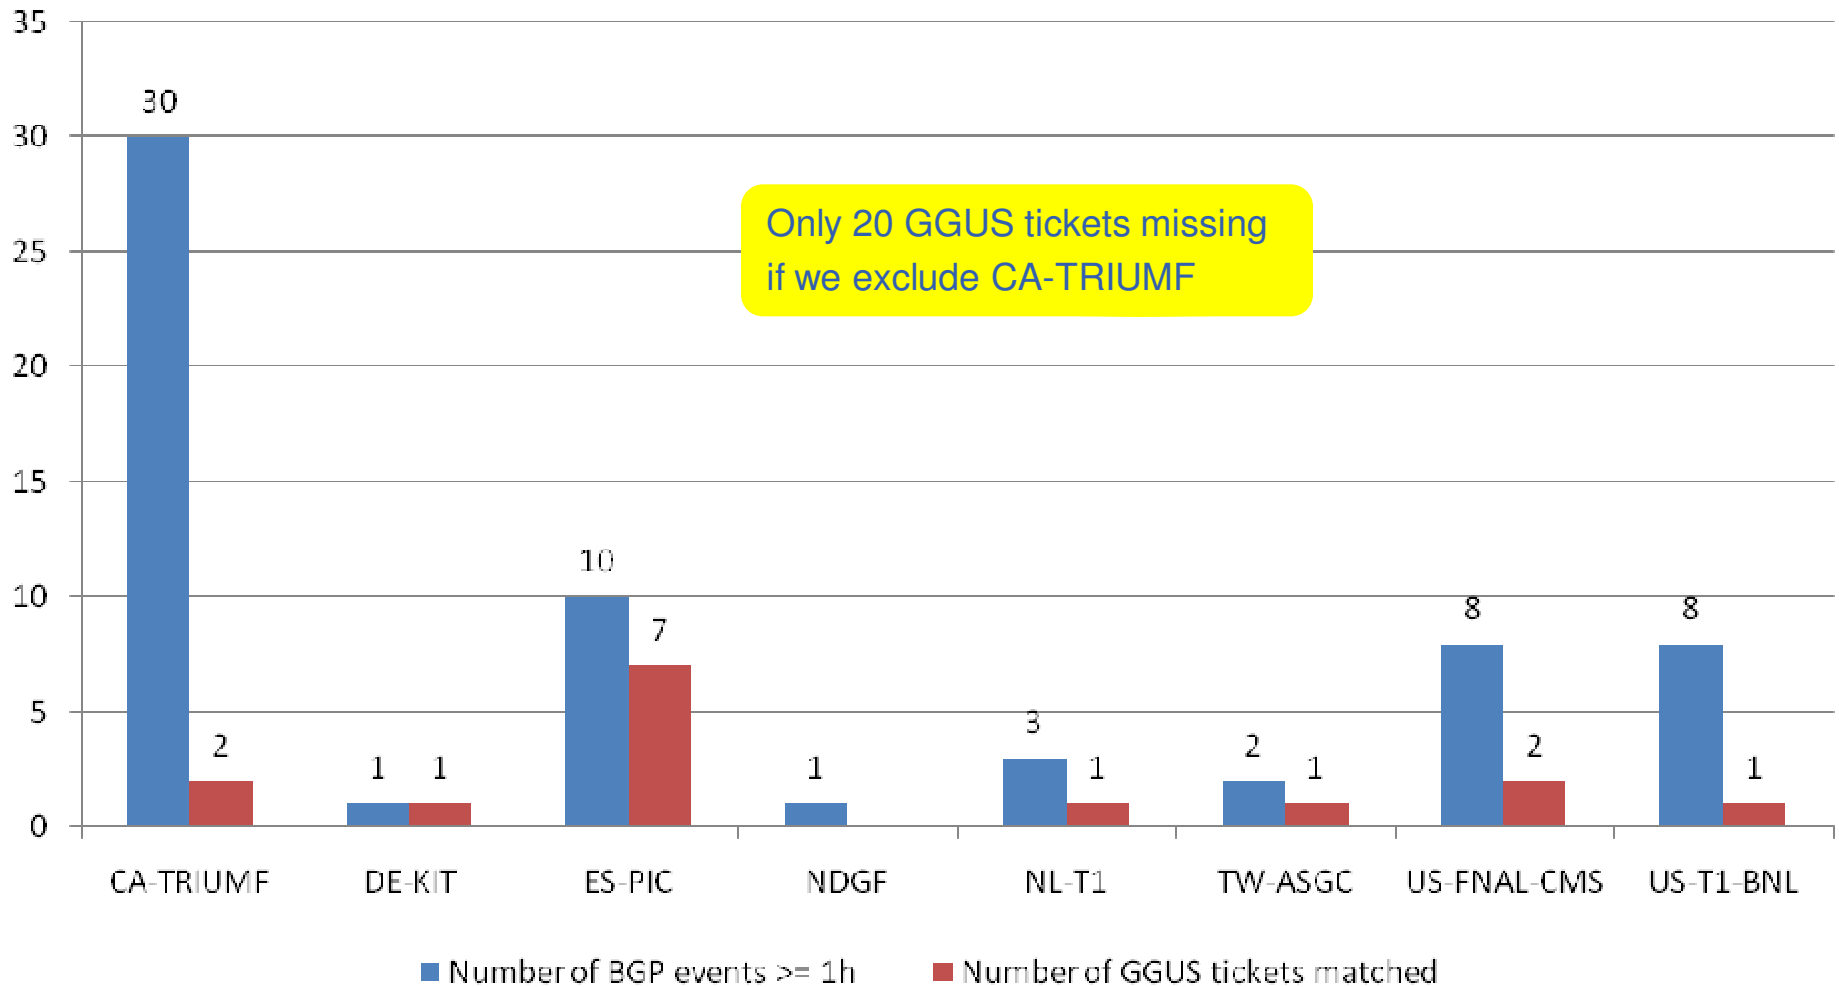
<https://ccenoc.in2p3.fr/LHCOPN/webcalendar/>

- **Daily computation shown on a webpage**
 - <https://ccenoc.in2p3.fr/LHCOPN/report/> [Auth with any valid CERT]
 - With automated correlation
 - **Fully** automated (= maybe errors & no pity on inaccuracy)
 - Show if events are included in time range of GGUS tickets

2009-08					
58	CA-TRIUMF	2009-08-28 13:00:01	2009-08-28 14:42:01	1 hour 42 minutes	
54	CA-TRIUMF	2009-08-26 07:30:01	2009-08-26 17:02:01	9 hours 32 minutes	
52	US-T1-BNL (Backup)	2009-08-23 02:17:02	2009-08-23 04:40:01	2 hours 22 minutes 59 seconds	
51	CA-TRIUMF	2009-08-23 02:15:01	2009-08-23 04:40:01	2 hours 25 minutes	
50	ES-PIC	2009-08-15 06:05:01	2009-08-15 07:37:01	50 minutes 1 second	[#GGUS-50816]
48	CA-TRIUMF (Backup)	2009-08-15 06:05:01	2009-08-15 07:37:01	54 minutes 59 seconds	[#GGUS-50863] [#GGUS-50851]
47	CA-TRIUMF	2009-08-15 06:05:01	2009-08-15 07:37:01	17 hours 20 minutes	
44	ES-PIC	2009-08-15 06:05:01	2009-08-15 07:37:01	1 day 10 minutes	
39	TW-ASGC	2009-08-19 08:17:01	2009-08-19 08:55:01	38 minutes	
27	ES-PIC	2009-08-15 06:05:01	2009-08-15 07:37:01	1 hour 32 minutes	[#GGUS-50949]
24	US-T1-BNL (Backup)	2009-08-14 04:15:01	2009-08-14 05:15:01	1 hour	
9	US-FNAL-CMS (Backup)	2009-08-06 03:42:01	2009-08-06 15:25:01	11 hours 43 minutes	[#GGUS-50798]
8	US-T1-BNL (Backup)	2009-08-06 03:42:01	2009-08-06 15:25:01	11 hours 43 minutes	[#GGUS-50798]
7	ES-PIC	2009-08-04 13:27:01	2009-08-05 15:45:01	1 day 2 hours 18 minutes	[#GGUS-50777]
6	CA-TRIUMF	2009-08-04 04:15:01	2009-08-04 11:47:01	7 hours 32 minutes	

06:03 - 07:53 06:05 - 07:37
 50949: Incident L2, Connectivity - Interoute fiber cut
 BGP peer down on ES-PIC





Only 20 GGUS tickets missing if we exclude CA-TRIUMF

Between 2009-08-01 and 2009-11-30

- **No deep details, only to give global trends**
- **Only 25% of events are correctly matched**
 - If we exclude CA-TRIUMF, 40% are matched
 - 56 tickets not matched: Unnecessary tickets?
 - To be mitigated, a lot are related to T1-T1 issues
- **Confirming some feelings**
 - Maybe too many efforts spent reporting unnecessary things
 - L2 maintenances? Non service impacting events?
 - While some key events are not reported
 - Information improperly relayed or no information?

- **Link monitoring systems with TTS**
 - Automatically create tickets for some key events
 - After ½ hour? 1 hour?
 - Skipped if tickets is still existing
 - But how to avoid killing pro-activity of sites?
 - *Late is better than never?*
 - *Maintenances must be registered in advance*
- **Real time KPIs delivery**
 - Live Ops dashboard to show sites what they miss
 - Done: <https://ccenoc.in2p3.fr/LHCOPN/report/>
- **Review during Ops phoneconf correlation and open tickets**
 - But only post mortem

- **Be sure network teams are doing the strict useful minimum for operations**
 - And give this feeling to them
 - Improve thresholds – balance laziness and risks
 - Really taking account resilience possibilities
 - EU / Non EU sites, L2 / L3 events, Incident/Maintenance, Service impacting / non impacting
 - Can we say we no longer report non service impacting events?
 - Not relevant for WLCG
 - But we need to have reliable backup paths
 - *Only report non service impacting events > 5 hours?*

- **Are shape of processes or tools stoppers?**
 - Submit ticket in one click, ease authentication etc.
 - Better integration heavily required?

- **Some efforts are duplicated as we did not report events at source**
 - Fibre cut in AMS, GVA, etc: Bunch of links affected
 - But several sites need to report after being warned

- **Avoid some steps, increase number of actors**
 - Enable particular providers to report on behalf of sites?
 - Those having end to end vision: USLHCNET...

- **Reliable and accurate network monitoring**
 - Link and service view missing making operations harder
 - Backup through generic IP may completely hide LHCOPN failures
 - Ping etc. won't see LHCOPN failure, delays variation are not reliable
 - BGP won't see backup through generic IP

- **Review of twiki to be finalised**
- **Ops phoneconf, 2010-01-14**
 - Review ops behaviour
- **Ops Working Group phoneconf or meeting**
 - Work around possible improvements
 - KPIs to be accurately precised and regularly computed

- **Tools and processes supporting operations in place and maturing**
 - But still of production quality
- **Assessment of operations to be improved**
 - Heavily relying on correlation with accurate network monitoring
- **We are learning from first production period**
 - Definitely a place for improvements