



LHCOPN: Operations and monitoring status

Guillaume Cessieux

CC-IN2P3 network team / LHCOPN Ops WG leader

Guillaume . Cessieux @ cc.in2p3.fr

LCG France, Lyon, 2010-11-22

WLCG
Worldwide LHC Computing Grid

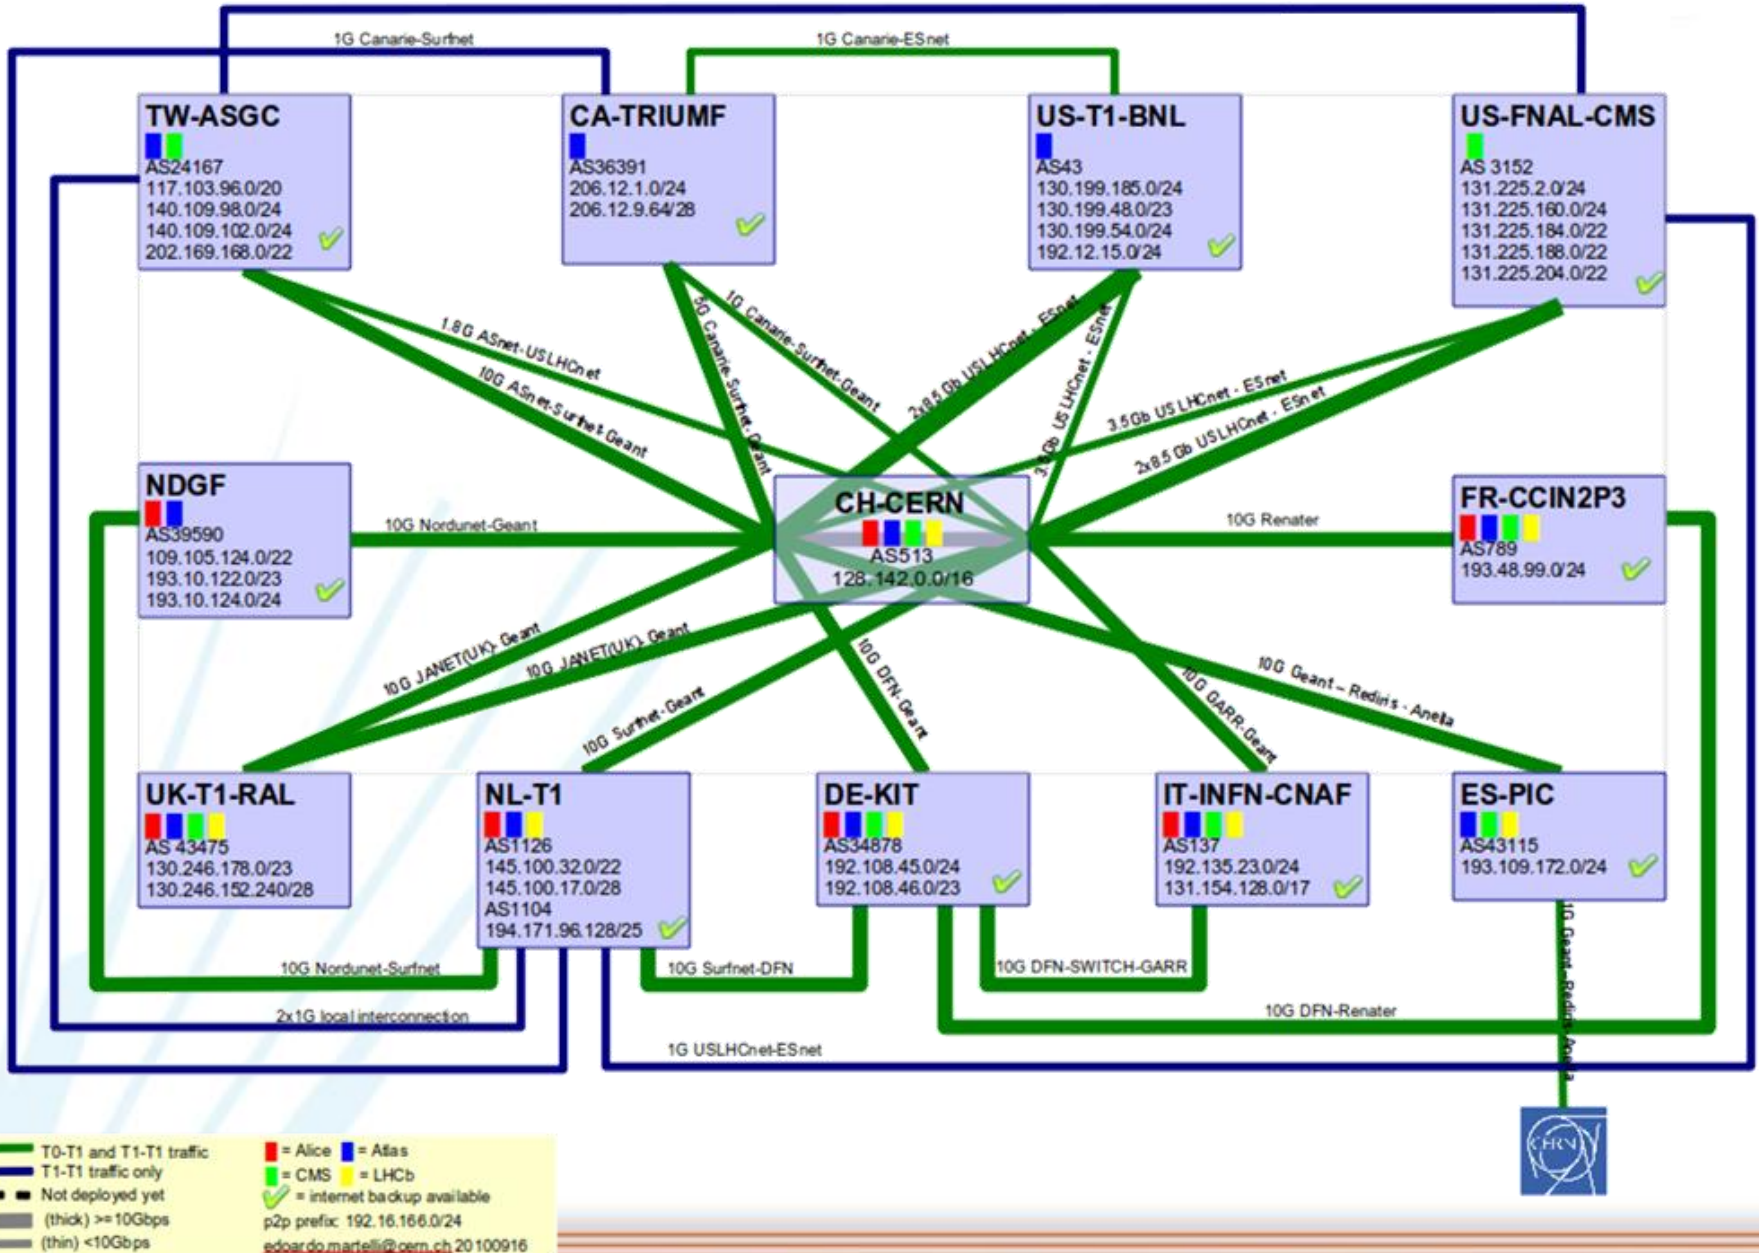


lrfu
cea
saclay

LHCOPN: LHC optical private network

- Missions
 - T0-T1 transfers, T1-T1 as best effort
 - Missions won't be extended (No T2 traffic into)
- Users: Data managers on Tiers 0/1
- Capacity: Mainly 10G links
- High level of resilience and redundancy achieved
 - Traffic able to survive major fibre cut (demonstrated!)
 - Now 28 links in production
 - Very few unexpected service impacting events
 - 4 this year (CH-CERN, US-T1-BNL, TW-ASGC, IT-INFN-CNAF)
- 4 WG: Routing / Security / **Operations** / **Monitoring**

Infrastructure status



Operation status (1/3)

- Lot of work last years
 - Through LHCOPN Ops WG
 - Designing processes, setting up all tools
 - TTS: Dedicated helpdesk with GGUS
 - <http://ggus.org/lhcopn>
 - web repositories: CERN's twiki
 - <https://twiki.cern.ch/twiki/bin/view/LHCOPN/WebHome>
 - Now procedures agreed, implemented and followed
 - Regular review during quarterly Ops phoneconf

Dedicated helpdesk in GGUS

User information

Name E-Mail

CC to

incl. DANTE operation (for monitoring change only)
 ENOC

Notification mode on every change
 on solution

Problem information

Problem Start Date / UTC [▶Current time at T0/T1s](#)

Problem End Date / UTC [▶Set to start date + 1 hour](#)

Ticket Category Priority

Short description (required)

Describe your problem providing the information listed here [?](#)

Impacted sitenames [▶select all](#) [▶unselect all](#) Link ID [Details on Link IDs?](#)

<input type="checkbox"/> CH-CERN	<input type="checkbox"/> FR-CCIN2P3	<input type="checkbox"/> TW-ASGC
<input type="checkbox"/> CA-TRIUMF	<input type="checkbox"/> IT-INFN-CNAF	<input type="checkbox"/> UK-T1-RAL
<input type="checkbox"/> DE-KIT	<input type="checkbox"/> NDGF	<input type="checkbox"/> US-FNAL-CMS
<input type="checkbox"/> ES-PIC	<input type="checkbox"/> NL-T1	<input type="checkbox"/> US-T1-BNL

n/a

ASGC-FERMI-LHCOPN-001 - T1 - 2G

BNL-TRIUMF-LHCOPN-001 - 10G

CERN-ASGC-LHCOPN-003 - P - 10G

CERN-ASGC-LHCOPN-004 - B - 1.8G

CERN-BNL-LHCOPN-001 - P - 10G

No notification to site

Type of network problem Assign it to


Expected impact on L3 services delivered to WLCG

This is a change add information to the CMDB

Dedicated helpdesk in GGUS

LHCOPN - Dashboard

Page loaded: 2010-11-19 09:13 UTC

show/save search result as PDF  | CSV

Show tickets in status to view only changes

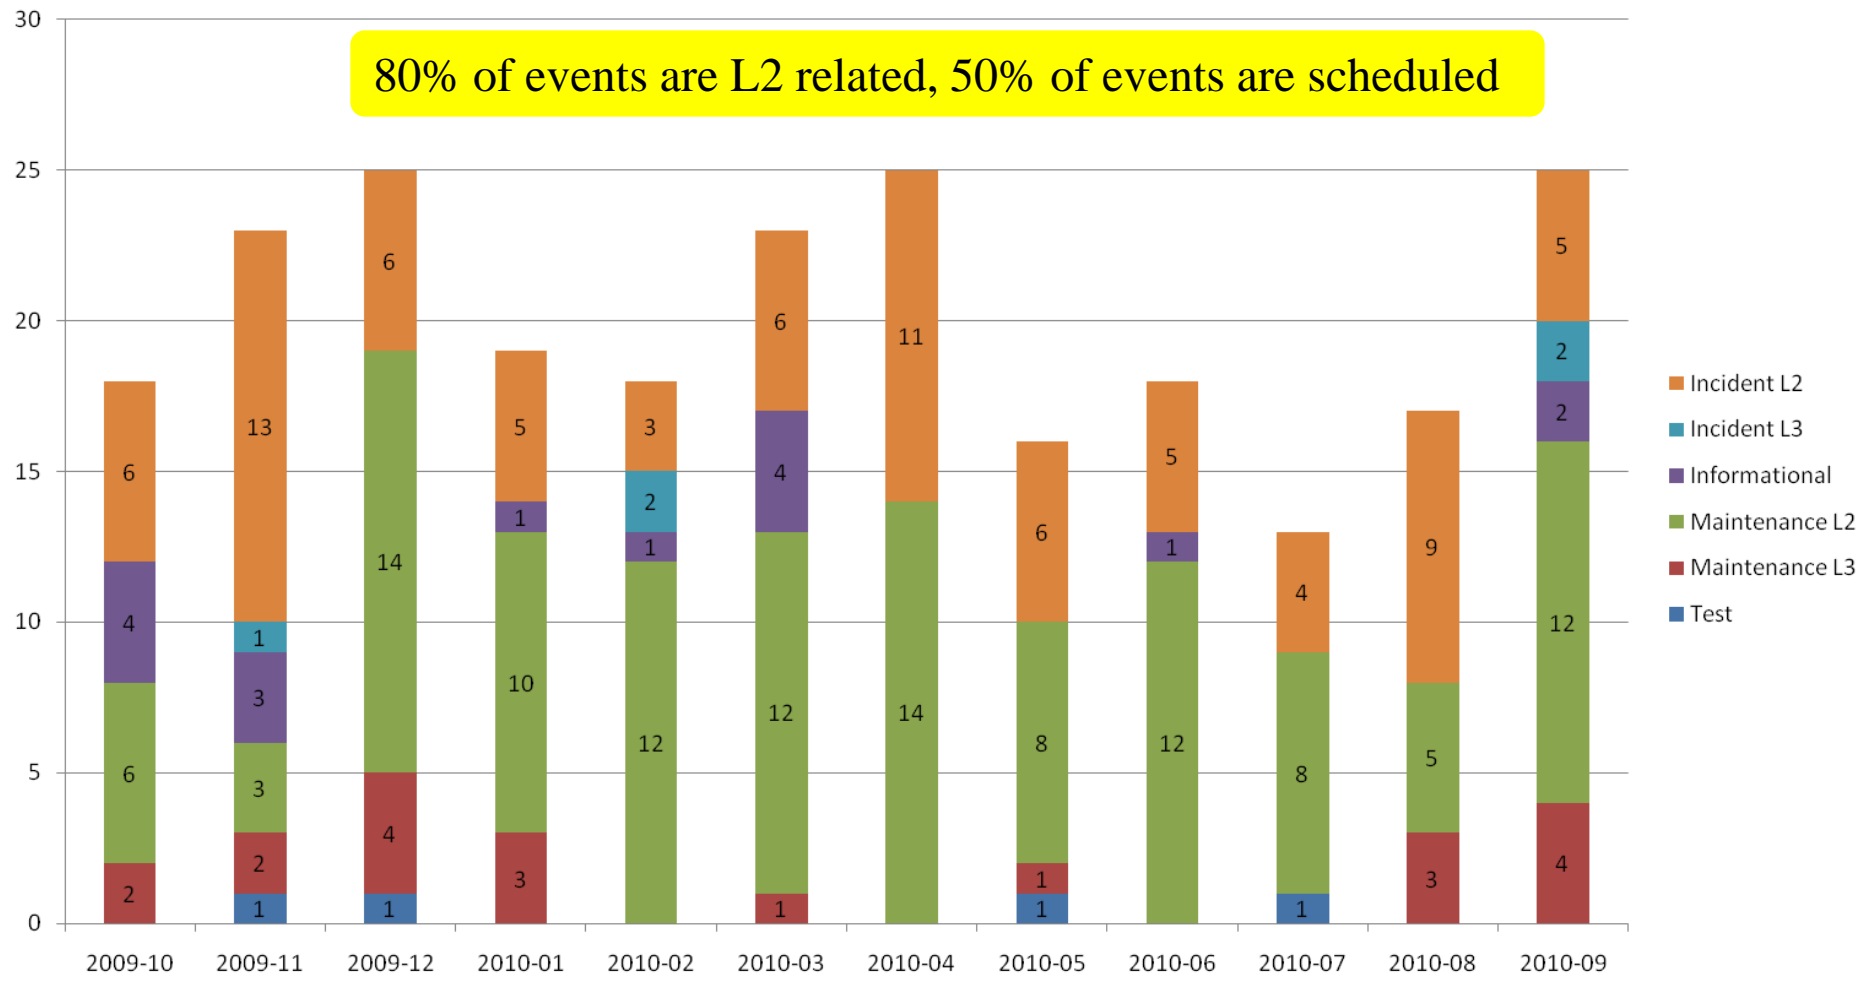
kind of impact ordered how

14 tickets found.

ID	Status	Prob Start	Prob End	Category*	Problem*	Assigned to	LinkID	Imp.SiteName	Info
▶ 64334	in progress	2010-11-29 23:00	2010-11-30 05:00	ML3	Connect	TW-ASGC	TW-ASGC-AMS-CHI-LHCOPN-001 - 2.5G	CA-TRIUMF CH-CERN DE-KIT ES-PIC FR-CCIN2P3 NDGF NL-T1 UK-T1-RAL	Network Maintenance
▶ 64321	in progress	2010-11-18 12:00	2010-11-18 13:00	ML2	Connect	US-FNAL-CMS	ASGC-FERMI-LHCOPN-001 - T1 - 2G CERN-FERMI-LHCOPN-001 - P - 10G CERN-FERMI-LHCOPN-002 - B - 3.5G CERN-FERMI-LHCOPN-003 - S - 7G SARA-FNAL-LHCOPN-001 - T1 - 1G	CH-CERN DE-KIT ES-PIC FR-CCIN2P3 IT-INFN-CNAF NL-T1 TW-ASGC UK-T1-RAL	FNAL power outage affecting all off-site connectivity.
▶ 64272	in progress	2010-11-29 23:00	2010-11-30 05:00	ML2	Connect	NL-T1	CERN-SARA-LHCOPN-001 - P - 10G	CH-CERN NL-T1	Planned maintenance by carrier
▶ 64270	assigned	2010-11-15 15:18	n/a	Info	None	US-T1-BNL	n/a	US-T1-BNL	Documentation of routing policies
▶ 64269	assigned	2010-11-15 15:17	n/a	Info	None	US-FNAL-CMS	n/a	US-FNAL-CMS	Documentation of routing policies
▶ 64267	assigned	2010-11-15 15:15	n/a	Info	None	NDGF	n/a	NDGF	Documentation of routing policies
▶ 64266	assigned	2010-11-15 15:14	n/a	Info	None	IT-INFN-CNAF	n/a	IT-INFN-CNAF	Documentation of routing policies
▶ 64264	assigned	2010-11-15 15:12	n/a	Info	None	DE-KIT	n/a	DE-KIT	Documentation of routing policies
▶ 64246	in progress	2010-12-15 23:00	2010-12-16 05:00	ML2	Connect	DE-KIT	GRIDKA-SARA-LHCOPN-001 - 10G	DE-KIT NL-T1	DFN maintenance
▶ 64245	in progress	2010-12-15 22:00	2010-12-16 05:00	ML2	Connect	DE-KIT	CERN-GRIDKA-LHCOPN-001 - P - 10G	CH-CERN DE-KIT	DFN maintenance
▶ 63554	in progress	2011-01-11 08:00	2011-01-11 12:00	ML2	Connect	NL-T1	SARA-TRIUMF-LHCOPN-001 - T1 - 1G	CA-TRIUMF NL-T1	Planned maintenance by carrier

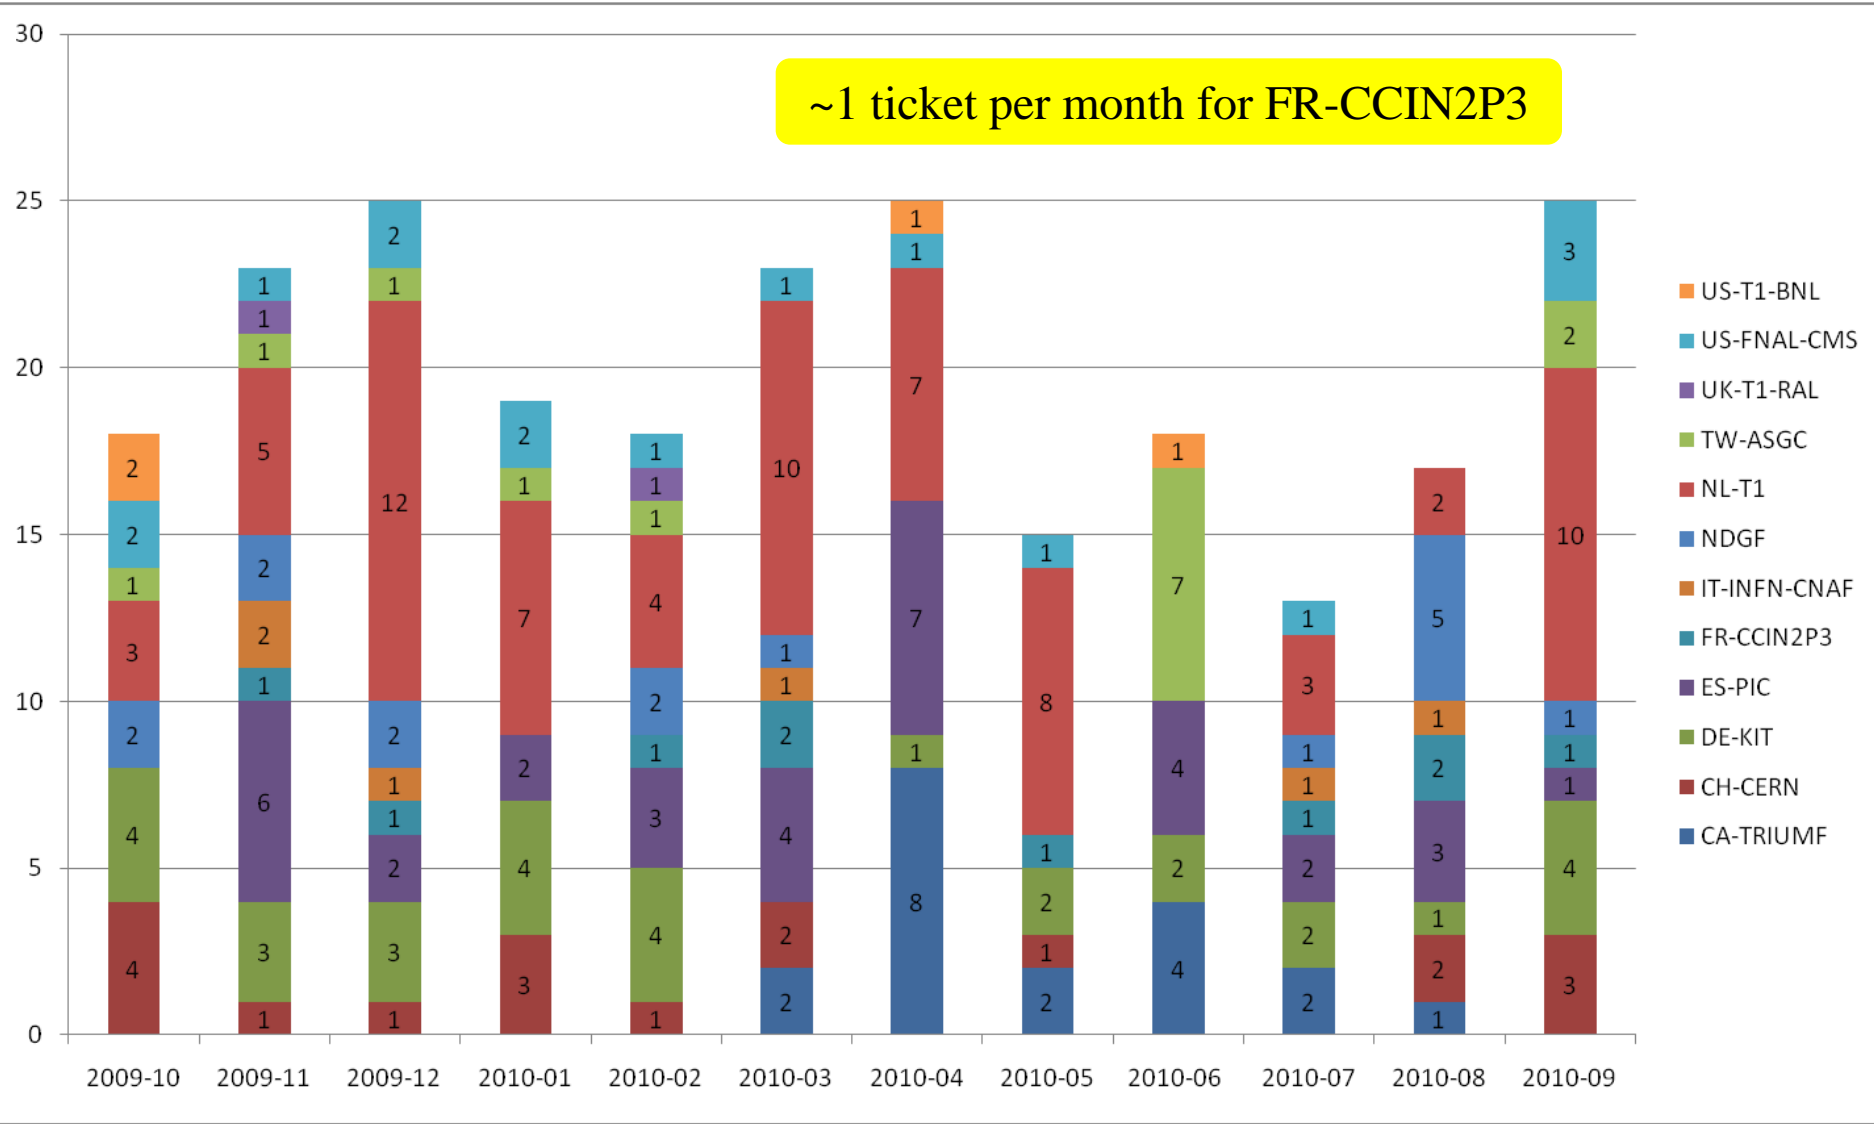
Kind of tickets per month

80% of events are L2 related, 50% of events are scheduled



Ownership of tickets per month per site

~1 ticket per month for FR-CCIN2P3



Operation status (2/3)

- Nothing major to highlight
 - Lot of regular events, few service affecting
 - Information delivered to WLCG must be filtered
 - Main complex network issues lie at end sites
 - Performance issues or LHCOPN bypassed
 - **Lot of network issues for WLCG are not anyhow related to the LHCOPN**
- 2 main weaknesses being addressed
 - Few operational relationships with WLCG
 - Mainly communication issue
 - Must rely on accurate network monitoring

Operation status (3/3)

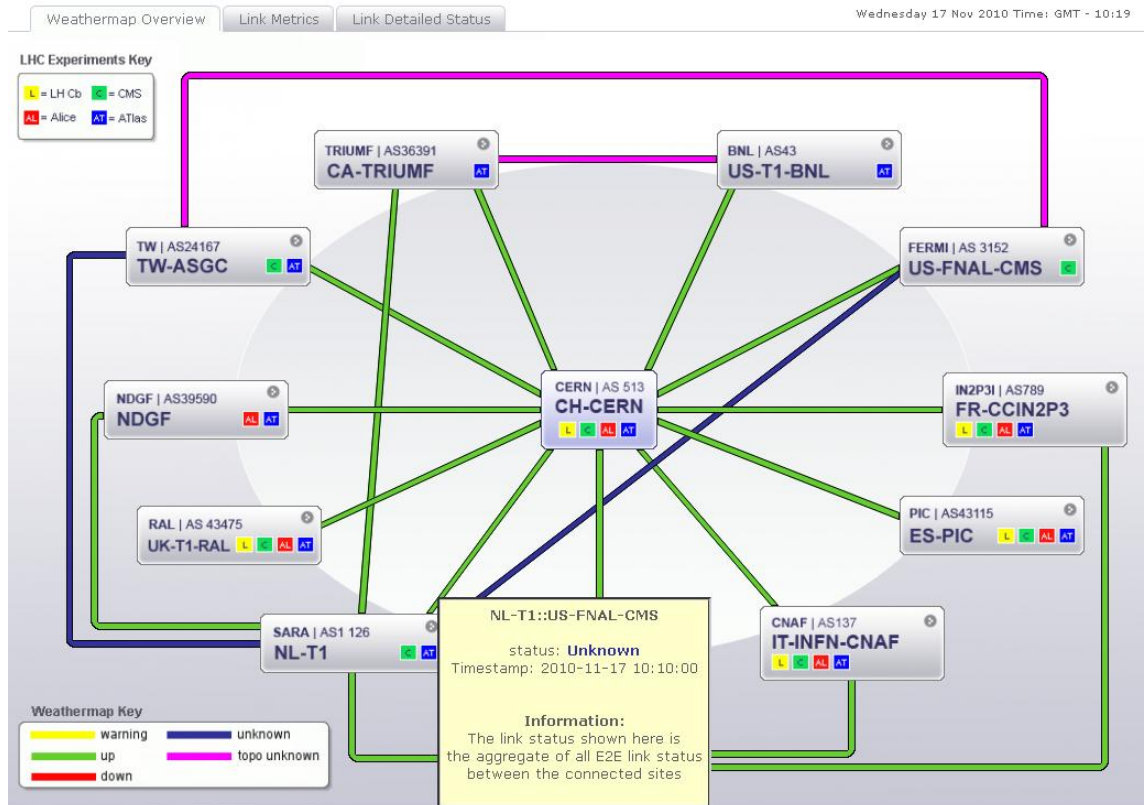
■ Ongoing work

- Extend and unify LHCOPN processes to cope with standard network issues on public networks
 - Ownership of trouble
 - Fit within current support scheme
 - Ensure timely resolution and regular updates
 - Clarify communication channels: network providers/sites/Project
- Single point of contact for WLCG
 - But distributed networking support unit behind
- Wait Tiers2 networking infrastructure to reach consensus before studying how LHCOPN experience can benefit

Monitoring status (1/5)


- Long standing topic ~4 years old
 - Not vital for network operations: Each site still monitoring its subpart
 - Missing: Global service view for experiments
- Task previously fully outsourced to GN2
 - Based on perfSONAR MDM solution especially designed and tuned
 - MDM: Multi domain Monitoring
 - Data gathering & presentation
- Hardware solution deployed on each Tiers 0/1
 - Heavy: 3 appliances per site remotely managed by GN2
 - 1 bandwidth measurement: Active tests, BWCTL, 1G injected each 30 min
 - 1 latency measurement: OWD
 - 1 scheduler and metrics database
 - Central collector currently hosted by DFN (GN2)

Monitoring status (2/5)




perfSONAR MDM weathermap


E2EMon aggregated status of abstract site-site link 'CH-CERN::TW-ASGC' [Info](#)



utilization of >-rftec-1.cern.ch::TenGigabitEthernet 0/2::192.16.166.1::stub' [Info](#)



errors of >-rftec-1.cern.ch::TenGigabitEthernet 0/2::192.16.166.1::stub' [Info](#)



Monitoring status (3/5)

Leadership changed 2010-07

- GN3 now only in charge of data gathering
- CERN leading revamped Monitoring WG
 - Including people from experiments
 - SARA (NL-T1) take the lead for the dashboard
 - CERN working on integration

Dashboard requirements

- Simple for users
- In-depth view for operators
- Reliability is vital

Draft proposal functional design LHCOPN dashboard

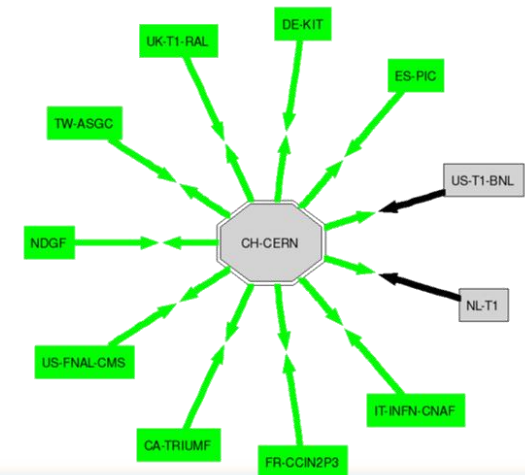
Draft proposal functional design
LHCOPN dashboard
LHCOPN Monitor Working Group

Document history:

Version:	Details:	Editor:	Date:
0.1	Draft version, NL-T1 internal	Hanno Piet (NL-T1/SARA)	8-9-2010
0.2	Draft version, comments SB and II	Hanno Piet (NL-T1/SARA)	10-9-2010
0.3	Draft version, comments YC 15-5-2010	Hanno Piet (NL-T1/SARA)	18-4-2010
0.4	Draft version, comments JM, SB and II	Hanno Piet (NL-T1/SARA)	22-5-2010
0.5	Draft version, comments JS and JM	Hanno Piet (NL-T1/SARA)	25-5-2010
0.6	Draft version, comments LHCOPN meeting Barcelona	Hanno Piet (NL-T1/SARA)	6-7-2010
0.7	Draft version, comments JS and ST	Hanno Piet (NL-T1/SARA)	15-7-2010

NL-T1 editor:

©2010 SARA Computing & Networking Services 1/15, 15-7-2010



Monitoring status (4/5)

- Tricky part is logic behind data presentation
 - SLD: What is a degraded situation?
 - Resilience working, thresholds etc.

- UP: (One-Way Delay +/- 15% / 5 minutes) && (Packet Loss < 0.1% / 5 min)
- DOWN: Packet Loss = 100% / 5min
- DEGRADED: All conditions between Up and DOWN.

- Ongoing

- Audit from GN2 about the solution deployed
 - What's working or not
- Dashboard prototype being prepared by NL-T1
 - Target: February 2011

Monitoring status (5/5)

Preview of dashboard

End to End IPv4 unicast connectivity availability (current view)	To CA-TRIUMF	To CH-CERN	To DE-KIT	To ES-PIC	To FR-CC-IN2P3	To IT-INFN-CNAF	To NDGF	To NL-T1	To TW-ASGC	To UK-T1-RAL	To US-BNL	To US-FNAL-CMS	
From CA-TRIUMF	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	
From CH-CERN	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	
From DE-KIT	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	
From ES-PIC	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	
From FR-CC-IN2P3	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	
From IT-INFN-CNAF	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	
From NDGF	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	
From NL-T1	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	
From TW-ASGC	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	
From UK-T1-RAL	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	
From US-BNL	50%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	
From US-FNAL-CMS	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	
Date and time: 17-6-2010 12:30 UTC													
What does "Normal" mean	End to End IPv4 unicast connectivity availability (last 5min view)												
What does "Degraded" mean	From CH-CERN	To CH-CERN	To CA-TRIUMF	To DE-KIT	To ES-PIC	To FR-CCIN2P3	To IT-INFN-CNAF	To NDGF	To NL-T1	To TW-ASGC	To UK-T1-RAL	To US-FNAL-CMS	To US-T1-BNL
What does "Down" mean	(89.84ms)	(90.77ms)	(6.142ms)	(16.50ms)	(1.832ms)	(6.618ms)	(11.71ms)	(12.85ms)	(152.3ms)	(9.934ms)	(53.38ms)	(53.38-53)	(53.38-53)
What does "No data" mean	(89.84-80)	(90.77-90)	(6.142-5.5)	(16.50-16)	(1.832-1.4)	(6.618-5.5)	(11.71-11)	(12.85-12)	(152.3-147)	(9.934-11)	(53.38-53)	(53.38-53)	(53.38-53)
	=9.849	=0.773	=0.642	=0.503	=0.432	=1.118	=0.715	=0.853	=5.353	=-1.06	=0.387	=0.387	=0.387
	=12.31%+Hoss	=0.859%+Hoss	=11.67%+Hoss	=3.148%+Hoss	=30.85%+Hoss	=20.34%+Hoss	=6.501%+Hoss	=7.109%+Hoss	=3.641%+Hoss	=9.689%+Hoss	=0.730%+Hoss	=0.730%+Hoss	=0.730%+Hoss
	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
	From CA-TRIUMF	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN
	(89.84ms)	(95.67ms)	(106.0ms)	(99.49ms)	(101.2ms)	(78.60ms)	(241.9ms)	(81.14ms)	(40.18ms)	(40.18-41)	(40.18-41)	(40.18-41)	(40.18-41)
	(89.84-80)	(95.67-90)	(106.0-100)	(99.49-95)	(101.2-101)	(78.60-80)	(241.9-237)	(81.14-79)	(40.18-41)	(40.18-41)	(40.18-41)	(40.18-41)	(40.18-41)
	=9.849	=5.678	=6.067	=4.498	=0.257	=-1.29	=4.920	=2.149	=-0.81	=-0.81	=-0.81	=-0.81	=-0.81
	=12.31%+Hoss	=6.309%+Hoss	=6.067%+Hoss	=4.734%+Hoss	=0.254%+Hoss	=1.737%+Hoss	=2.076%+Hoss	=2.720%+Hoss	=1.990%+Hoss	=1.990%+Hoss	=1.990%+Hoss	=1.990%+Hoss	=1.990%+Hoss
	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
	From DE-KIT	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN
	(6.907ms)	(101.5ms)	(27.47ms)	(4.770ms)	(6.229ms)	(20.60ms)	(7.097ms)	(160.0ms)	(14.40ms)	(67.31ms)	(67.31-58)	(67.31-58)	(67.31-58)
	(6.907-5.5)	(101.5-95)	(27.47-20)	(4.770-4.5)	(6.229-5.5)	(20.60-16)	(7.097-7)	(160.0-152)	(14.40-10)	(67.31-58)	(67.31-58)	(67.31-58)	(67.31-58)
	=-1.407	=6.521	=7.477	=0.270	=0.729	=4.604	=-0.097	=8.073	=4.405	=9.314	=9.314	=9.314	=9.314
	=25.59%+Hoss	=6.865%+Hoss	=37.38%+Hoss	=6.022%+Hoss	=13.27%+Hoss	=28.78%+Hoss	=1.385%+Hoss	=5.311%+Hoss	=44.05%+Hoss	=16.05%+Hoss	=16.05%+Hoss	=16.05%+Hoss	=16.05%+Hoss
	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
	From ES-PIC	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN
	(16.28ms)	(120.5ms)	(22.51ms)	(29.86ms)	(31.73ms)	(167.9ms)	(30.47ms)	(71.20ms)	(71.20-69)	(71.20-69)	(71.20-69)	(71.20-69)	(71.20-69)
	(16.28-15)	(120.5-100)	(22.51-20)	(29.86-21)	(31.73-28)	(167.9-148)	(30.47-21)	(71.20-69)	(71.20-69)	(71.20-69)	(71.20-69)	(71.20-69)	(71.20-69)
	=1.280	=20.56	=2.512	=8.860	=3.739	=19.96	=9.470	=2.205	=2.205	=2.205	=2.205	=2.205	=2.205
	=8.539%+Hoss	=20.56%+Hoss	=12.56%+Hoss	=42.19%+Hoss	=6.027%+Hoss	=13.35%+Hoss	=13.48%+Hoss	=45.09%+Hoss	=3.196%+Hoss	=3.196%+Hoss	=3.196%+Hoss	=3.196%+Hoss	=3.196%+Hoss
	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok
	From FR-CCIN2P3	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN	From CH-CERN
	(0.918ms)	(91.44ms)	(4.748ms)	(17.15ms)	(12.71ms)	(153.0ms)	(7.431ms)	(54.03ms)	(54.03-54)	(54.03-54)	(54.03-54)	(54.03-54)	(54.03-54)
	(0.918-1.3)	(91.44-90)	(4.748-5=-0.25)	(17.15-16)	(12.71-10)	(153.0-148)	(7.431-7.4)	(54.03-54)	(54.03-54)	(54.03-54)	(54.03-54)	(54.03-54)	(54.03-54)
	=-0.38	=1.447	=-5.037%+Hoss	=1.154	=-2.714	=5.023	=-0.031	=0.038	=0.038	=0.038	=0.038	=0.038	=0.038
	=29.31%+Hoss	=1.607%+Hoss	=-5.037%+Hoss	=7.218%+Hoss	=-27.14%+Hoss	=3.374%+Hoss	=-9.569%+Hoss	=3.394%+Hoss	=-0.419%+Hoss	=-0.419%+Hoss	=-0.419%+Hoss	=-0.419%+Hoss	=-0.419%+Hoss
	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok	ok

No strong effort on data presentation

- Data shared, up users to embed this wherever they want
- Alarms desired

Conclusion

- Infrastructure
 - Now fully deployed, heavily redundant
- Operations
 - Steady state, waiting monitoring & SLD for next improvements
 - Main weakness being addressed: Grid interactions
- Monitoring
 - Known as a key requirement from project
 - Metrics gathering on sites seems now ok
 - Global service view being made