CMS FacilitiesOps and IN2P3

[CMS visit to IN2P3 - Lyon, 23 Oct 09]

Daniele Bonacorsi, Peter Kreuzer
[CMS Facilities Ops]

Claudio Grandi, Chris Brew
[T1 coordination in CMS Facilities Ops]

Andrea Sciabà, Josep Flix
[Site Readiness in CMS Facilities Ops]

Nicolò Magini
[Data Transfer Operations and DDT in CMS Facilities Ops]

CMS FacilitiesOps

CMS FacilitiesOps weekly meetings

- ✦ To discuss status of T1 and T2 sites, and related items, over last 7 days
  - CMS contacts at T1’s asked to provide brief weekly reports
  - SAM and SiteReadiness status is reviewed, explanations are asked, discussion
- ✦ Weekly, Monday afternoon, 5pm GVA time

CMS attends WLCG Ops daily calls, 3pm GVA time

- ✦ Official WLCG official minutes:
  - https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOperationsMeetings
- ✦ Collection of CMS daily reports:
  - https://twiki.cern.ch/twiki/bin/view/CMS/FacOps_WLCGdailyreports

SAM Availability for CMS T1’s

CMS-specific SAM tests

- ✦ Complementary to WLCG SAM, to mimic real CMS workflows
  - Widely documented elsewhere

Overall SAM Availability ranking for CMS T1’s: goal is **90%**

- ✦ For all orangish/redish boxes we discuss at FacOps weekly meetings

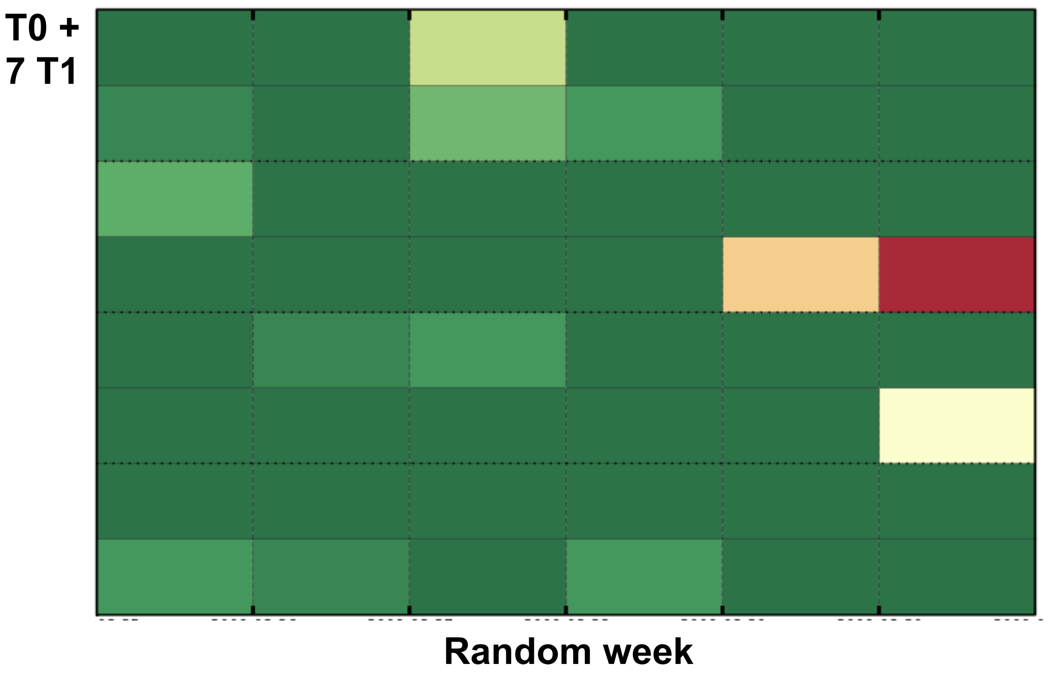

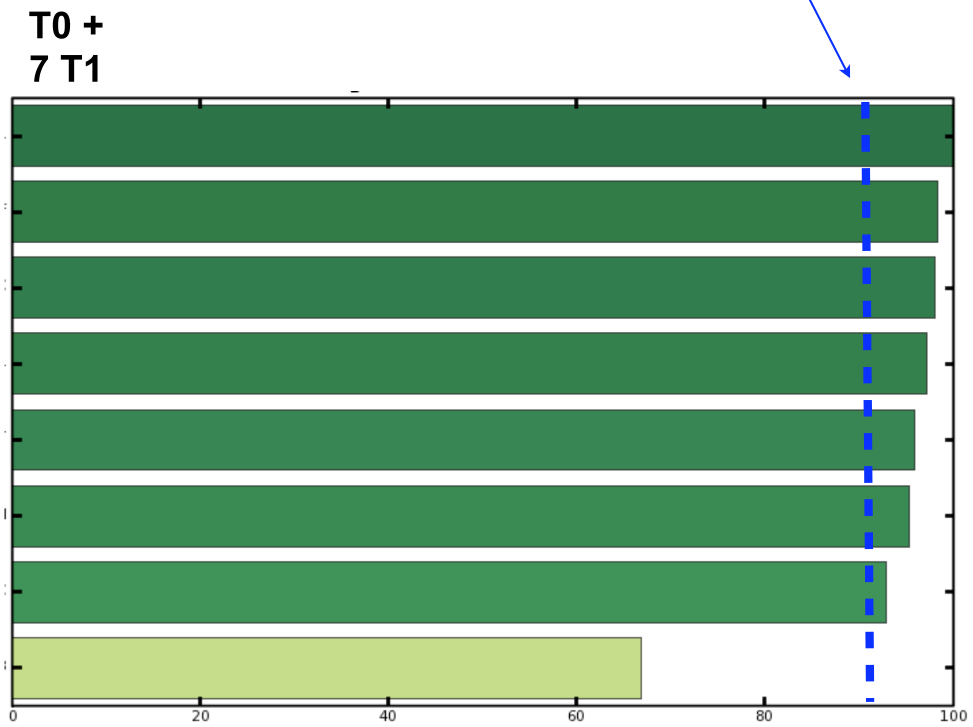

SAM Availability for ***IN2P3***

Looking to IN2P3 in CMS-specific SAM tests in 2009

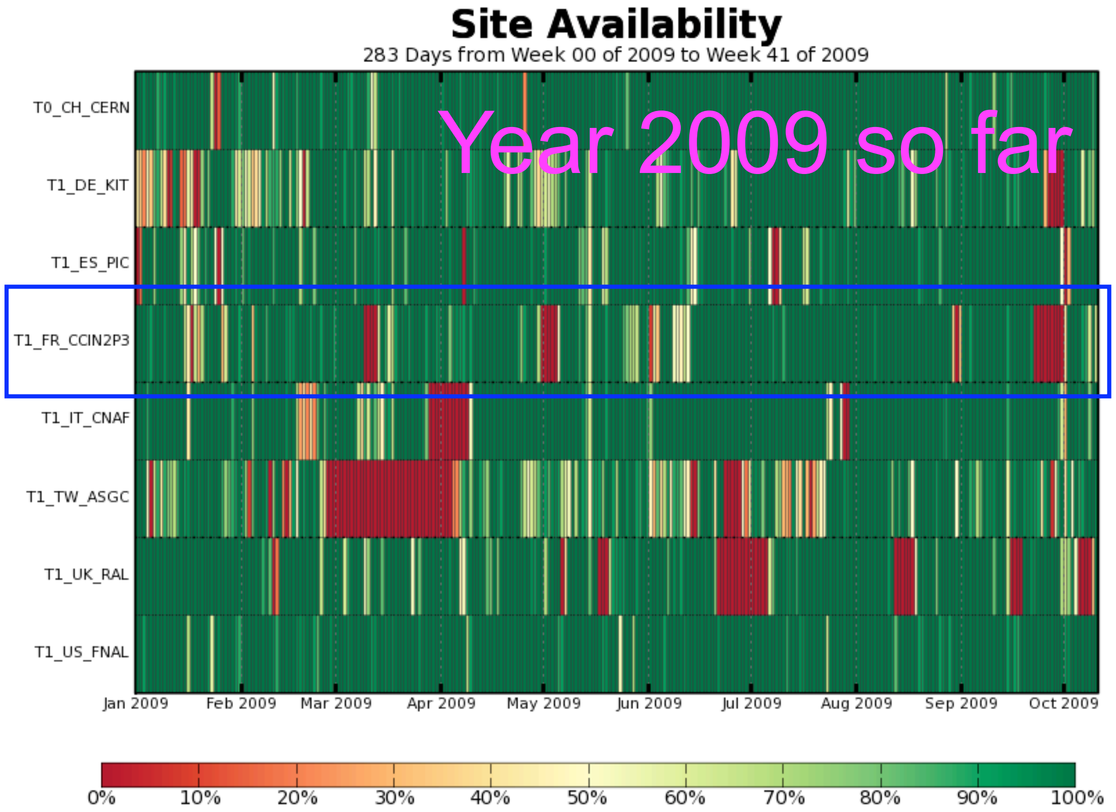

Year 2009 so far

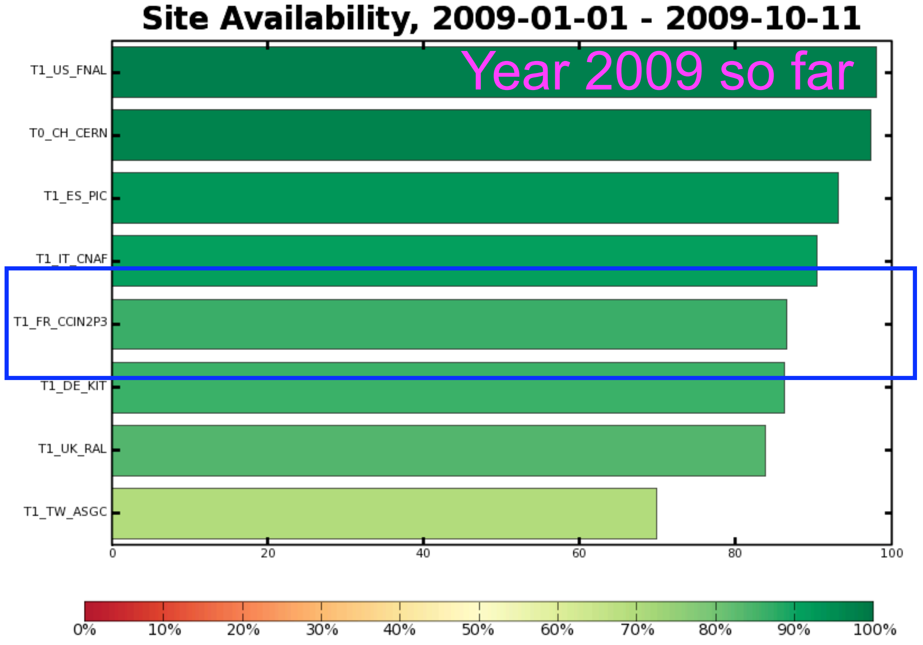

Year 2009 so far

CMS SiteReadiness

Global estimator in FacOps for the readiness of sites for daily operations

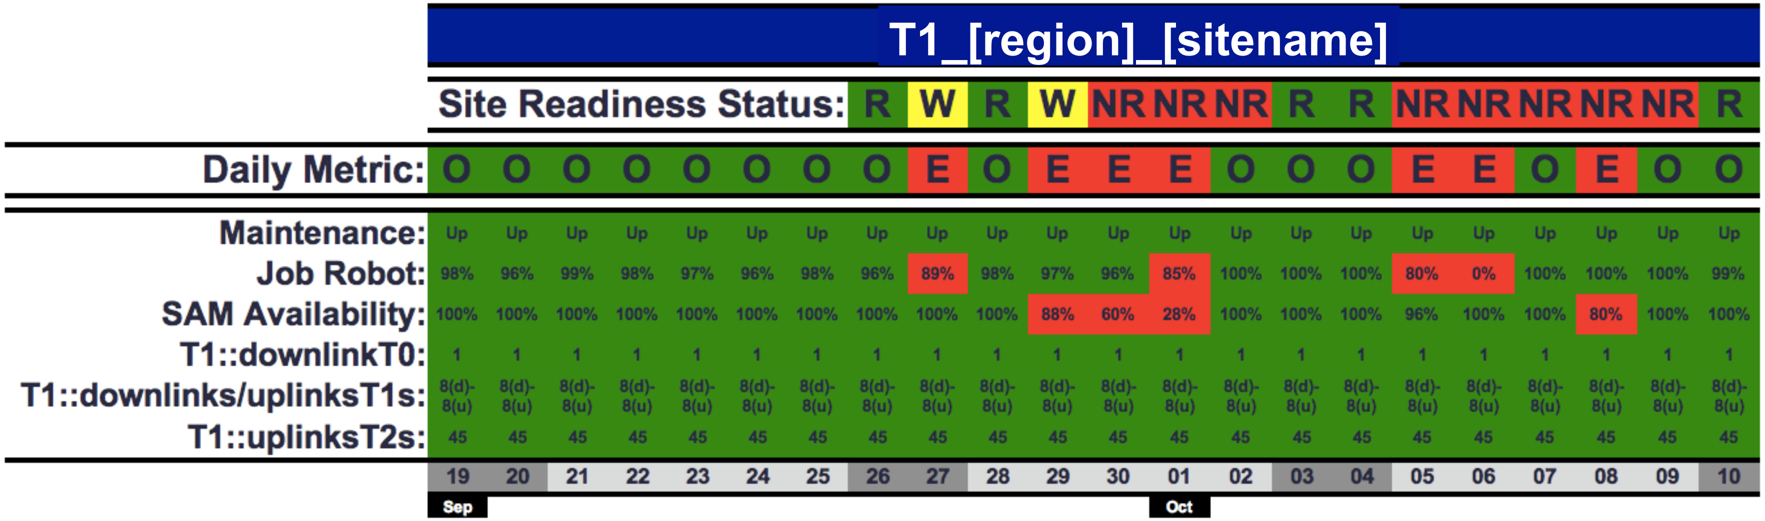

T1_[region]_[sitename]

Site Readiness Status:								R	W	R	W	NR	NR	NR	R	R	NR	NR	NR	NR	NR	R
Daily Metric:	O	O	O	O	O	O	O	O	E	O	E	E	E	O	O	O	E	E	O	E	O	O
Maintenance:	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up
Job Robot:	98%	96%	99%	98%	97%	96%	98%	96%	89%	98%	97%	96%	85%	100%	100%	100%	80%	0%	100%	100%	100%	99%
SAM Availability:	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	88%	60%	28%	100%	100%	100%	96%	100%	100%	80%	100%	100%
T1::downlinkT0:	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
T1::downlinks/uplinksT1s:	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)
T1::uplinksT2s:	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45	45
	19	20	21	22	23	24	25	26	27	28	29	30	01	02	03	04	05	06	07	08	09	10
	Sep												Oct									

https://twiki.cern.ch/twiki/bin/view/CMS/SiteCommRules

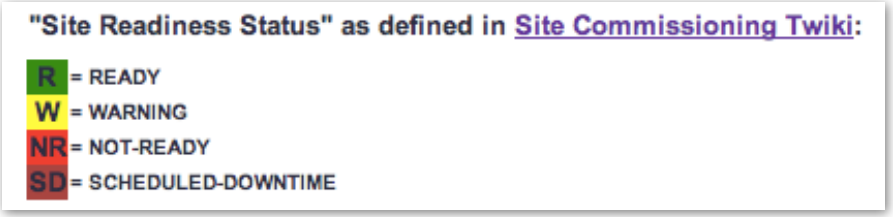

"Site Readiness Status" as defined in Site Commissioning Twiki:

R = READY
W = WARNING
NR = NOT-READY
SD = SCHEDULED-DOWNTIME

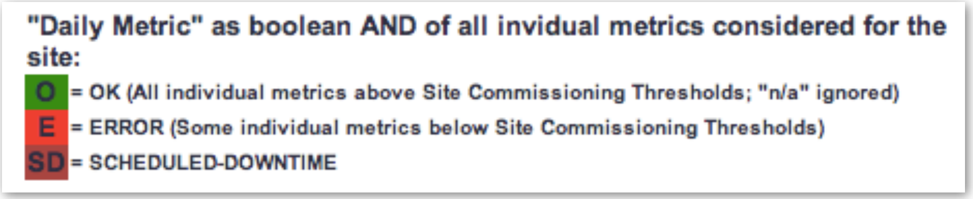

"Daily Metric" as boolean AND of all invidual metrics considered for the site:

O = OK (All individual metrics above Site Commissioning Thresholds; "n/a" ignored)
E = ERROR (Some individual metrics below Site Commissioning Thresholds)
SD = SCHEDULED-DOWNTIME

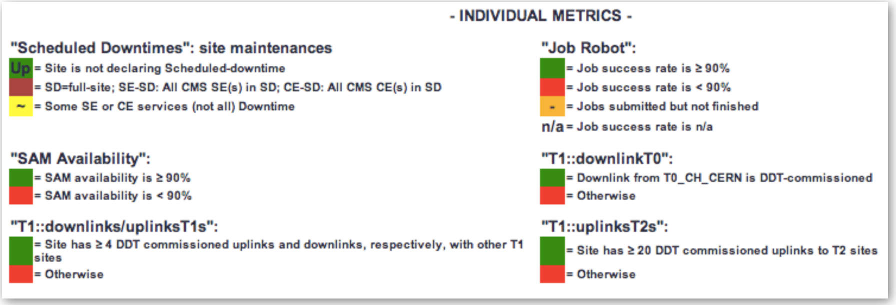

- INDIVIDUAL METRICS -

"Scheduled Downtimes": site maintenances
Up = Site is not declaring Scheduled-downtime
🟫 = SD=full-site; SE-SD: All CMS SE(s) in SD; CE-SD: All CMS CE(s) in SD
~ = Some SE or CE services (not all) Downtime

"SAM Availability":
🟩 = SAM availability is ≥ 90%
🟥 = SAM availability is < 90%

"T1::downlinks/uplinksT1s":
🟩 = Site has ≥ 4 DDT commissioned uplinks and downlinks, respectively, with other T1 sites
🟥 = Otherwise

"Job Robot":
🟩 = Job success rate is ≥ 90%
🟥 = Job success rate is < 90%
- = Jobs submitted but not finished
n/a = Job success rate is n/a

"T1::downlinkT0":
🟩 = Downlink from T0_CH_CERN is DDT-commissioned
🟥 = Otherwise

"T1::uplinksT2s":
🟩 = Site has ≥ 20 DDT commissioned uplinks to T2 sites
🟥 = Otherwise

From CMS Site Readiness metrics:

- Site availability: fraction of time all functional tests succeed
- JobRobot efficiency: fraction of successful “fake” analysis jobs
- Links: # of commissioned data transfer links

CMS SiteReadiness ranking for CMS T1's

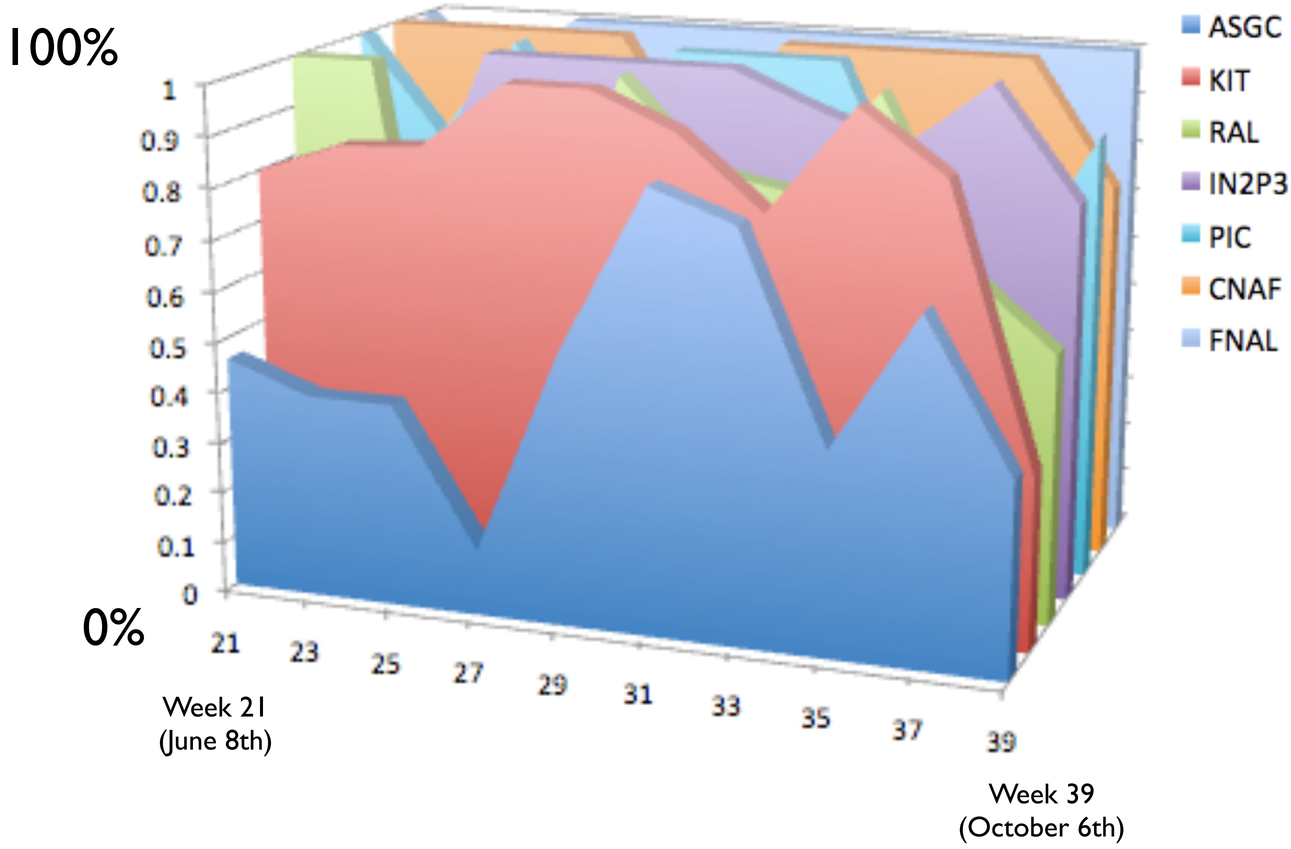

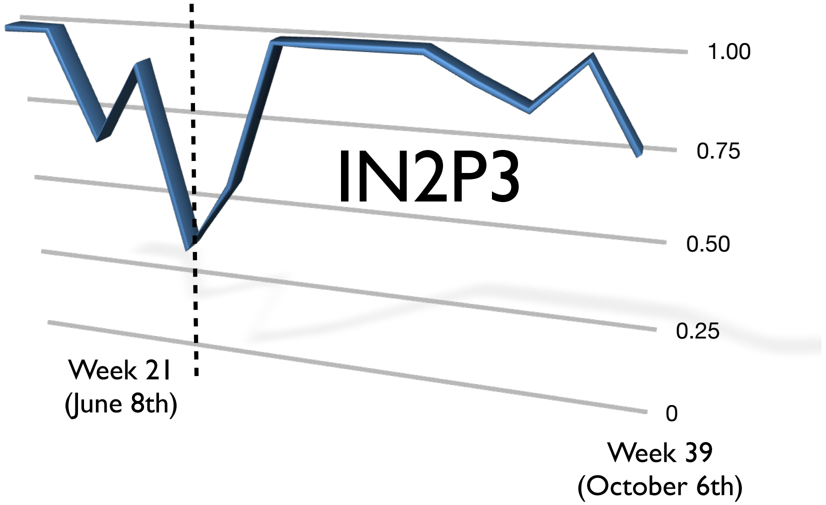

SiteReadiness goal for T1's: **90%**

Achieved averages in *Jun-Oct 2009*:

- {FNAL, CNAF} at {**99%**, **95%**}
- {PIC, IN2P3, KIT, RAL} at {**87%**, **86%**, **85%**, **73%**}
- ASGC at **50%**

WLCG SAM (ops) not the full picture

CMS-specific SAM not the full picture

SiteReadiness (even!) not the full picture

- Need high CPU eff, disk stability, MSS solidity and performance, ...

Readiness of sites: CMS requirements on Tiers [4/4]

Example: historical data on T1 and T2 sites

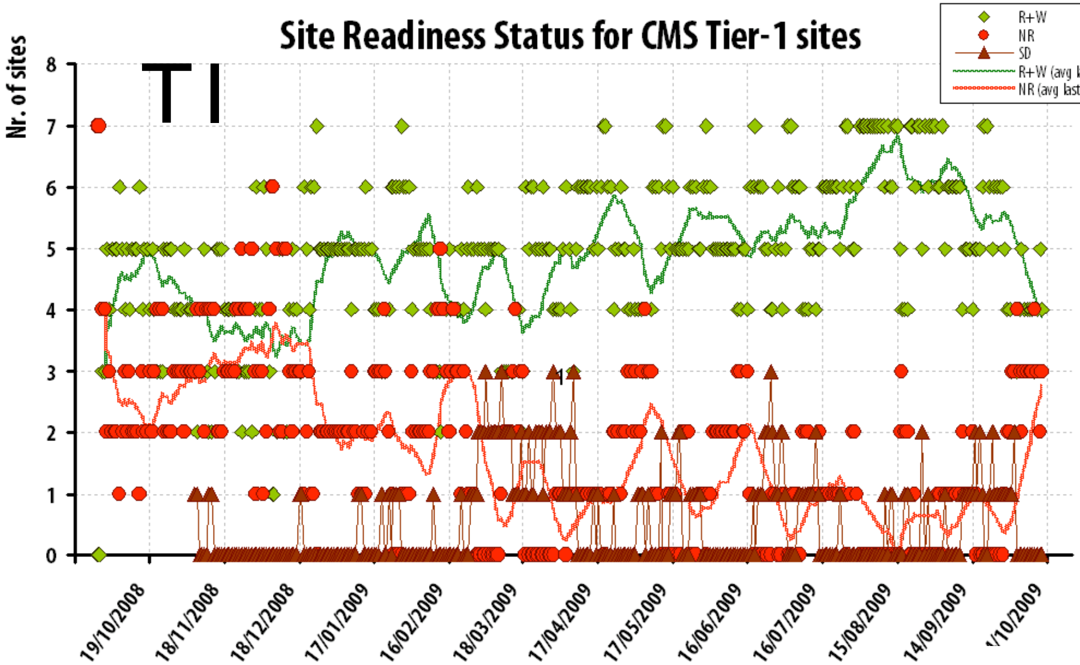

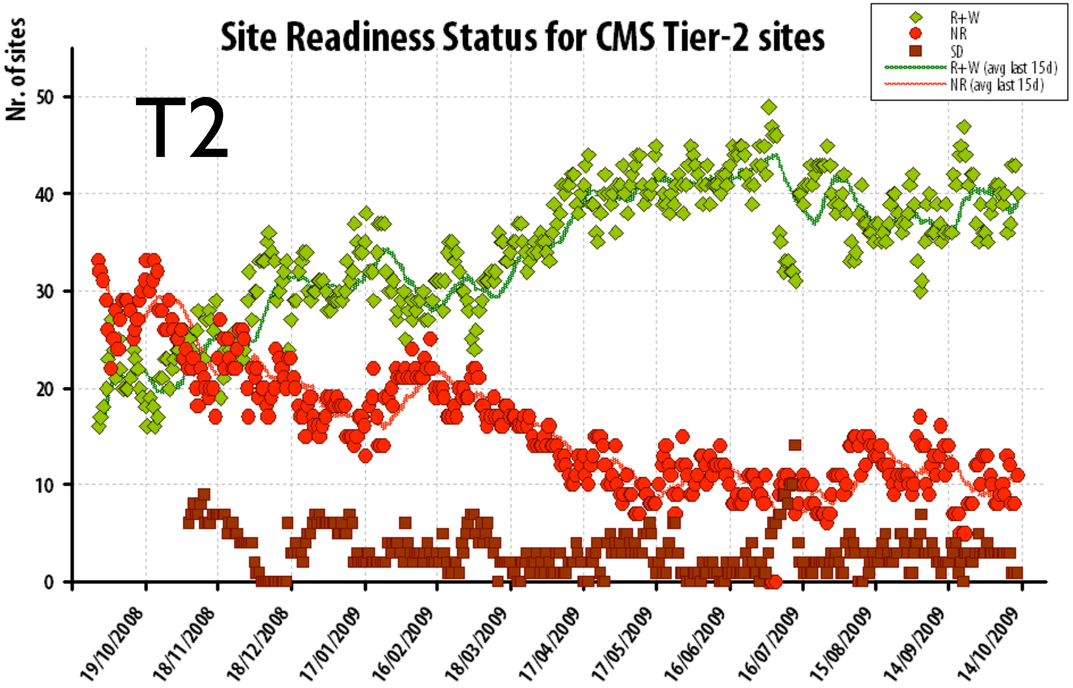

CMS SiteReadiness ranking for ***IN2P3***

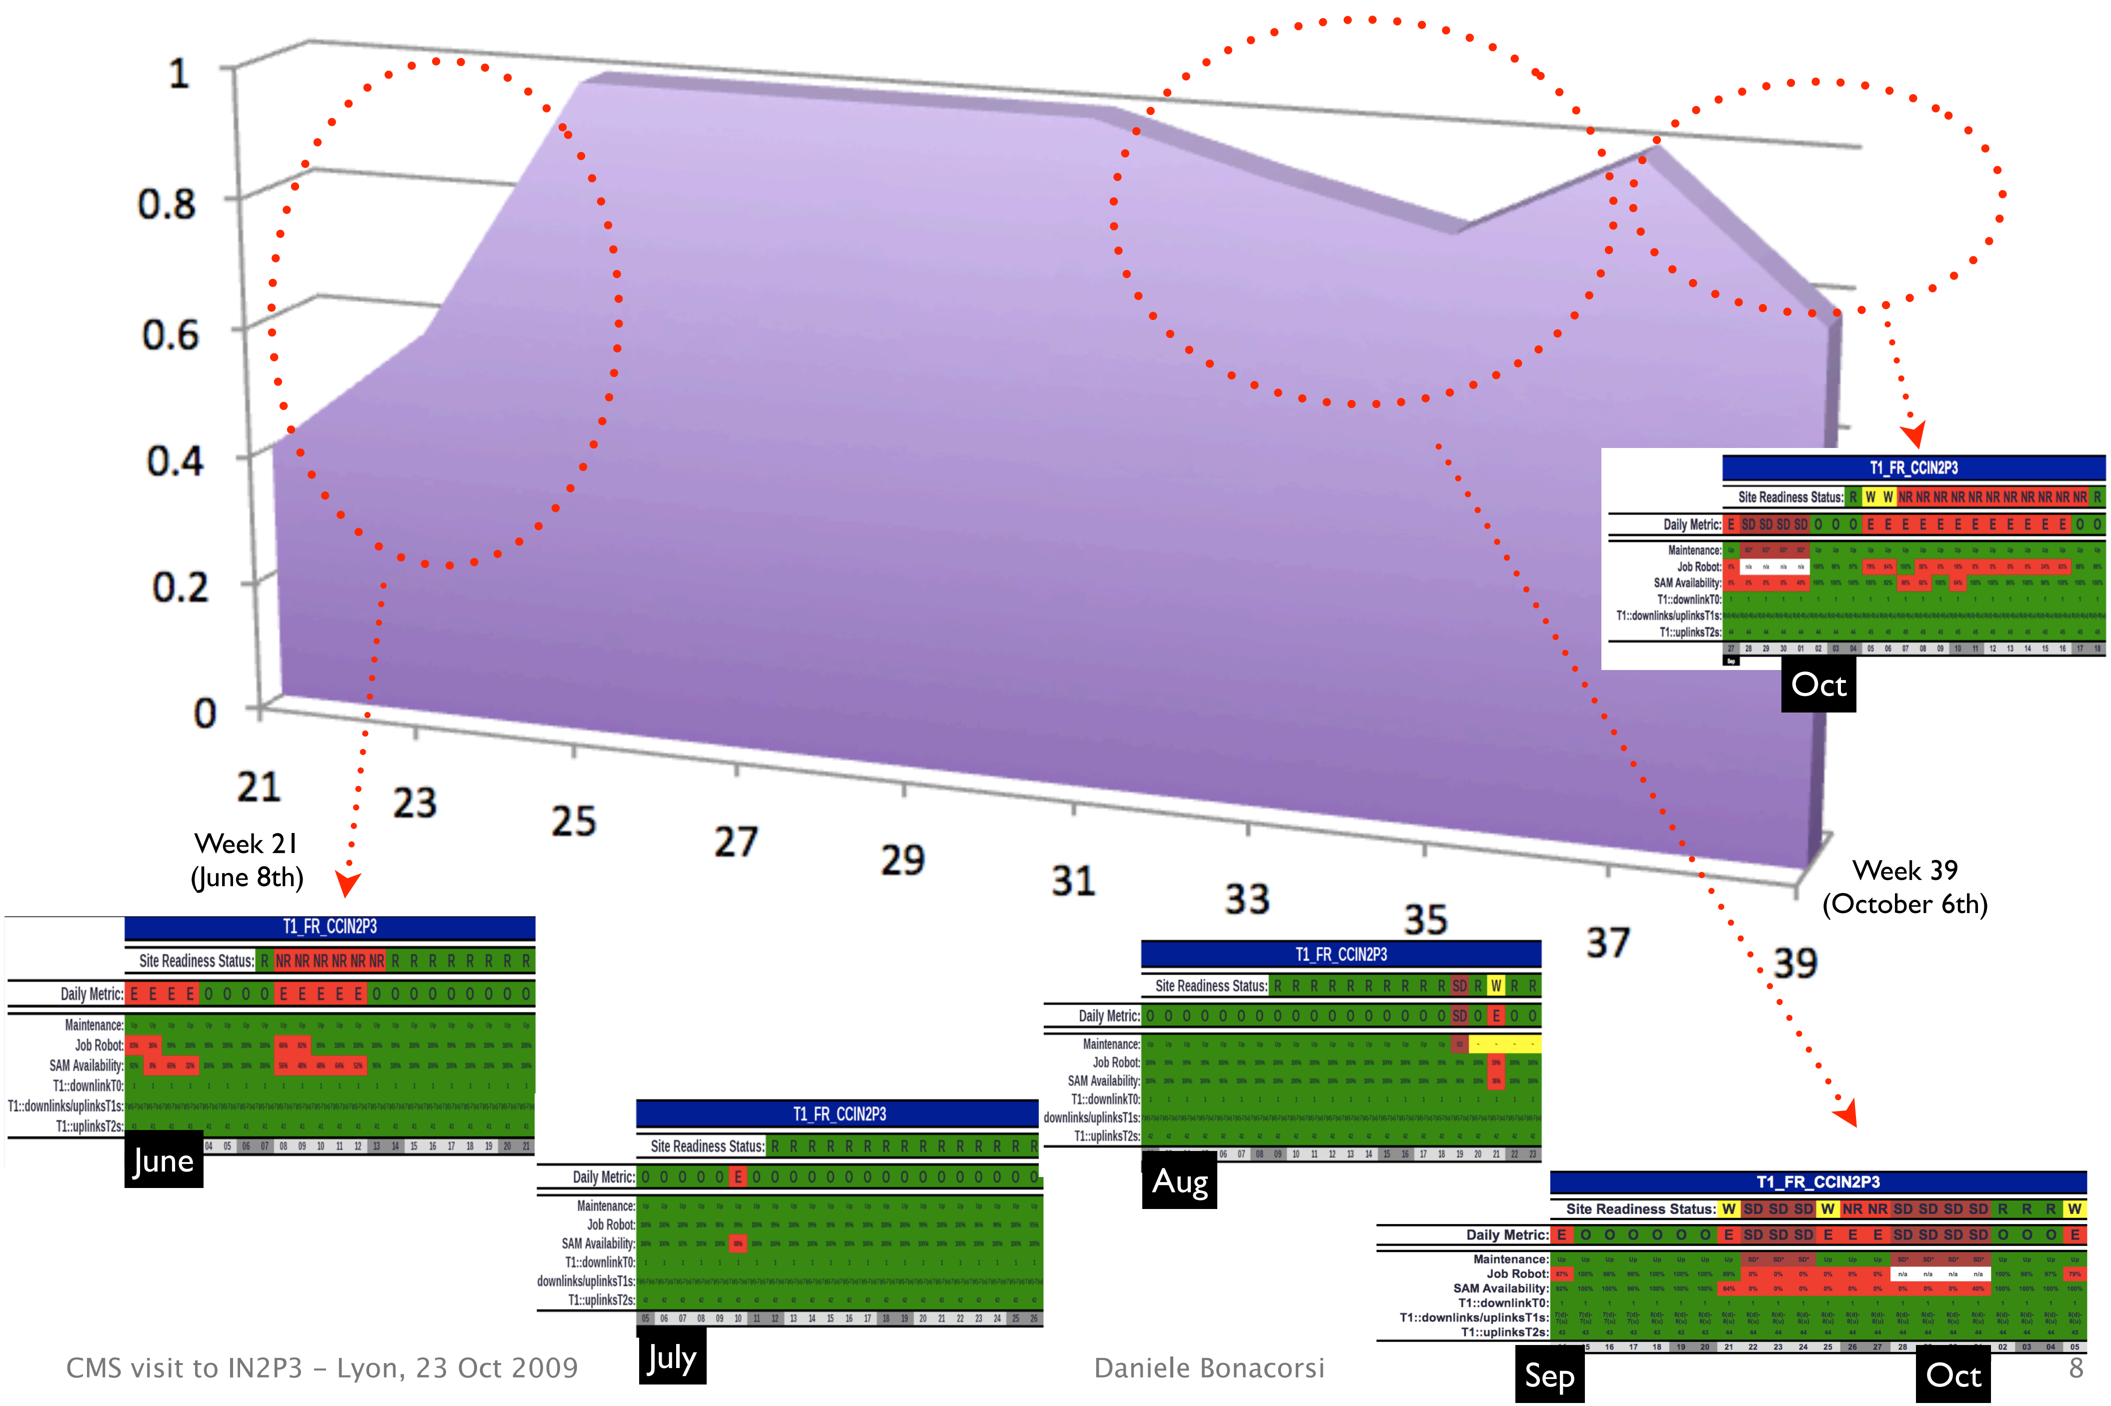

T1_FR_CCIN2P3 (June)

Site Readiness Status:								R	NR	NR	NR	NR	NR	NR	R	R	R	R	R	R	R	R
Daily Metric:	E	E	E	E	O	O	O	O	E	E	E	E	E	O	O	O	O	O	O	O	O	O
Maintenance:	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up
Job Robot:	83%	16%	99%	100%	98%	100%	100%	100%	66%	82%	99%	100%	100%	100%	99%	100%	100%	99%	100%	100%	100%	100%
SAM Availability:	92%	8%	60%	32%	100%	100%	100%	100%	56%	48%	48%	64%	52%	96%	100%	100%	100%	100%	100%	100%	100%	100%
T1::downlinkT0:	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
T1::downlinks/uplinksT1s:	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)
T1::uplinksT2s:	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41
					04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21

June

T1_FR_CCIN2P3 (July)

Site Readiness Status:								R	R	R	R	R	R	R	R	R	R	R	R	R	R	R
Daily Metric:	O	O	O	O	O	E	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O
Maintenance:	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up
Job Robot:	100%	100%	100%	100%	98%	99%	100%	99%	100%	99%	99%	99%	100%	100%	100%	99%	100%	100%	96%	98%	100%	95%
SAM Availability:	100%	100%	92%	100%	100%	88%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
T1::downlinkT0:	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
downlinks/uplinksT1s:	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)
T1::uplinksT2s:	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42
	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26

July

T1_FR_CCIN2P3 (Aug)

Site Readiness Status:								R	R	R	R	R	R	R	R	R	R	SD	R	W	R	R
Daily Metric:	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	O	SD	O	E	O	O
Maintenance:	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	SD	-	-	-	-
Job Robot:	100%	99%	99%	98%	100%	99%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	99%	99%	100%	59%	100%	100%
SAM Availability:	100%	100%	100%	100%	96%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	96%	100%	36%	100%	100%
T1::downlinkT0:	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
downlinks/uplinksT1s:	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)
T1::uplinksT2s:	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42
					06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23

Aug

T1_FR_CCIN2P3 (Sep)

Site Readiness Status:								W	SD	SD	SD	W	NR	NR	SD	SD	SD	SD	R	R	R	W
Daily Metric:	E	O	O	O	O	O	O	E	SD	SD	SD	E	E	E	SD	SD	SD	SD	O	O	O	E
Maintenance:	Up	Up	Up	Up	Up	Up	Up	Up	SD*	SD*	SD*	Up	Up	Up	SD*	SD*	SD*	SD*	Up	Up	Up	Up
Job Robot:	87%	100%	98%	99%	100%	100%	100%	99%	0%	0%	0%	0%	0%	0%	n/a	n/a	n/a	n/a	100%	98%	97%	79%
SAM Availability:	92%	100%	100%	96%	100%	100%	100%	84%	0%	0%	0%	0%	0%	0%	0%	0%	0%	40%	100%	100%	100%	100%
T1::downlinkT0:	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
T1::downlinks/uplinksT1s:	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)
T1::uplinksT2s:	43	43	43	43	43	43	43	44	44	44	44	44	44	44	44	44	44	44	44	44	44	45
		15	16	17	18	19	20	21	22	23	24	25	26	27	28				02	03	04	05

Sep

Oct

T1_FR_CCIN2P3 (Oct)

Site Readiness Status:								R	W	W	NR	NR	NR	NR	NR	NR	NR	NR	NR	NR	NR	R
Daily Metric:	E	SD	SD	SD	SD	O	O	O	E	E	E	E	E	E	E	E	E	E	E	E	O	O
Maintenance:	Up	SD*	SD*	SD*	SD*	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up
Job Robot:	0%	n/a	n/a	n/a	n/a	100%	98%	97%	79%	84%	100%	50%	0%	16%	0%	0%	0%	0%	24%	83%	99%	99%
SAM Availability:	0%	0%	0%	0%	40%	100%	100%	100%	100%	92%	88%	60%	100%	64%	100%	100%	96%	100%	96%	100%	100%	100%
T1::downlinkT0:	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
T1::downlinks/uplinksT1s:	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)	8(d)-8(u)
T1::uplinksT2s:	44	44	44	44	44	44	44	44	45	45	45	45	45	45	45	45	45	45	45	45	45	45
	27	28	29	30	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18
	Sep																					

Oct

SiteReadiness breakdown for ***IN2P3***

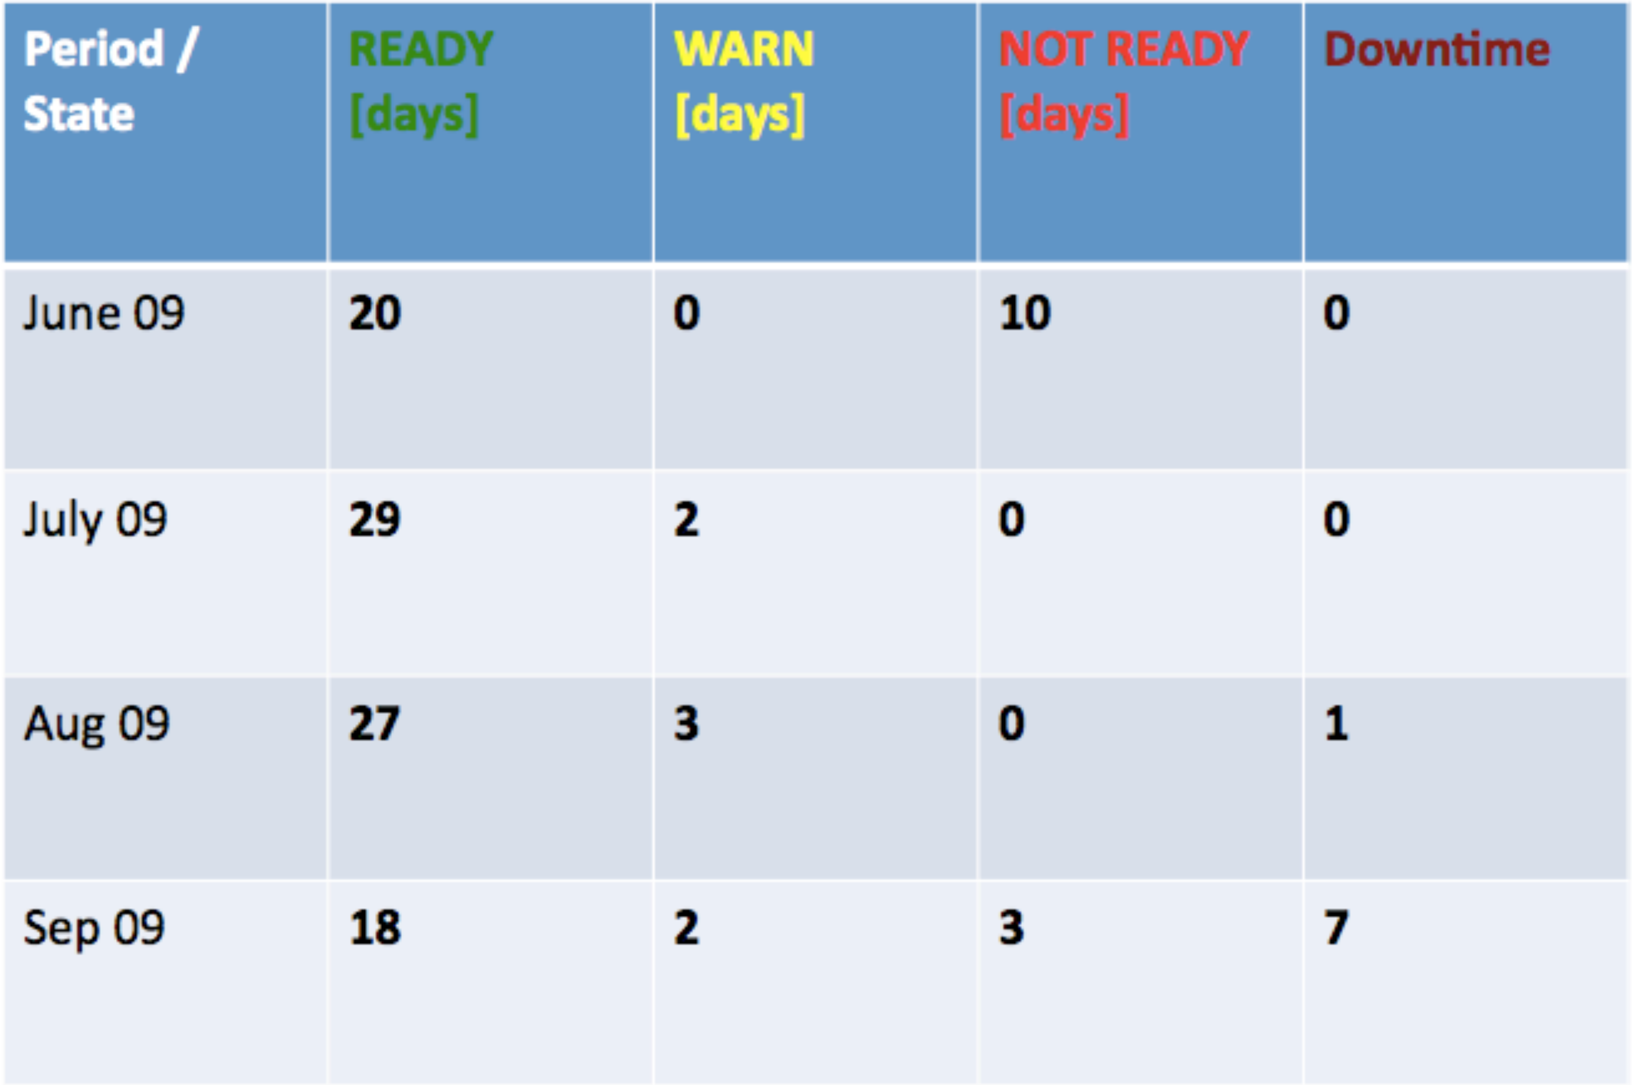

Period / State	READY [days]	WARN [days]	NOT READY [days]	Downtime
June 09	20	0	10	0
July 09	29	2	0	0
Aug 09	27	3	0	1
Sep 09	18	2	3	7

NOTE: SiteReadiness has lately suffered from SSB instabilities when tracing scheduled downtimes. The September IN2P3 downtime was corrected on SiteReadiness tables as announced here:

- https://hypernews.cern.ch/HyperNews/CMS/get/sc4/1969.html

Transfer rates: T0 -> IN2P3

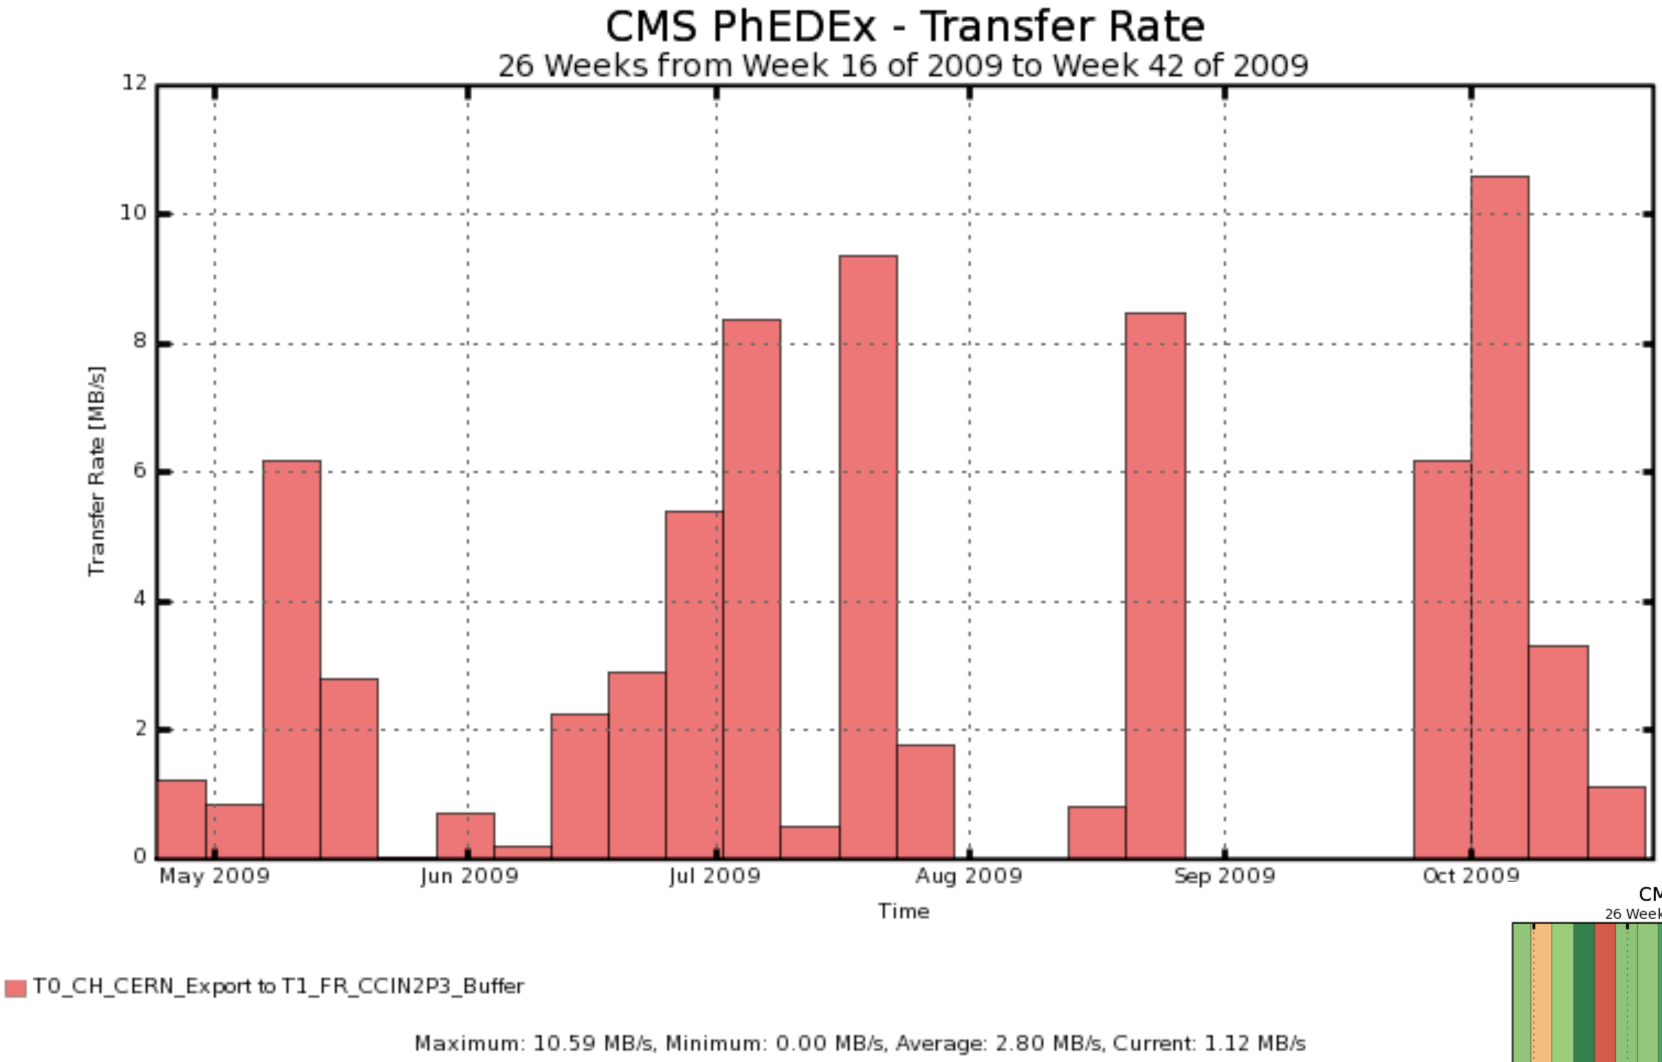

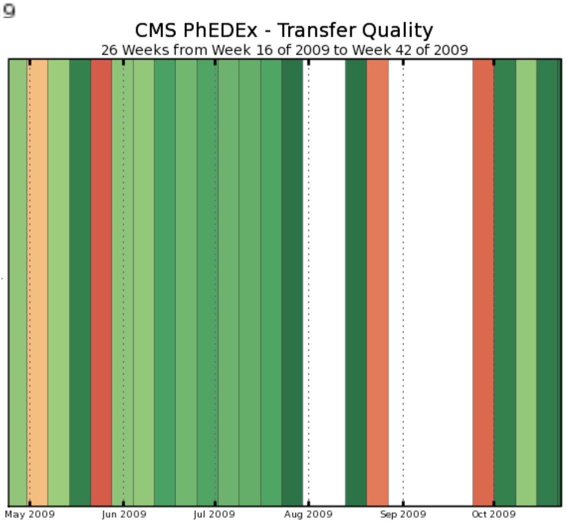

Little activity in the PhEDEx /Prod instance

- few datasets from T0 assigned to IN2P3 as custodial site...

Transfer rates: IN2P3 -> *

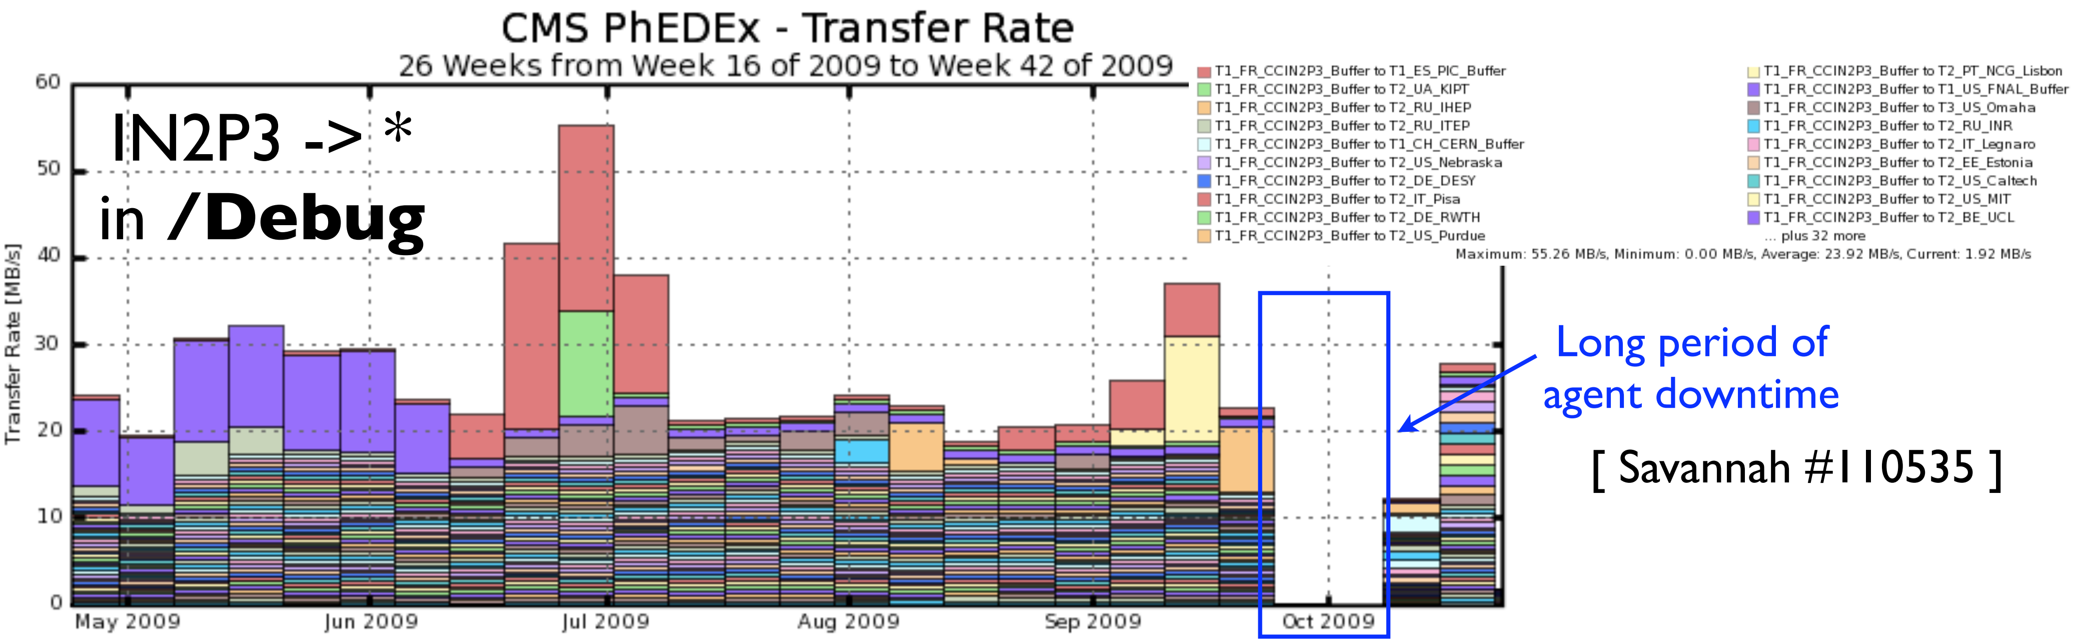

IN2P3 -> *
in **/Debug**

Long period of
agent downtime

[Savannah #110535]

IN2P3 -> *
in **/Prod**

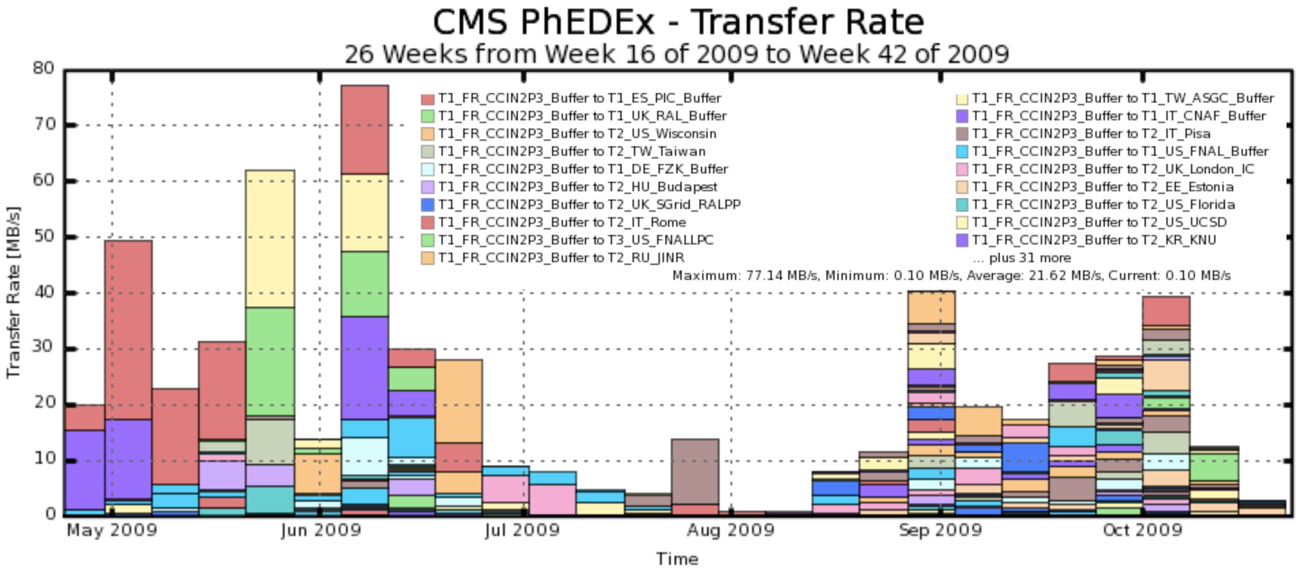

Transfer quality: IN2P3 -> *

IN2P3 -> * in **/Debug**

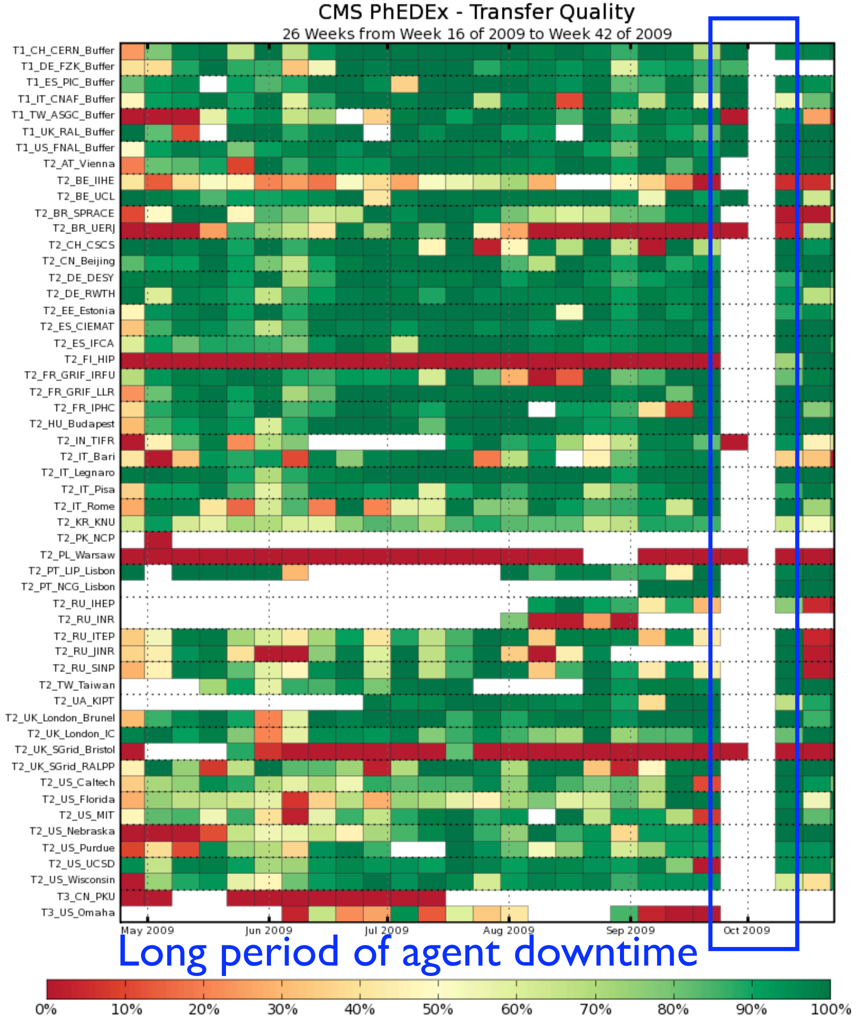

Long period of agent downtime

IN2P3 -> * in **/Prod**

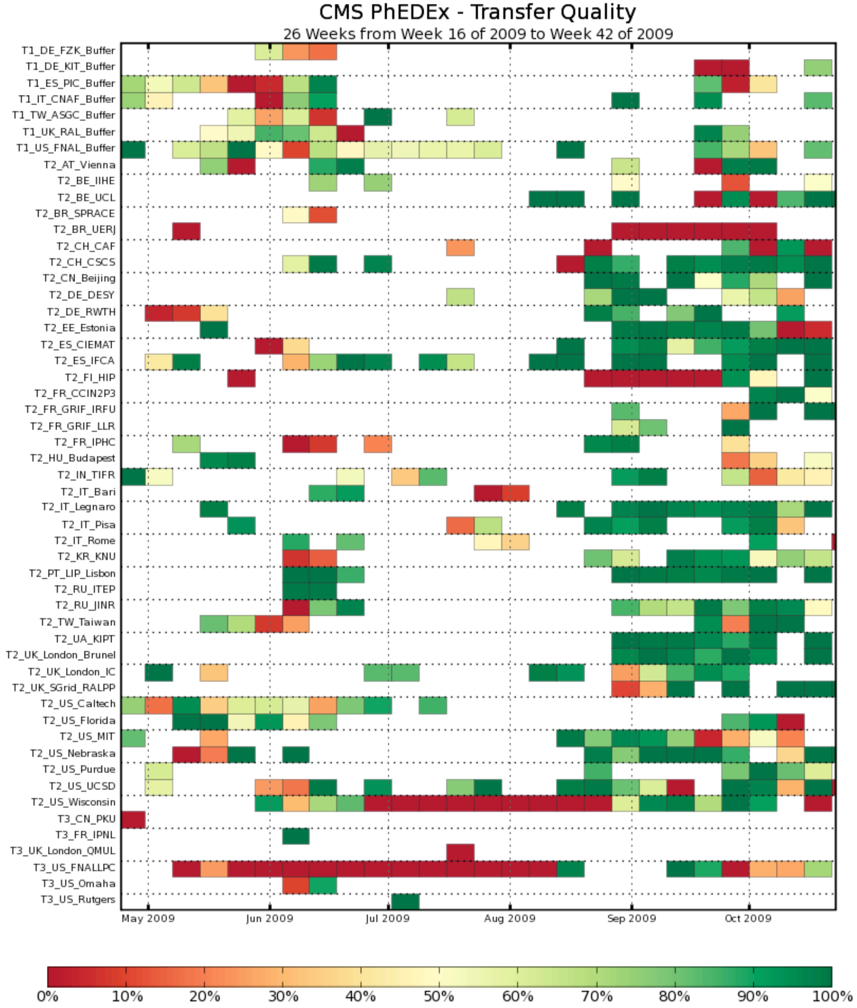

Generally OK in /Debug

- apart from the agent downtime in late Sept

Not too bad in /Prod

- Most frequent problem is transfer expirations due to FTS channel congestion - these are invisible in the plots...

Transfer rates: * -> IN2P3

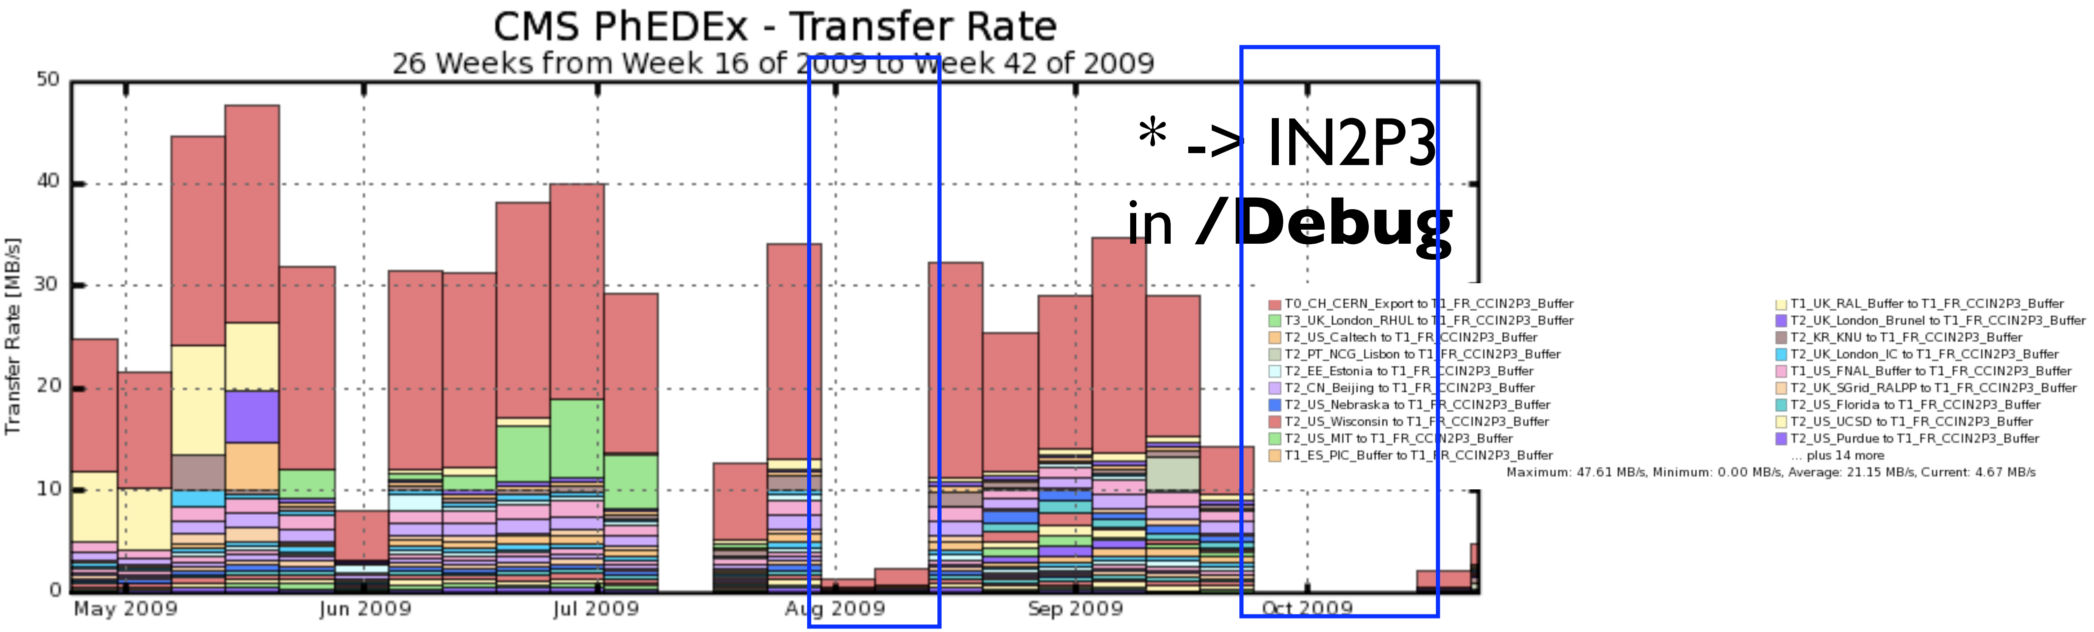

* -> IN2P3
in **/Debug**

* -> IN2P3
in **/Prod**

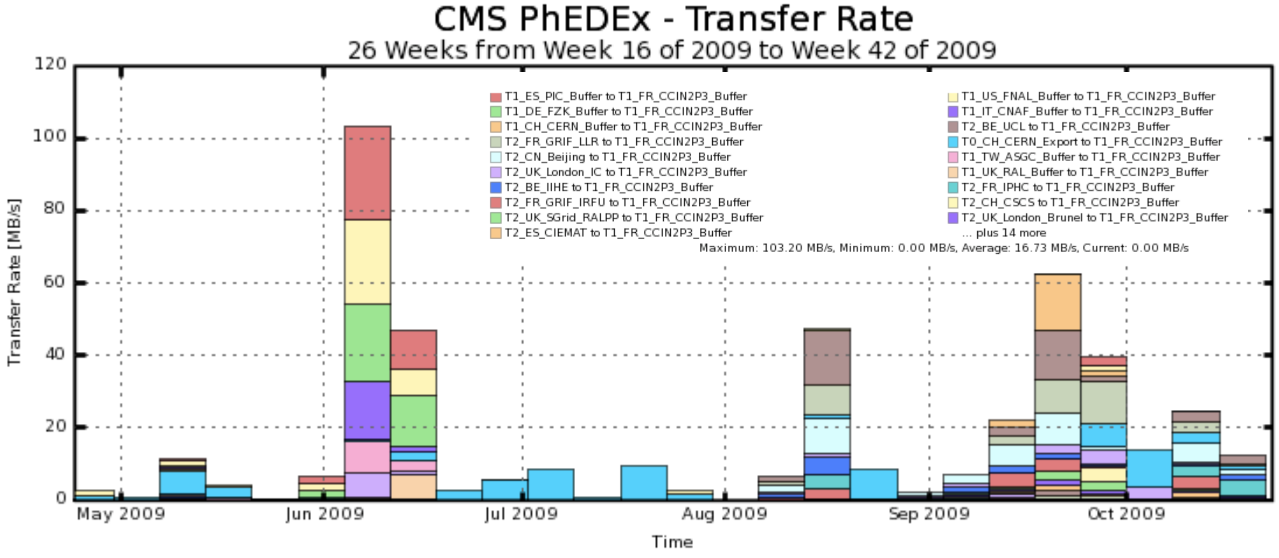

Transfer quality: * -> IN2P3

* -> IN2P3 in **/Debug**

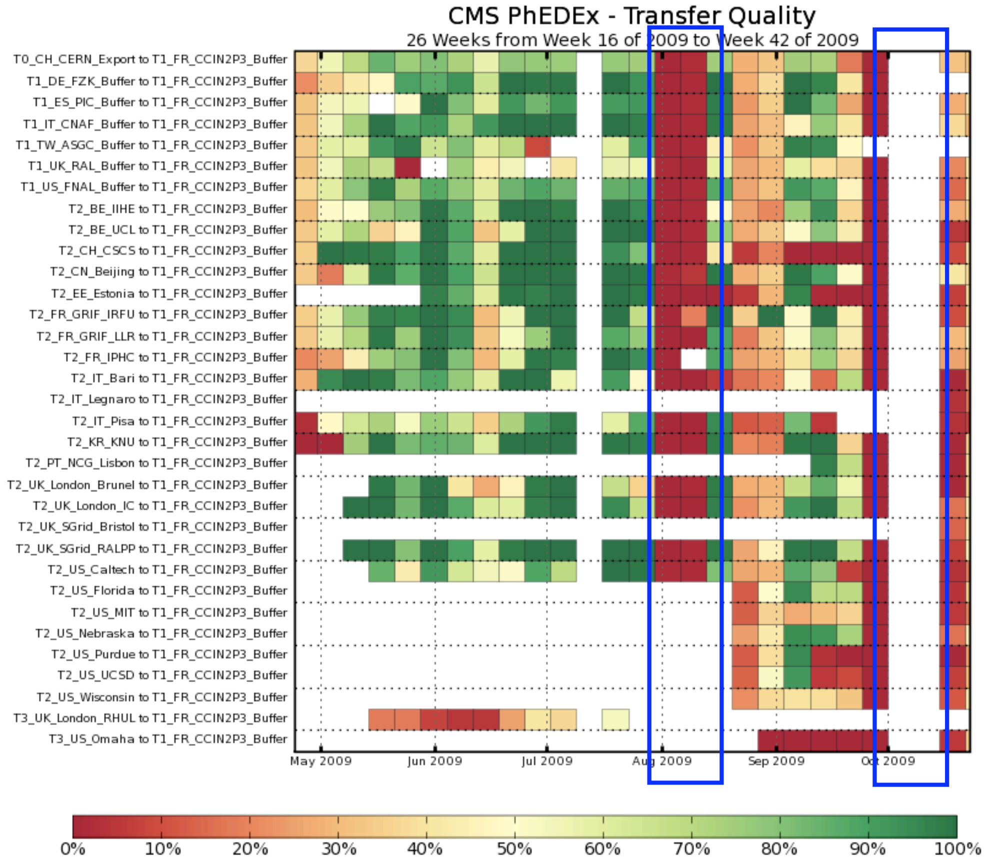

* -> IN2P3 in **/Prod**

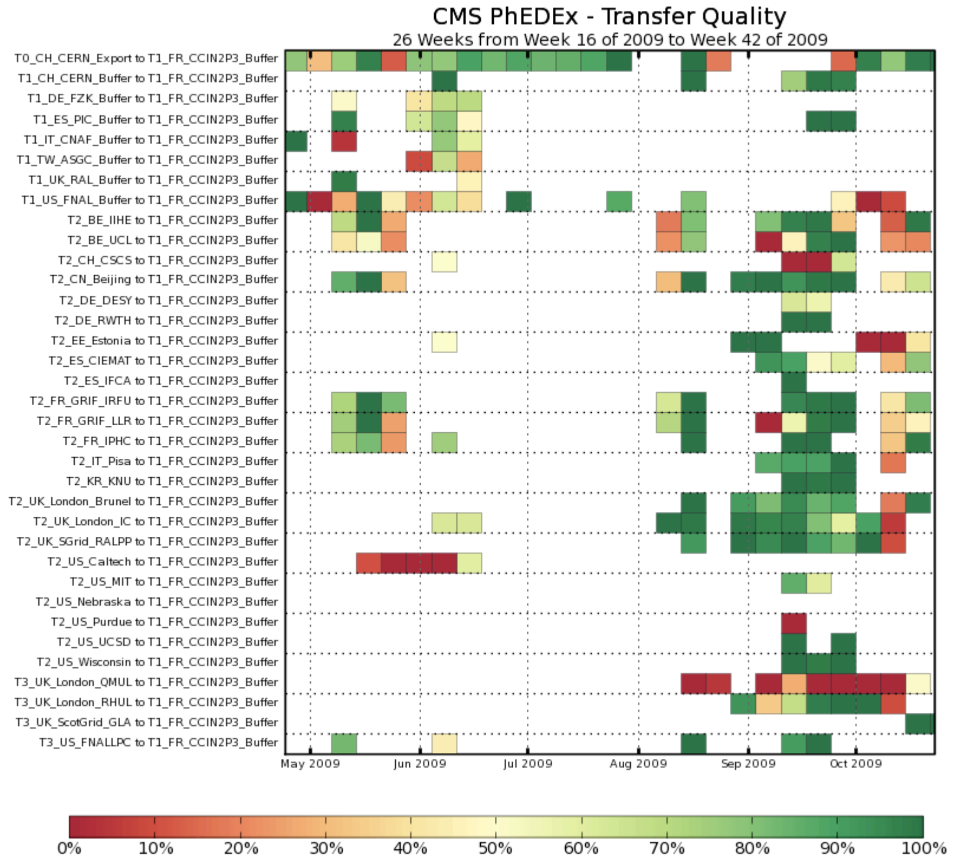

The import in the /Debug instance are more frequently in overall bad health

- agents down for long periods of time
- Relatively bad transfer quality in imports since summer

A large source of errors is "*Already have 1 record(s) with pnfsPath=[...]"

- probably a cleanup of the LoadTest target area would improve things...

Transfers: IN2P3 <-> T1's [1/2]

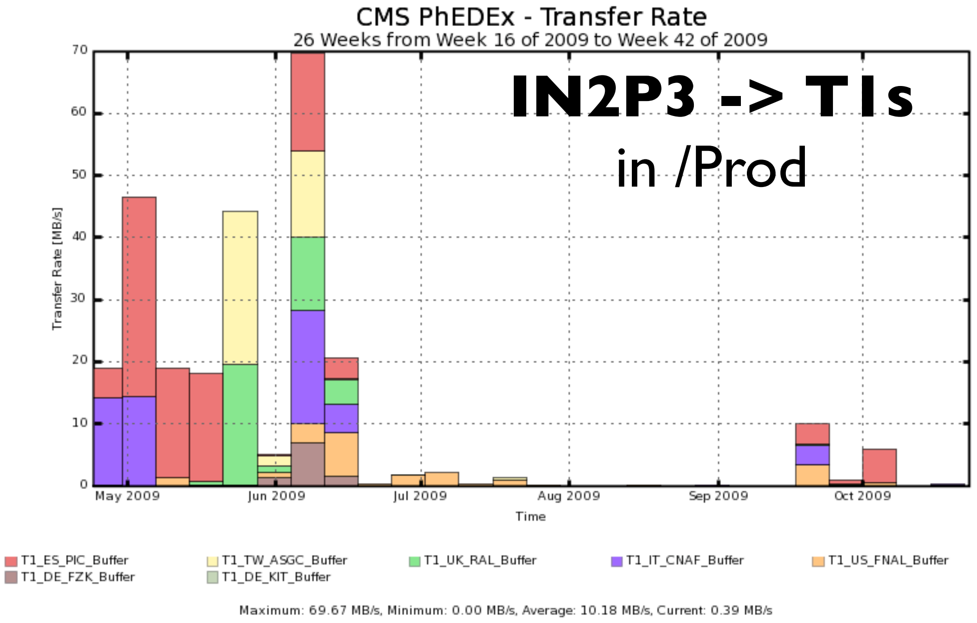

IN2P3 -> T1s
in /Prod

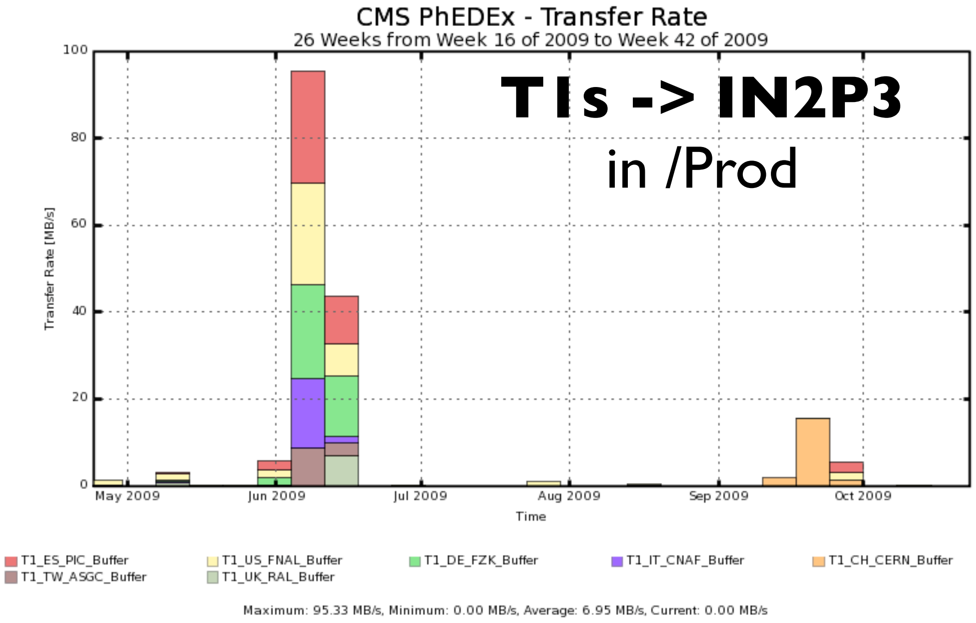

T1s -> IN2P3
in /Prod

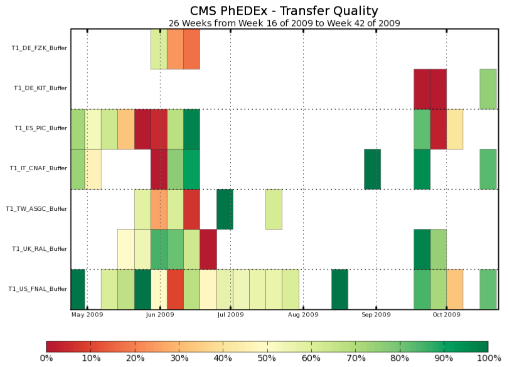

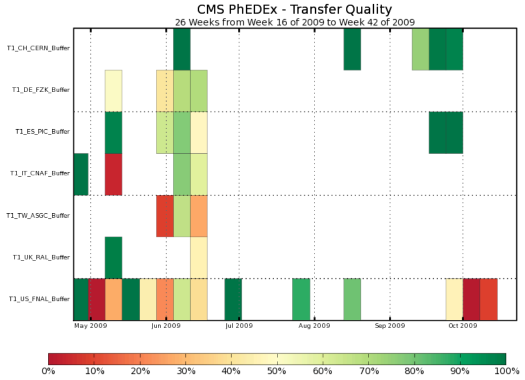

Almost no activity outside STEP09

During STEP09, very good rates *(more in the back-up slides)*

✦ Targets (assuming no rerouting in PhEDEx) were 185 MB/s in, 105 MB/s out - exceeded in one day

Transfers: IN2P3 <-> T1's [2/2]

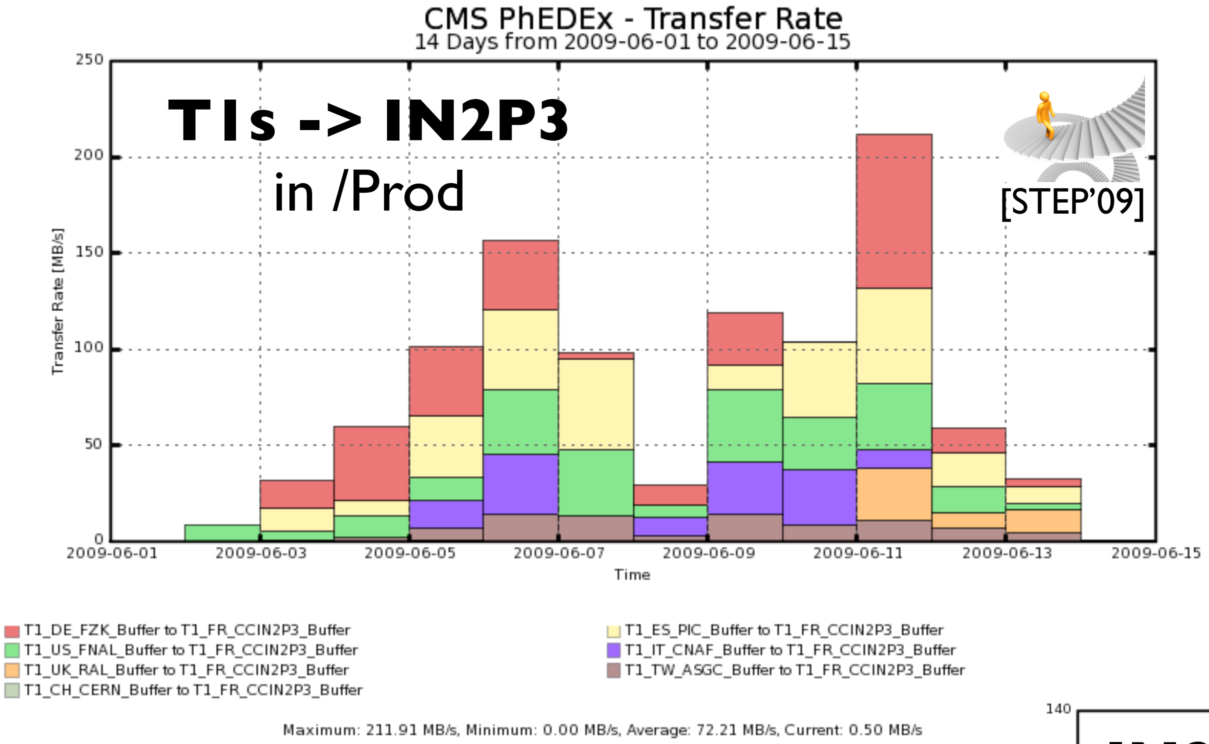

T1s -> IN2P3
in /Prod

[STEP'09]

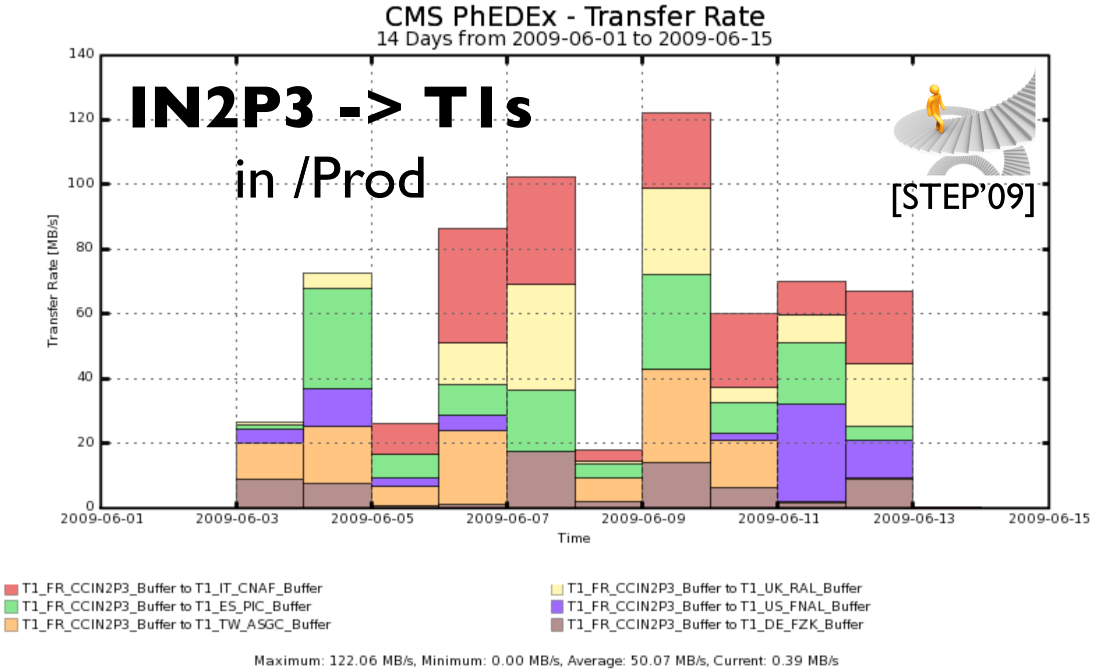

IN2P3 -> T1s
in /Prod

[STEP'09]

Transfers: IN2P3 -> T2's

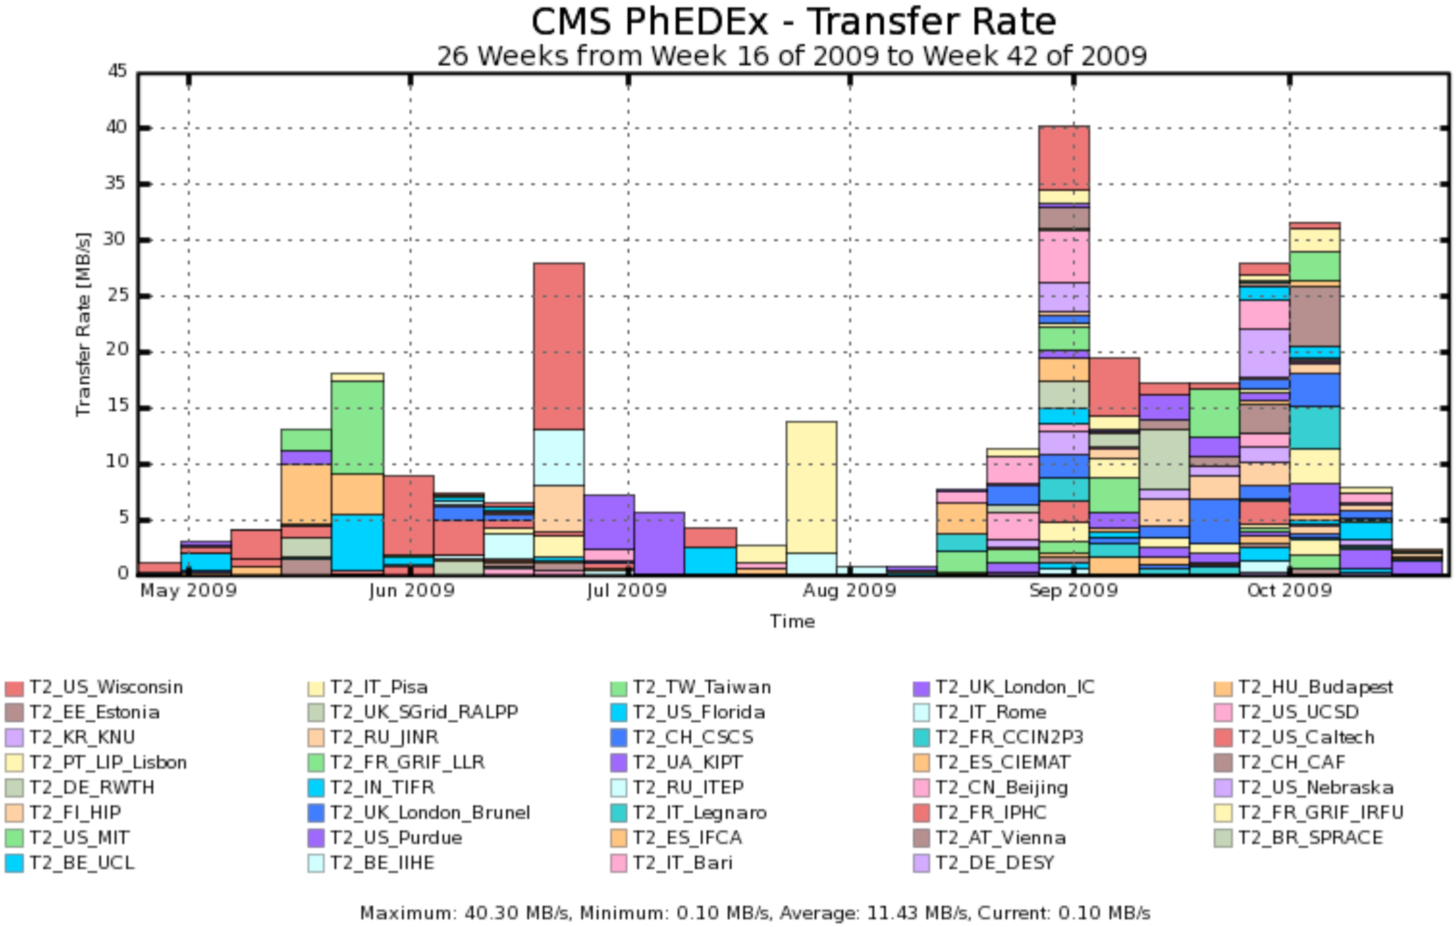

Constant activity since September

Target in CCRC08 was ~80 MB/s averaged over a long period

- still below, despite lots of data custodial at T1_FR_CCIN2P3 (~600 TB).

DataOps scheduled a round of 'DDT-style' tests IN2P3->T2_* last week to measure export rates

Transfers: T2’s -> IN2P3 and other T1’s

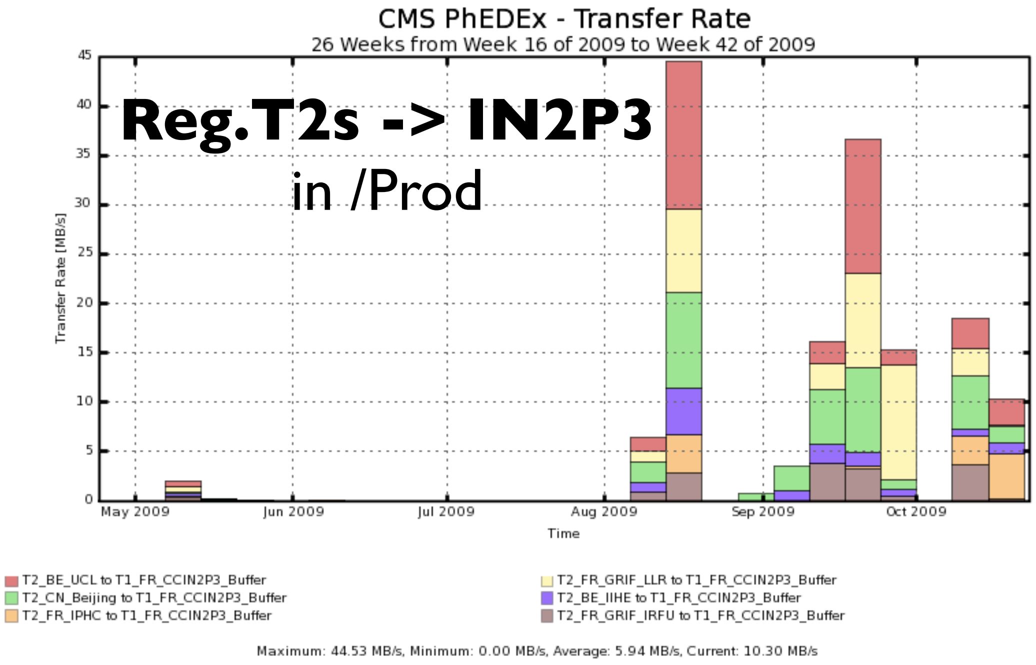

Reg.T2s -> IN2P3
in /Prod

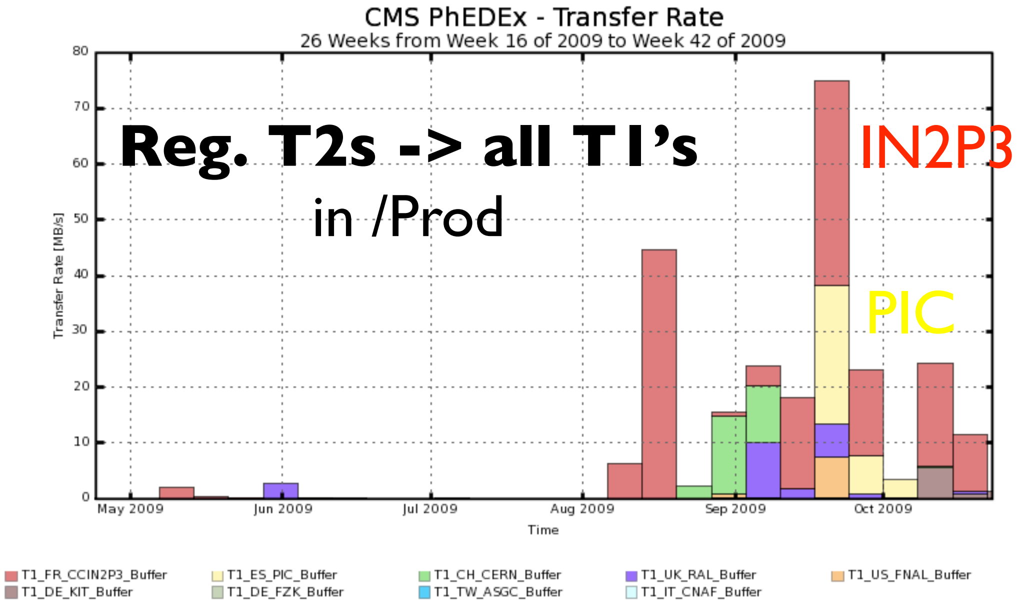

Reg. T2s -> all T1’s
in /Prod

IN2P3

PIC

Assuming CCRC’08 targets:

- the MC production rate from T2s in the France/Belgium/China region averaged over a long period should be 7.4 MB/s

We are way higher than that after the summer

Link commissioning status

http://lhcweb.pic.es/cms/CommLinksReports/CommissionedLinks_Sites.html

T1_FR_IN2P3:

- Export links commissioned, except for some in T2_RU/T2_TR region (rate limitations)
  - http://cmsweb.cern.ch/phedex/prod/Components::Links?from_filter=T1_FR&andor=and&to_filter=.*&Update=Update#
- Import links OK, also some non-regional links (not all of them, though)
  - http://cmsweb.cern.ch/phedex/prod/Components::Links?from_filter=&andor=and&to_filter=T1_FR&Update=Update#

T2's in France/Belgium/China region:

- All fully equipped with downlinks and with many backup uplinks
- T2_FR_CCIN2P3 exports still inactive during namespace migration
- Remarkably, T2_FR_GRIF_LLR also has lots of T2<->T2 links

Ops efficiency and Communication

A good coverage of CMS Ops includes:

- ✦ Fulfill your site contact responsibilities
  - Good summary in DataOps slides (next talk)
- ✦ Attend regularly the Ops weekly meetings
  - Provide the brief weekly report every Monday to FacOps
  - Come prepared and discuss current issues on SAM, JR, ... in full depth
  - Give feedback to DataOps on production activities
- ✦ Give complete and precise answers to questions by FacOps and DataOps
  - Meetings, HN, private communications, ...
- ✦ Ask questions yourself !

Savannah somehow gives a feeling of the rate of issues notifications

- ✦ No Savannah gets opened if a problem is monitored, seen, fixed by CMS contacts onsite before any operator / shifter / user sees it
  - http://snipurl.com/savannah-in2p3
    - IN2P3 102, CNAF 84, ASGC 79, FNAL 74, RAL 48, PIC 30, [KIT 8 - before: FZK, no history]

We strongly rely on CMS contacts at T1 sites for efficient operations

Back-up

CMSSW deployment

CMSSW installed via Grid job on EGEE and OSG sites

- ✦ Basic strategy: use RPM (with apt-get) in CMS SW area

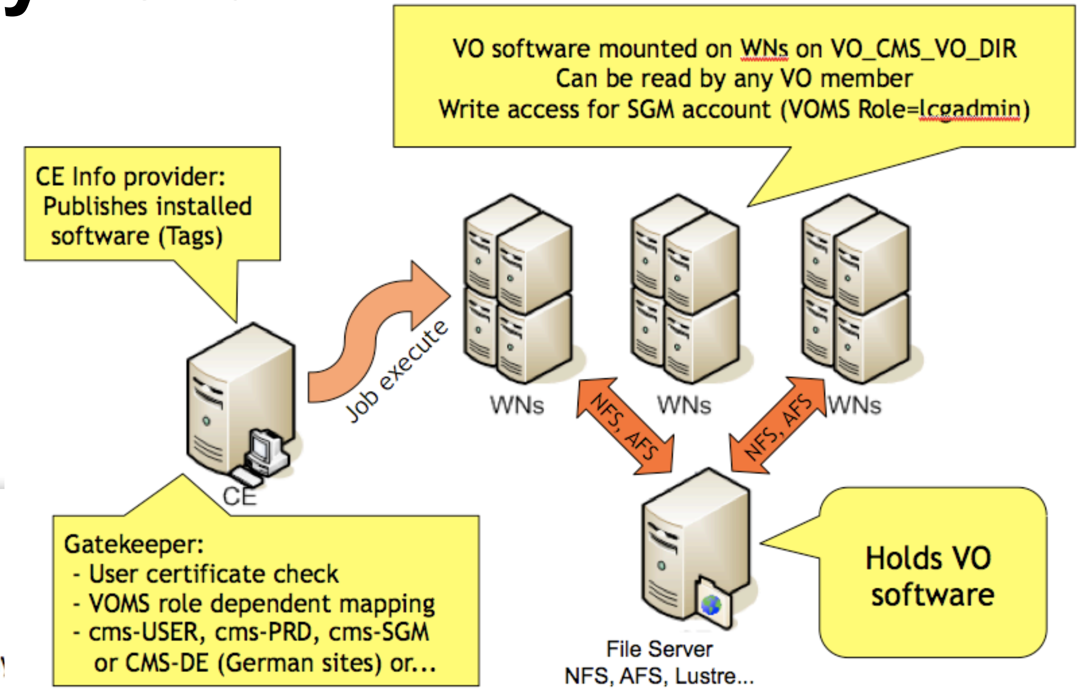

VO software mounted on WNs on VO_CMS_VO_DIR
Can be read by any VO member
Write access for SGM account (VOMS Role=lcgadmin)
CE Info provider: Publishes installed software (Tags)
Job execute
WNs
WNs
WNs
NFS, AFS
NFS, AFS
CE
Gatekeeper:
- User certificate check
- VOMS role dependent mapping
- cms-USER, cms-PRD, cms-SGM or CMS-DE (German sites) or...
Holds VO software
File Server
NFS, AFS, Lustre...

CMSSW_314 deployment

- release announcement on Saturday Oct 3rd at 13h22
  - first installation jobs submitted at 13h41
- status: CMSSW_314 release deployed and ready for Oct-X start-up on day
  - EGEE: submitted to 51 Computing Elements (CE), 44 were DONE after few hrs (see plot); started to follow up on tails over the weekend already
  - OSG: release not tagged into the tag collector, so installed manually; smoothly and quickly completed in most OSG T2/T3

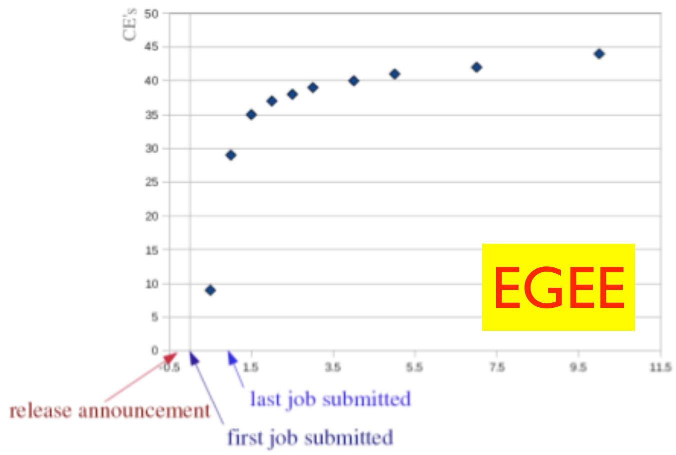

Credits:
CMSSW deployment teams in Facilities Ops:
Bockjoo Kim,
Joris Maes,
Lukas Vanelderen,
Petra Van Mulders,
Ilaria Villella,
Christoph Wissing

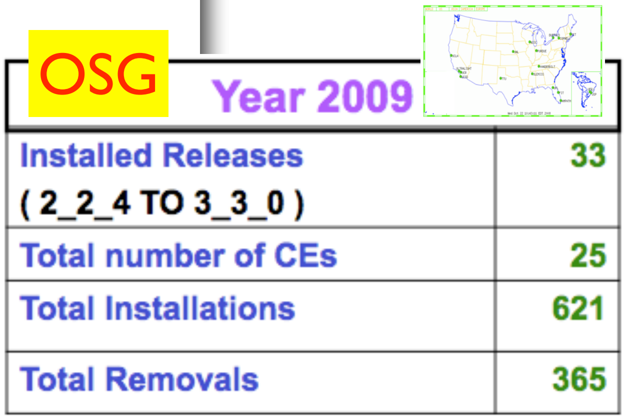

Year 2009	
Installed Releases (2_2_4 TO 3_3_0)	33
Total number of CEs	25
Total Installations	621
Total Removals	365

Credits: Bockjoo Kim

On EGEE and OSG:

- ✦ CMSSW releases get routinely installed smoothly in most sites within few hrs from the release

Ticketing systems [1/2]

GGUS

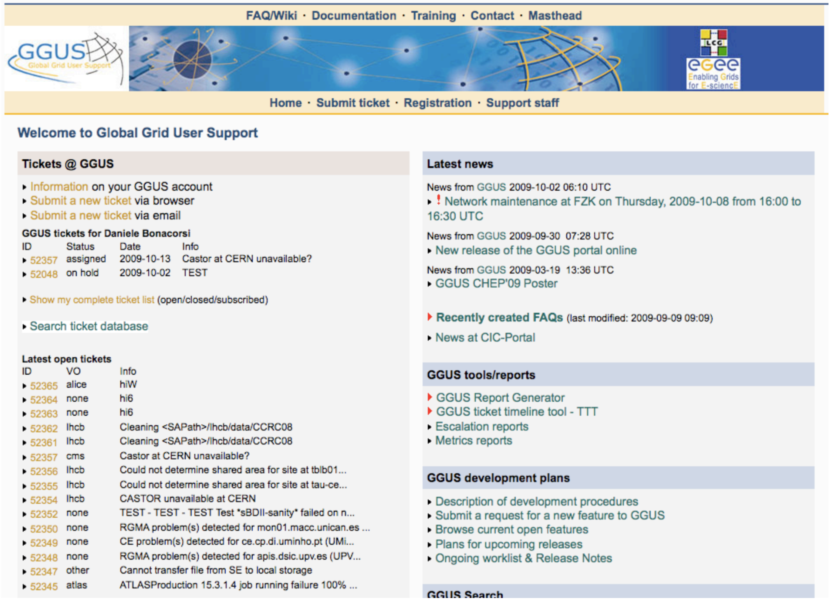
FAQ/Wiki · Documentation · Training · Contact · Masthead
GGUS Global Grid User Support
LCG
egee Enabling Grids for E-sciencE
Home · Submit ticket · Registration · Support staff

Welcome to Global Grid User Support

Tickets @ GGUS

- Information on your GGUS account
- Submit a new ticket via browser
- Submit a new ticket via email

GGUS tickets for Daniele Bonacorsi

ID	Status	Date	Info
52357	assigned	2009-10-13	Castor at CERN unavailable?
52048	on hold	2009-10-02	TEST

- Show my complete ticket list (open/closed/subscribed)
- Search ticket database

Latest open tickets

ID	VO	Info
52365	alice	hiW
52364	none	hi6
52363	none	hi6
52362	lhcb	Cleaning <SAPath>/lhcb/data/CCRC08
52361	lhcb	Cleaning <SAPath>/lhcb/data/CCRC08
52357	cms	Castor at CERN unavailable?
52356	lhcb	Could not determine shared area for site at tblb01...
52355	lhcb	Could not determine shared area for site at tau-ce...
52354	lhcb	CASTOR unavailable at CERN
52352	none	TEST - TEST - TEST Test *sBDII-sanity* failed on n...
52350	none	RGMA problem(s) detected for mon01.macc.unican.es ...
52349	none	CE problem(s) detected for ce.cp.di.uminho.pt (UMi...
52348	none	RGMA problem(s) detected for apis.dsic.upv.es (UPV...
52347	other	Cannot transfer file from SE to local storage
52345	atlas	ATLASProduction 15.3.1.4 job running failure 100% ...

Latest news

News from GGUS 2009-10-02 06:10 UTC
- ! Network maintenance at FZK on Thursday, 2009-10-08 from 16:00 to 16:30 UTC

News from GGUS 2009-09-30 07:28 UTC
- New release of the GGUS portal online

News from GGUS 2009-03-19 13:36 UTC
- GGUS CHEP'09 Poster

- **Recently created FAQs** (last modified: 2009-09-09 09:09)
- News at CIC-Portal

GGUS tools/reports

- GGUS Report Generator
- GGUS ticket timeline tool - TTT
- Escalation reports
- Metrics reports

GGUS development plans

- Description of development procedures
- Submit a request for a new feature to GGUS
- Browse current open features
- Plans for upcoming releases
- Ongoing worklist & Release Notes

GGUS Search

Savannah

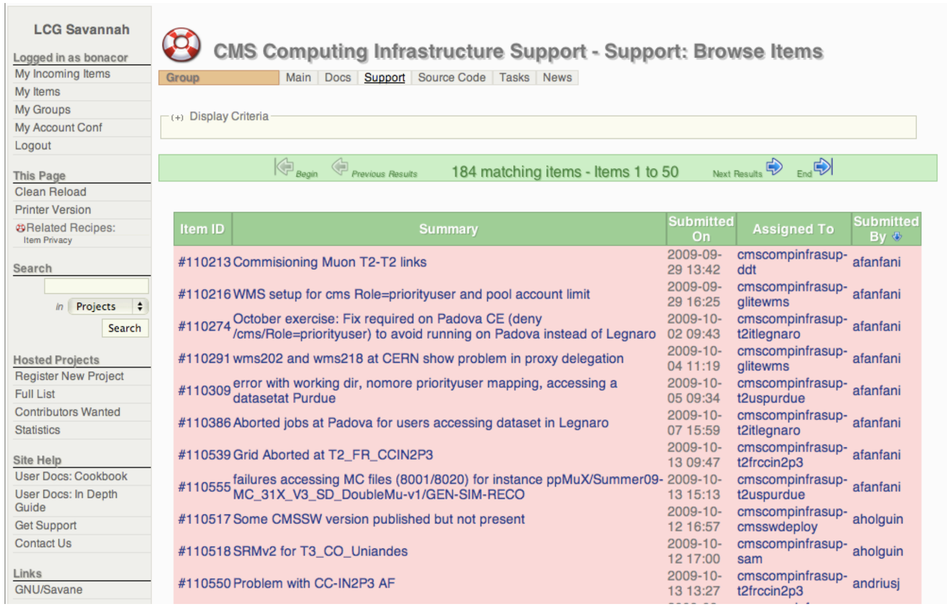

CMS Computing Infrastructure Support - Support: Browse Items

184 matching items - Items 1 to 50

Item ID	Summary	Submitted On	Assigned To	Submitted By
#110213	Commisioning Muon T2-T2 links	2009-09-29 13:42	cmscompinfrasup-ddt	afanfani
#110216	WMS setup for cms Role=priorityuser and pool account limit	2009-09-29 16:25	cmscompinfrasup-glitewms	afanfani
#110274	October exercise: Fix required on Padova CE (deny /cms/Role=priorityuser) to avoid running on Padova instead of Legnaro	2009-10-02 09:43	cmscompinfrasup-t2itlegnaro	afanfani
#110291	wms202 and wms218 at CERN show problem in proxy delegation	2009-10-04 11:19	cmscompinfrasup-glitewms	afanfani
#110309	error with working dir, nomore priorityuser mapping, accessing a datasetat Purdue	2009-10-05 09:34	cmscompinfrasup-t2uspurdue	afanfani
#110386	Aborted jobs at Padova for users accessing dataset in Legnaro	2009-10-07 15:59	cmscompinfrasup-t2itlegnaro	afanfani
#110539	Grid Aborted at T2_FR_CCIN2P3	2009-10-13 09:47	cmscompinfrasup-t2frccin2p3	afanfani
#110555	failures accessing MC files (8001/8020) for instance ppMuX/Summer09-MC_31X_V3_SD_DoubleMu-v1/GEN-SIM-RECO	2009-10-13 15:13	cmscompinfrasup-t2uspurdue	afanfani
#110517	Some CMSSW version published but not present	2009-10-12 16:57	cmscompinfrasup-cmsswdeploy	aholguin
#110518	SRMv2 for T3_CO_Uniandes	2009-10-12 17:00	cmscompinfrasup-sam	aholguin
#110550	Problem with CC-IN2P3 AF	2009-10-13 13:27	cmscompinfrasup-t2frccin2p3	andriusj

GGUS

- ✦ Long tradition of the standard Global Grid User Support system
  - Reaches the WLCG site-admins and the fabric-level experts

Savannah

- ✦ Problem tracking, troubleshooting reference, statistics, …
  - Reaches ‘squads’ easy to define: CMS contacts at Tiers, tools/services experts, …
  - More: baseline tool for Offline Computing shifts, integrated with other CMS projects, ...

Ticketing systems [2/2]

Wouldn't a single ticketing system be preferable?

- Of course. BUT: is there one with all the features CMS uses for Ops?

CMS requested a Savannah-to-GGUS bridging

- Work finalized. Now ready to be used. Start soon to gain experience in Ops
  - Thanks to Guenter Grein (GGUS), Yves Perrin (LCG/SPI) and Simon Metson (CMS) for their great efforts in the technical implementation and testing

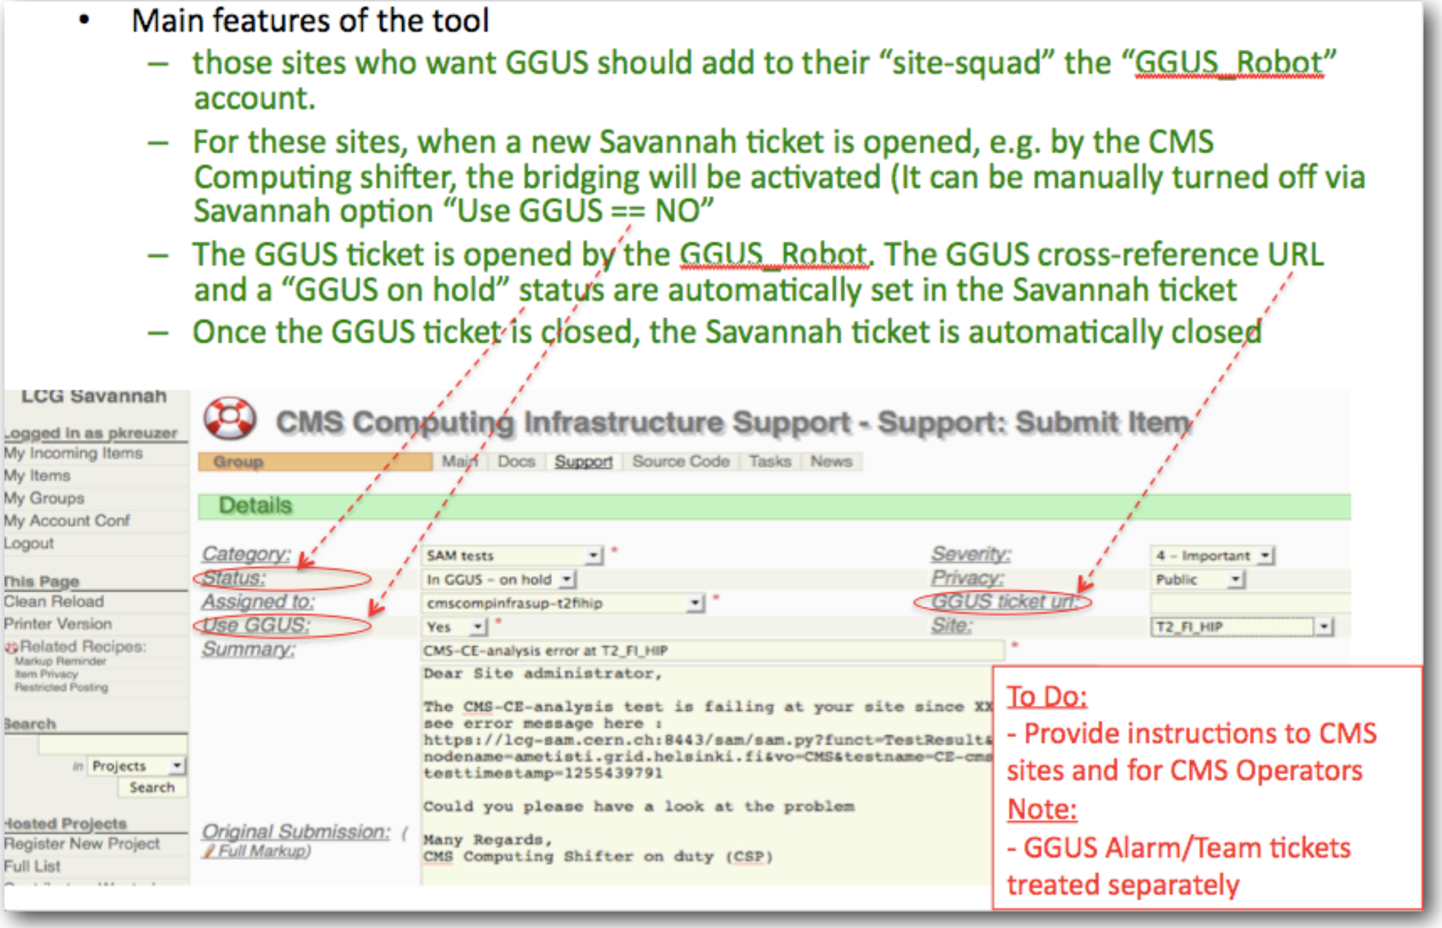

- Main features of the tool
  - those sites who want GGUS should add to their "site-squad" the "GGUS_Robot" account.
  - For these sites, when a new Savannah ticket is opened, e.g. by the CMS Computing shifter, the bridging will be activated (It can be manually turned off via Savannah option "Use GGUS == NO"
  - The GGUS ticket is opened by the GGUS_Robot. The GGUS cross-reference URL and a "GGUS on hold" status are automatically set in the Savannah ticket
  - Once the GGUS ticket is closed, the Savannah ticket is automatically closed

CMS Computing Infrastructure Support - Support: Submit Item

Details

Category: SAM tests *
Status: In GGUS - on hold
Assigned to: cmscompinfrasup-t2fihip *
Use GGUS: Yes *
Summary: CMS-CE-analysis error at T2_FI_HIP *
Severity: 4 - Important
Privacy: Public
GGUS ticket url:
Site: T2_FI_HIP

Original Submission: (Full Markup)

```
Dear Site administrator,

The CMS-CE-analysis test is failing at your site since XX
see error message here :
https://lcg-sam.cern.ch:8443/sam/sam.py?funct=TestResult&
nodename=ametisti.grid.helsinki.fi&vo=CMS&testname=CE-cms
testtimestamp=1255439791

Could you please have a look at the problem

Many Regards,
CMS Computing Shifter on duty (CSP)
```

To Do:
- Provide instructions to CMS sites and for CMS Operators

Note:
- GGUS Alarm/Team tickets treated separately

STEP’09 :: IN2P3

Pre-staging started on June 8-12th due to scheduled HPSS upgrade

- ✦ Site-operated pre-staging approach was chosen *(1)*
- ✦ HPSS v.6.2 interfaced to TReqs interface was used
  - files sorting based on the file position on tape

Sizable multi-VO activity throughout STEP’09

High loads observed on HPSS (June 8-13th) *(2)*:

- ✦ Due to all CMS activities simultaneously, in particular CMS analysis at the T2, and also other VOs activities
  - Decided to suspend T2 analysis activity during STEP’09

Reprocessing

- ✦ High reprocessing load by CMS and other VOs *(4)*
  - Failures mainly due to stage-out
- ✦ File distribution per tape on a typical day averages at ~10 *(3)*

Transfer

- ✦ Relatively smooth
  - some structure (in quality) to be cured, mainly in T1-T1 *(5)*

STEP'09 :: IN2P3 in plots

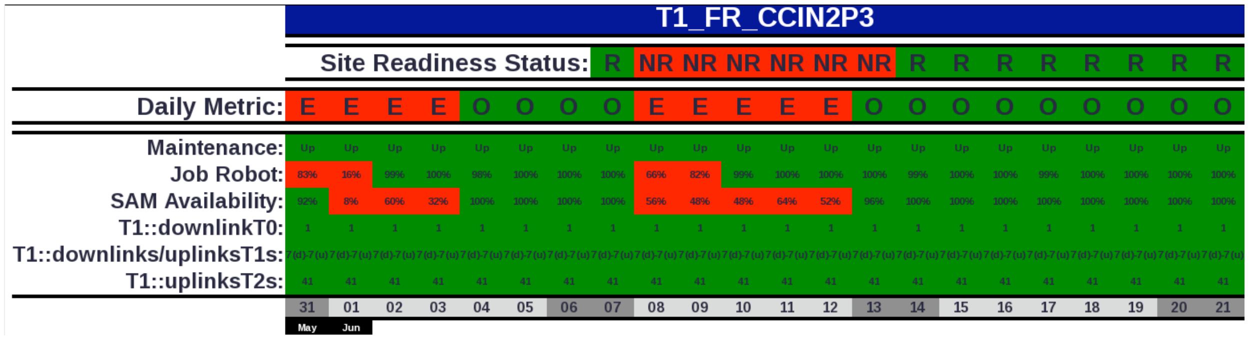

T1_FR_CCIN2P3																						
Site Readiness Status:								R	NR	NR	NR	NR	NR	NR	R	R	R	R	R	R	R	R
Daily Metric:	E	E	E	E	O	O	O	O	E	E	E	E	E	O	O	O	O	O	O	O	O	O
Maintenance:	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up
Job Robot:	83%	16%	99%	100%	98%	100%	100%	100%	66%	82%	99%	100%	100%	100%	99%	100%	100%	99%	100%	100%	100%	100%
SAM Availability:	92%	8%	60%	32%	100%	100%	100%	100%	56%	48%	48%	64%	52%	96%	100%	100%	100%	100%	100%	100%	100%	100%
T1::downlinkT0:	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
T1::downlinks/uplinksT1s:	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)	7(d)-7(u)
T1::uplinksT2s:	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41
	31	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21
	May	Jun																				

(3)

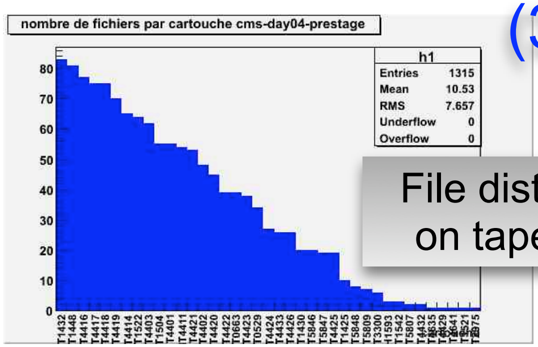

File distrib on tapes

(1)

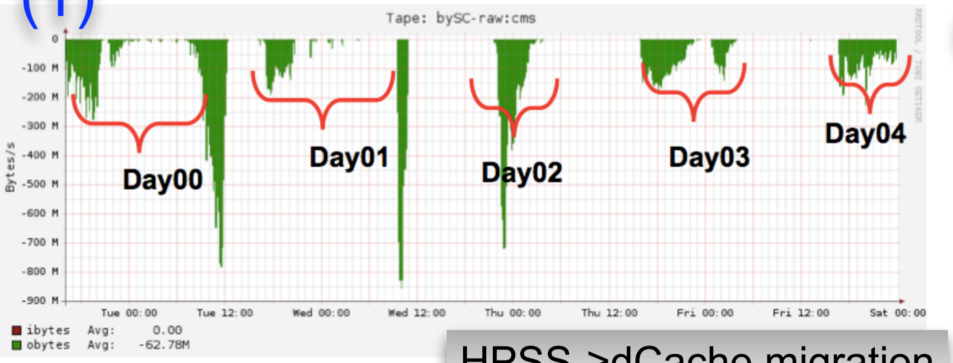

HPSS->dCache migration

(4)

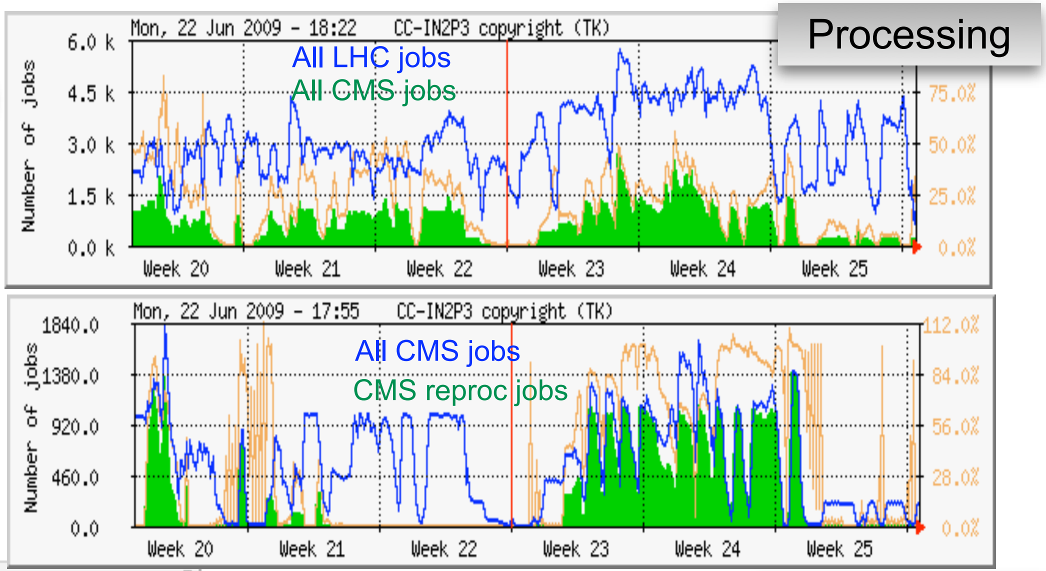

All LHC jobs

All CMS jobs

All CMS jobs

CMS reproc jobs

Processing

(2)

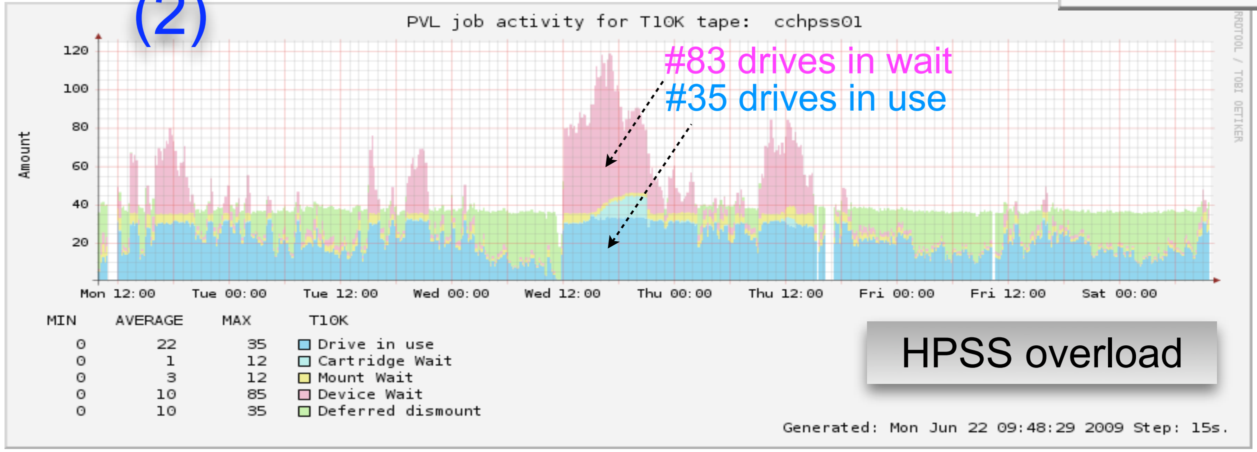

#83 drives in wait

#35 drives in use

MIN	AVERAGE	MAX	T10K
0	22	35	□ Drive in use
0	1	12	□ Cartridge Wait
0	3	12	□ Mount Wait
0	10	85	□ Device Wait
0	10	35	□ Deferred dismount

Generated: Mon Jun 22 09:48:29 2009 Step: 15s.

HPSS overload

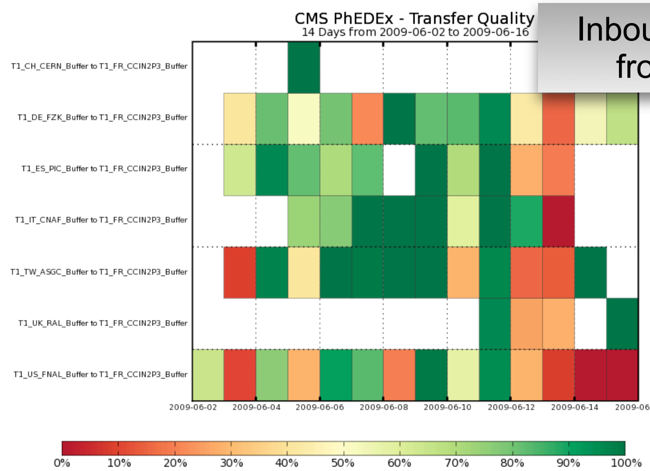

Inbound traffic from T1s

(5)

STEP'09 :: IN2P3 *[pre-staging]*

	Site-operated	Central SRM script	PhEDEX agent
ASGC			X
CNAF		X	*Castor + StoRM: need additional work in PhEDEx*
FNAL	X		
FZK	X	*Tape issues: preferred manual*	*Tape issues: preferred manual*
IN2P3	X	*HPSS downtime on week-1: preferred manual*	*HPSS downtime on week-1: preferred manual*
PIC			X
RAL			X

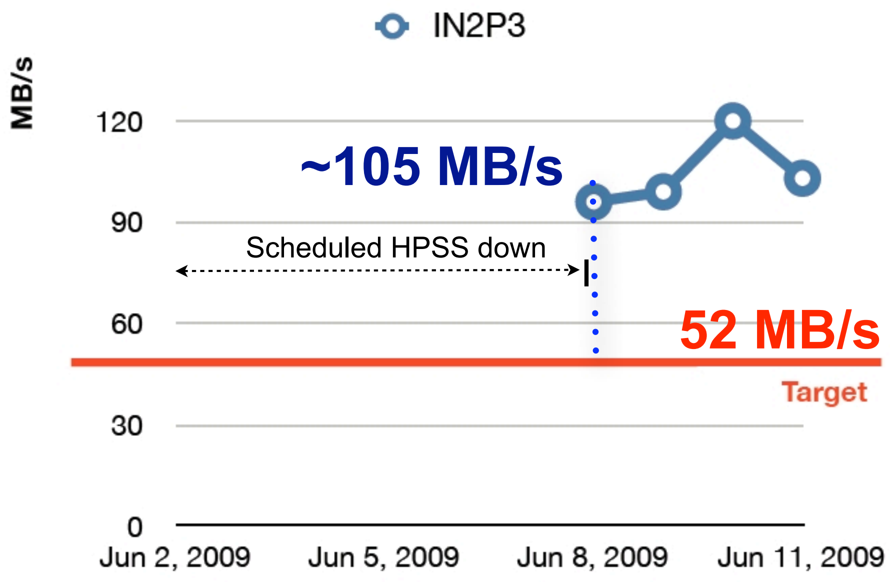

	Target [MB/s]	2-Jun	3-Jun	4-Jun	5-Jun	6-Jun	7-Jun	8-Jun	9-Jun	10-Jun	11-Jun
ASGC	73	*Digesting migration*	140	170	190	160	145	150	140	150	220
CNAF	56	380	300	160	240	240	270	105	80	125	240
FNAL	242	280	200	200	120	*Still staging previous day*	*Recovering from backlog*		379	380	400
FZK	85	*Tape system not available [unscheduled downtime]*							*Participated in pre-staging but performance not clear*		
IN2P3	52	*Tape system not available [scheduled downtime]*						96	99	120	103
PIC	50	60	61	106	83	*Samples not purged*	*Samples partially on*	99	142	123	142
RAL	40	250	230	160	140	135	190	170	100	220	180

STEP’09 :: IN2P3 *[CPU efficiency]*

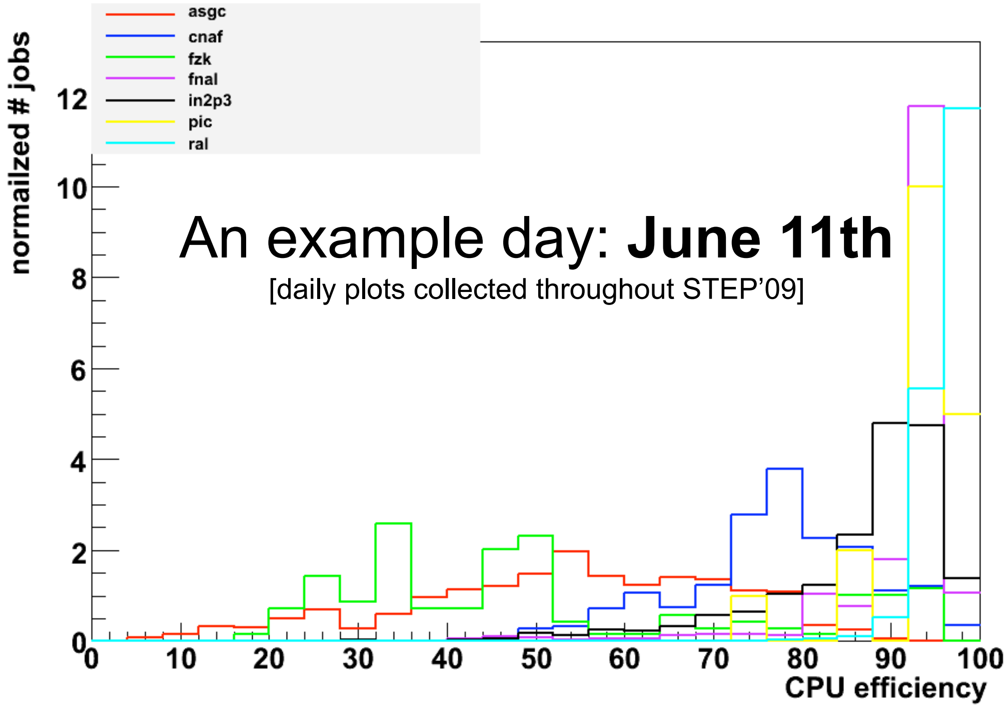

An example day: **June 11th**

[daily plots collected throughout STEP’09]

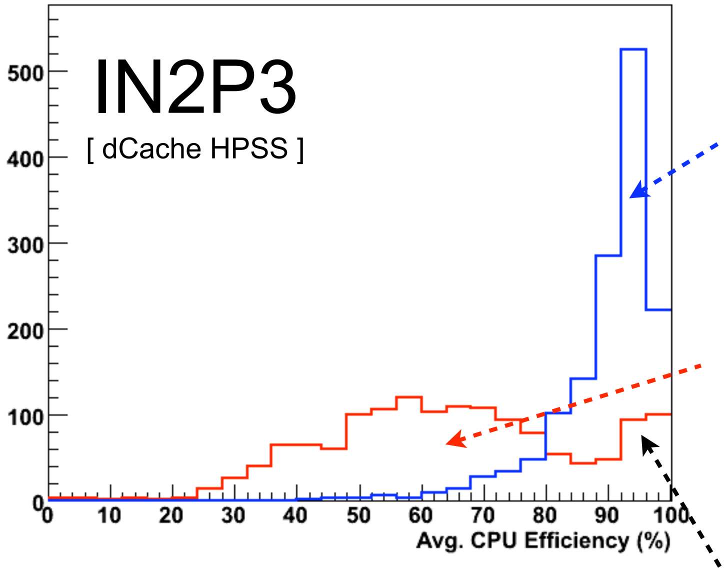

IN2P3

[dCache HPSS]

with pre-staging

without pre-staging

Previously failed jobs might have already triggered the pre-staging

Measured every day, at each T1 site. Mixed results:

- Very good CPU efficiency for **FNAL**, **IN2P3**, (**PIC**), **RAL**
- Not so good CPU efficiency for **ASGC**, **CNAF**
- Test not significant for **FZK**

Current understanding:

- Test demonstrated the significant effect of pre-staged data for processing
- Site specifics to be investigated: **IN2P3 not one of these**