




:: CMS AAA federation, monitoring and operations

Marian Zvada on behalf of CMS AAA

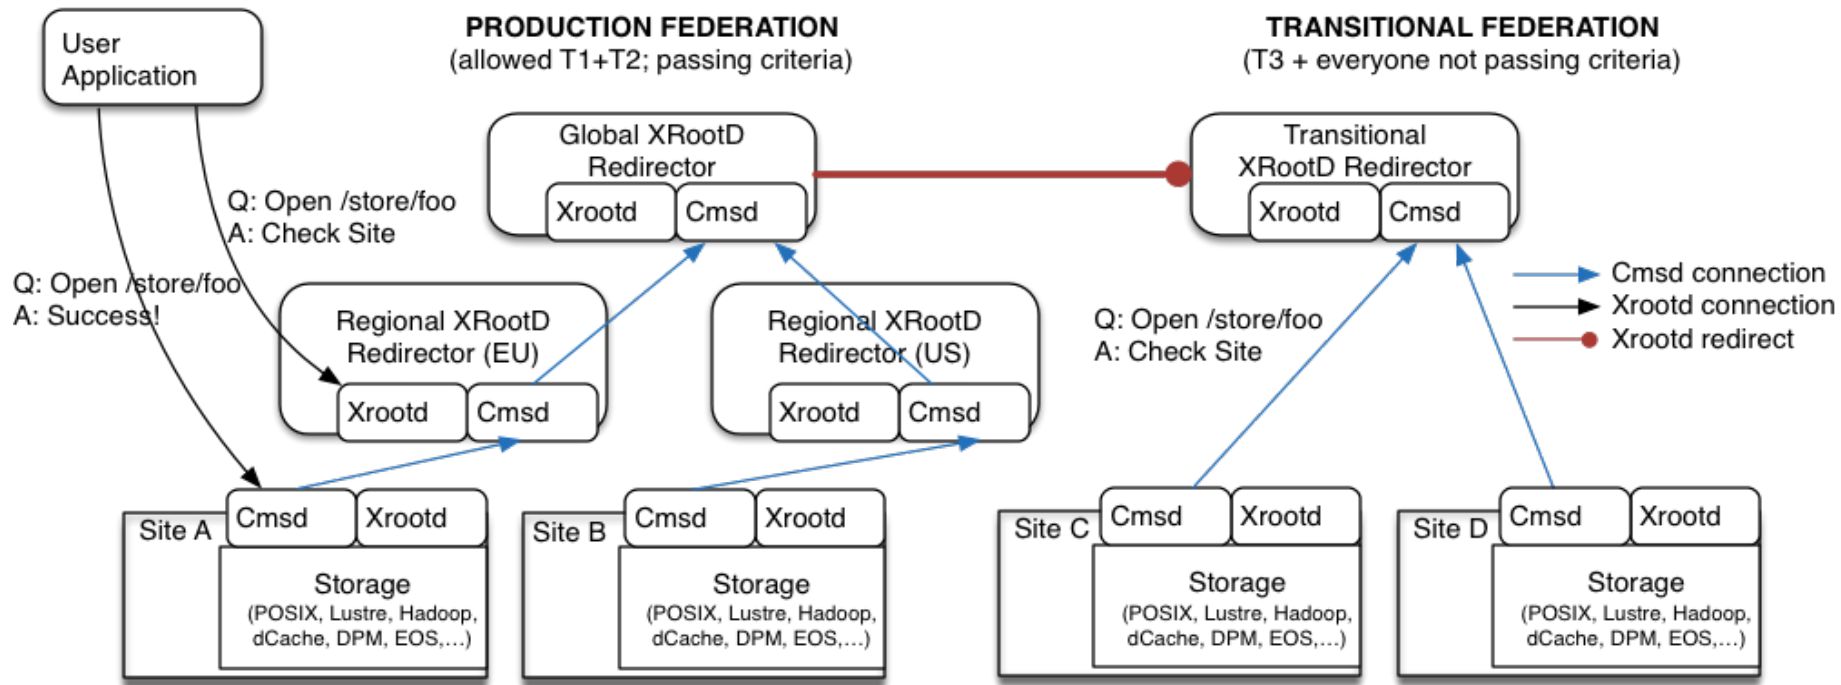
DPM2018 @ Prague
June 1, 2018



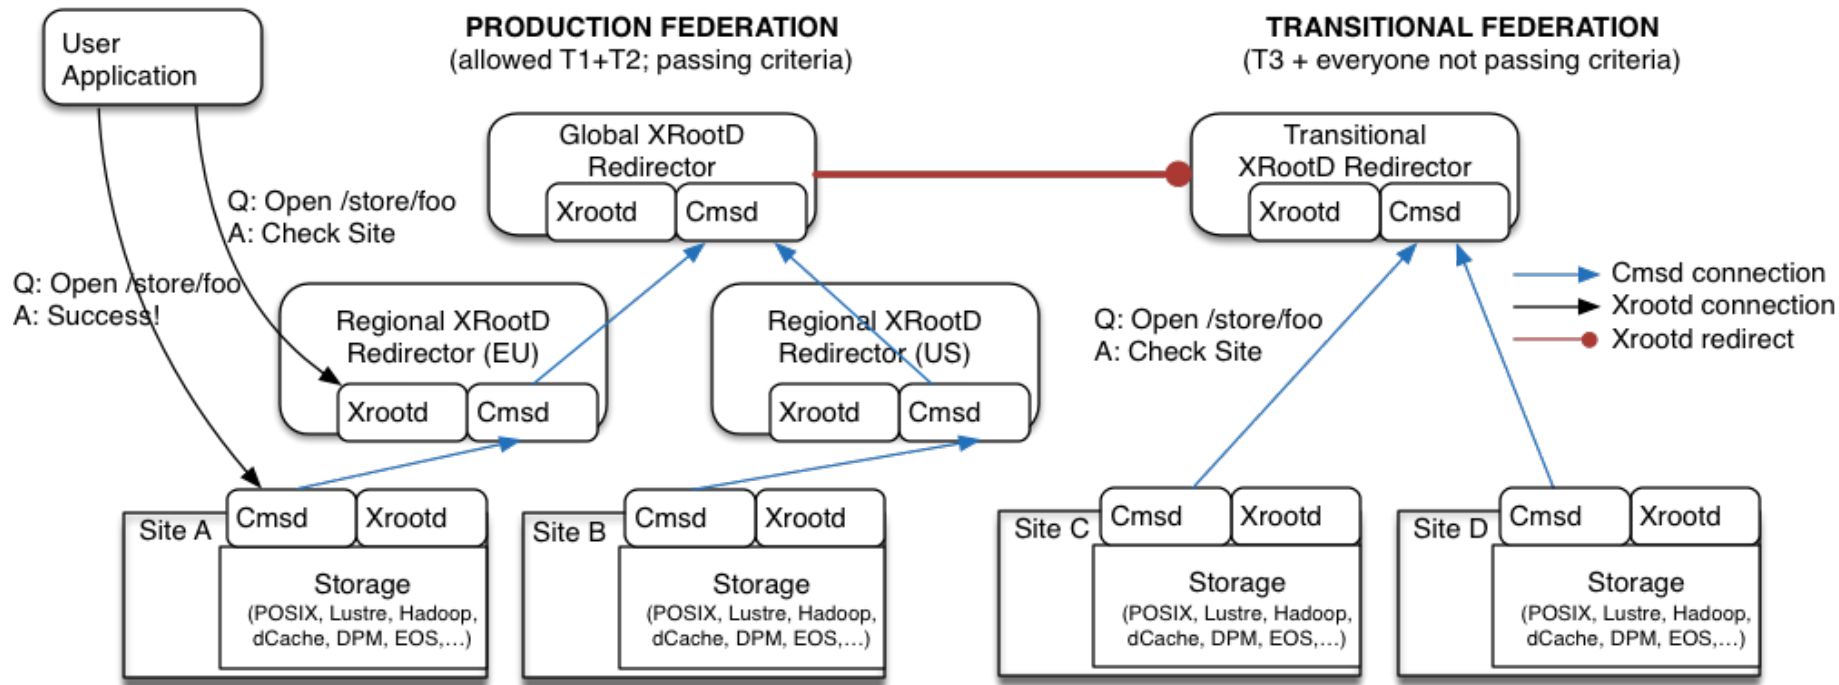
Refresher about AAA

- **AAA = Any data, Anytime, Anywhere**
 - similar concept ATLAS uses for FAX
- An effort to create a storage federation of the CMS sites
- AAA makes CMS data access transparent toward users at any CMS sites
- Sites' data content is federated on the fly using the native clustering of the xrootd framework
- AAA is adopted by CMS as a system to access remote data if not available locally.
 - Its usage continues to grow past the original use case (fallback) from >5 years ago.
- Distinguishing between Production and Transitional Federation

Production and Transitional Federation



Production and Transitional Federation

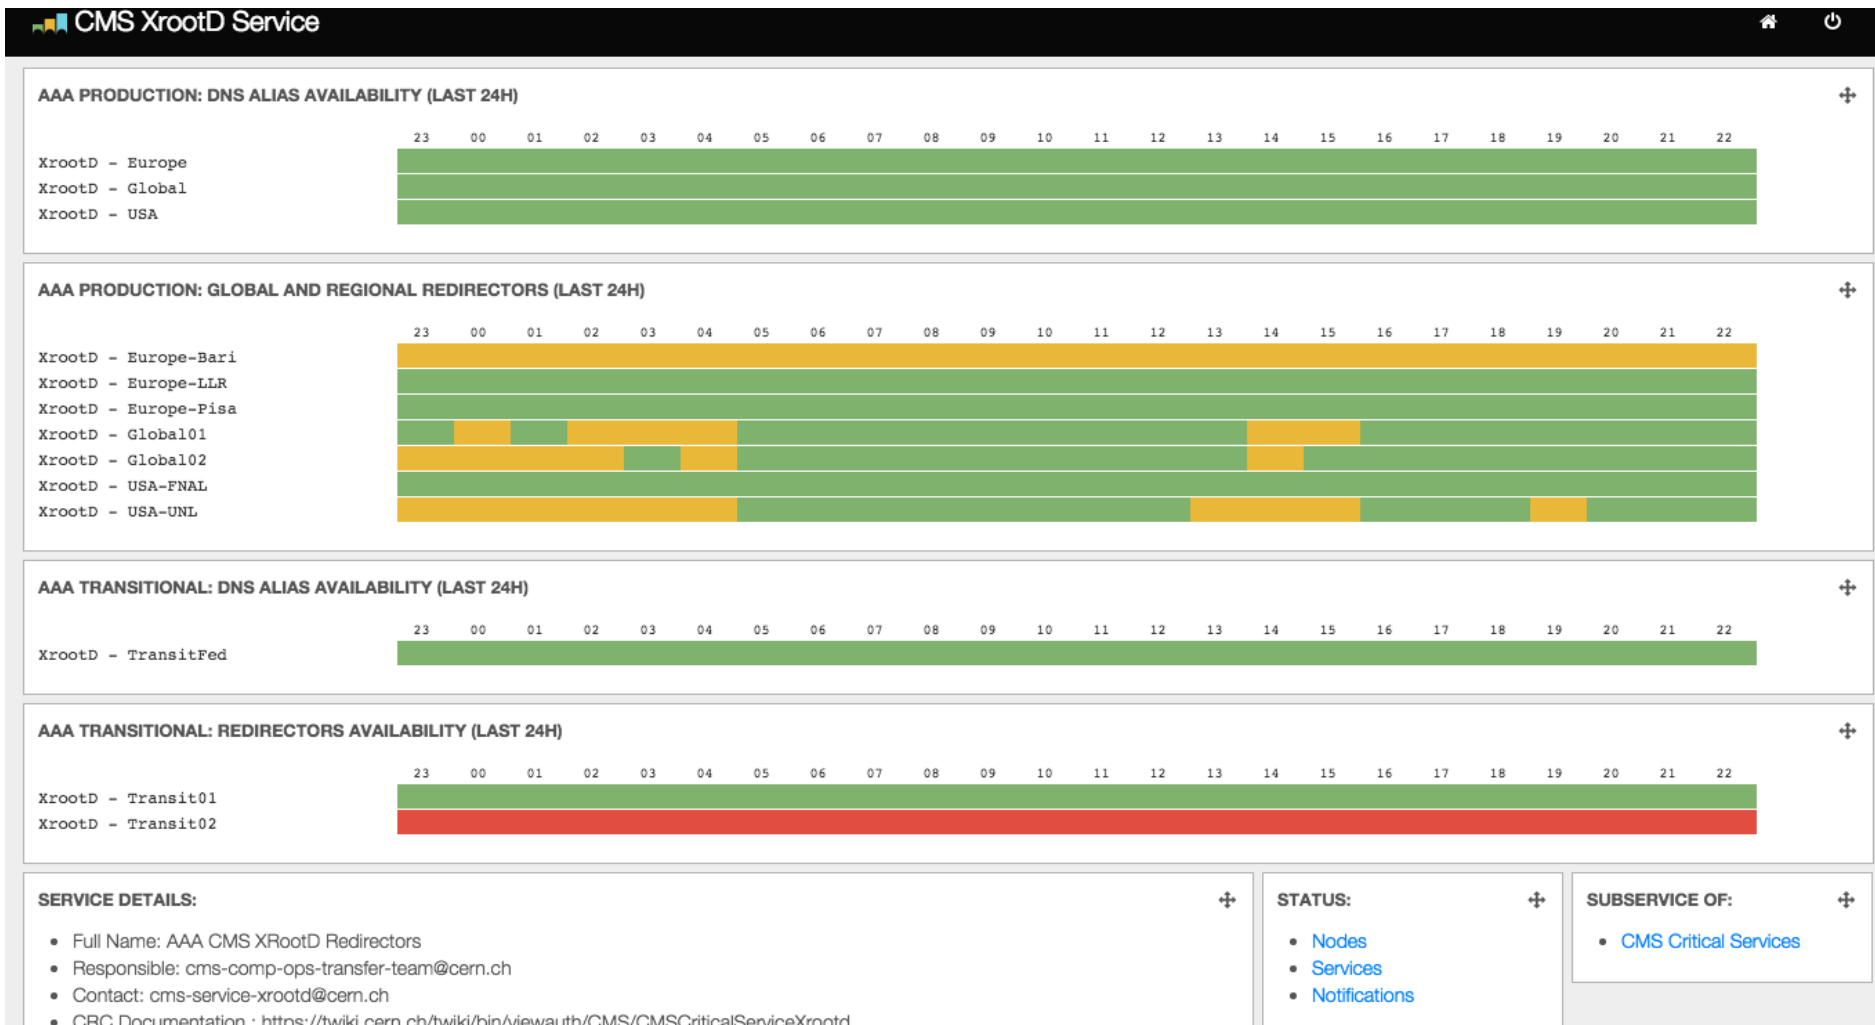


- based on scaling tests and other criteria → isolate sites which might affect overall AAA functionality, e.g. sites with weak availability and reliability
- keep production activities intact:
 - even if site in production federation shifted to transitional federation, data unique to the site is still accessible
- production->transitional fallback transparent to users

AAA in the CMS critical service map



Top level monitoring – regional redirectors





Kibana:

https://meter.cern.ch/public/_plugin/kibana/#/dashboard/temp/CMS::XrootD

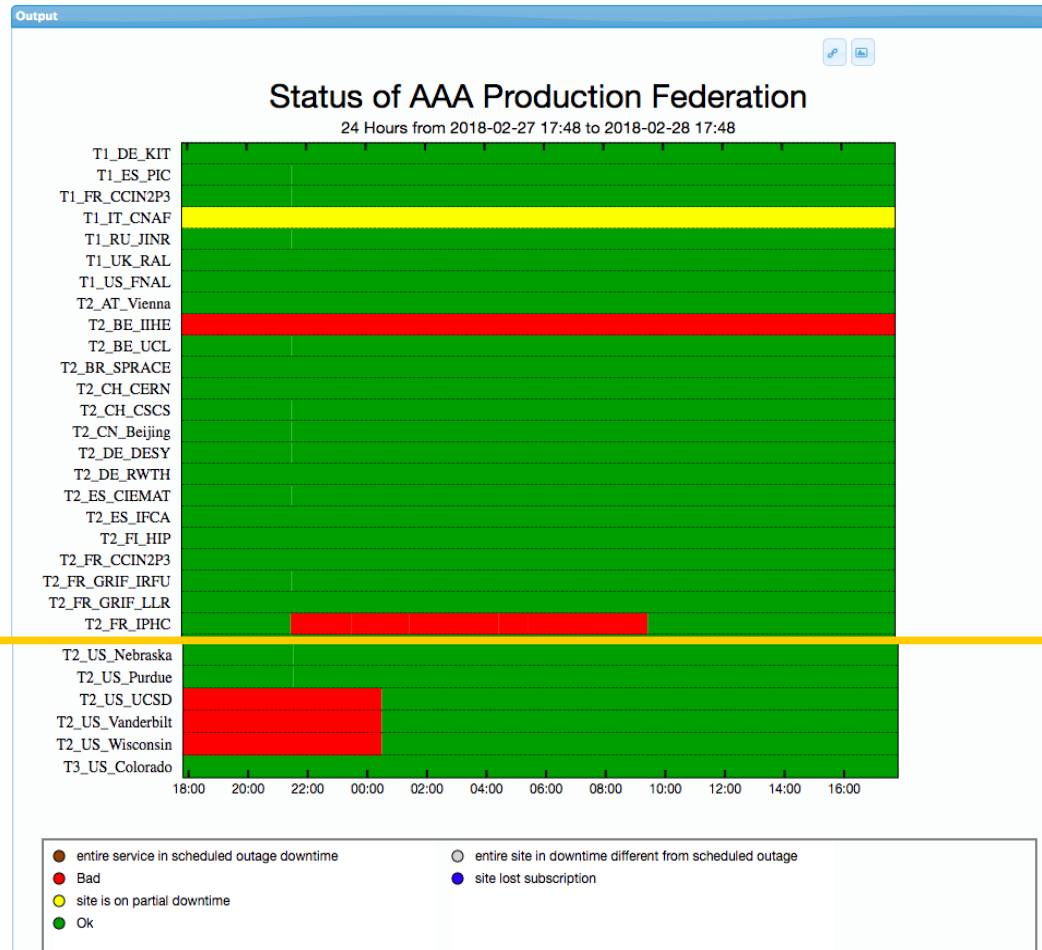
Monitoring access list of federations:

Production and Transitional Federation Sites



Monitoring – detailed view prod federation

<https://dashb-ssb.cern.ch/dashboard/request.py/siteviewhistory?columnid=224>



Criteria taken into account:

- AAA-related ticket in GGUS open for longer than two weeks.
- SAM xrootd access test < 50% for two weeks.
- Hammer Cloud (HC) test success rate < 80% for two weeks

Monitoring – detailed view per site

→ click on particular site (status bar) from previous slide see details

AAA Report, updated on 2018-02-28 at 18:29:06




Condition

```
IF (HC::xrootd rate > 70% OR site is tier 3) AND SAM::xrootd-access > 50 AND site has no AAA related ticket :
    site is good/green
ELSE:
    site is bad/red
```

Explanation

- SAM (xrootd-access) and HC (xrootd) rows represent **average success rate of last 2 weeks** starting from the given date
 - Values are scaled between red(low) and green(high) color
- GGUS row shows number of AAA related tickets. If there is no ticket, it is represented as green for the given date
- Time window of the table is **2 weeks** (please go over one of the rows and see the date and value)

Result

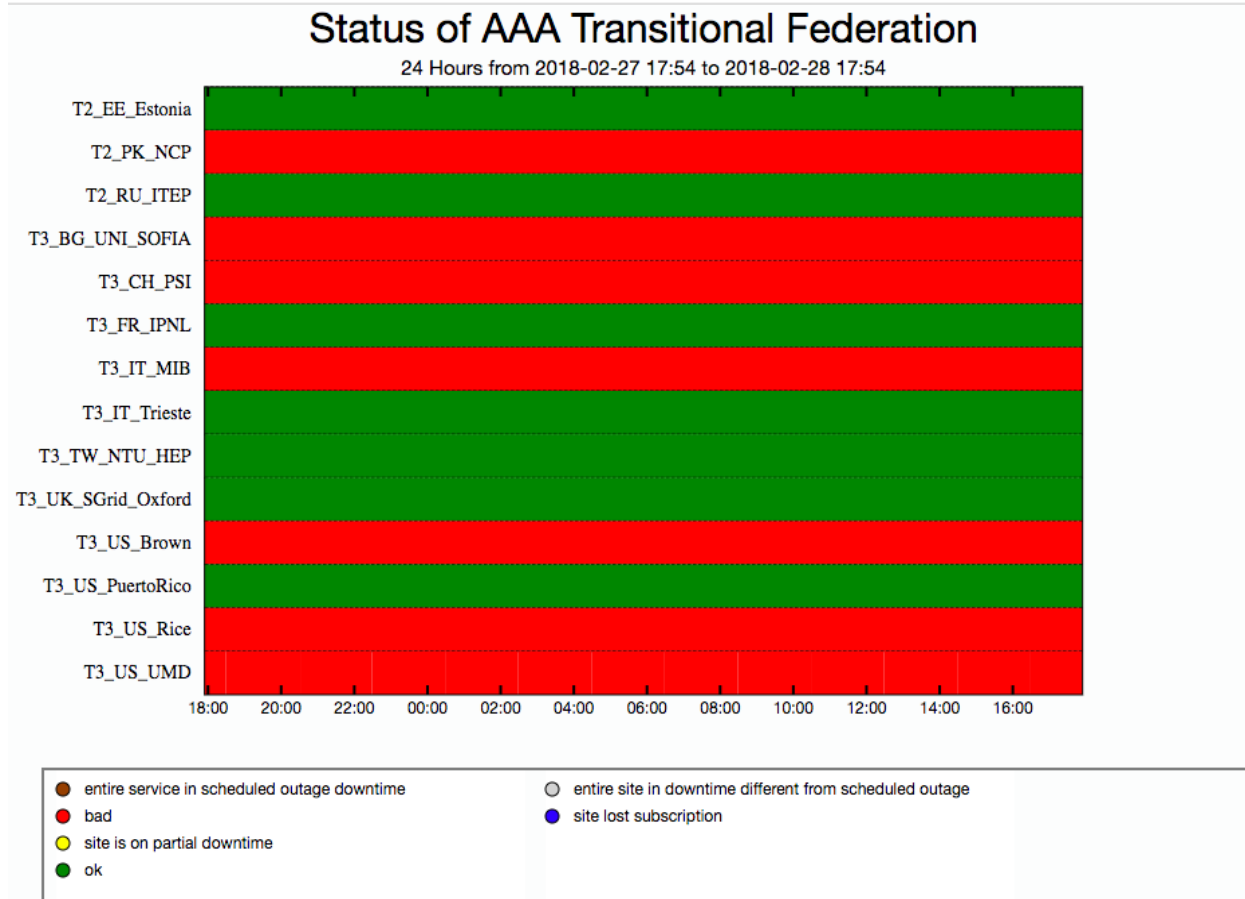
T2_US_Wisconsin	
SAM	
HC	
GGUS	

Criteria taken into account:

- AAA-related ticket in GGUS open for longer than two weeks.
- SAM xrootd access test < 50% for two weeks.
- Hammer Cloud (HC) test success rate < 80% for two weeks

Monitoring – detailed view transit federation

<https://dashb-ssb.cern.ch/dashboard/request.py/siteviewhistory?columnid=219>



Criteria taken into account:

- AAA-related ticket in GGUS open for longer than two weeks.
- SAM xrootd access test < 50% for two weeks.
- Hammer Cloud (HC) test success rate < 80% for two weeks

Monitoring – detailed view per site

→ click on particular site (status bar) from previous slide see details

Secure | https://cmsst.web.cern.ch/cmsst/aaa/T3_US_Brown_report.html ☆

AAA Report, updated on 2018-02-28 at 18:29:06

Condition

IF (HC::xrootd rate > 70% OR site is tier 3) AND SAM::xrootd-access > 50 AND site has no AAA related ticket:
 site is good/green
 ELSE:
 site is bad/red

Explanation

- SAM (xrootd-access) and HC (xrootd) rows represent **average success rate of last 2 weeks** starting from the given date
 - Values are scaled between red(low) and green(high) color
- GGUS row shows number of AAA related tickets. If there is no ticket, it is represented as green for the given date
- Time window of the table is **2 weeks** (please go over one of the rows and see the date and value)

Result

T3_US_Brown	
SAM	
HC	
GGUS	

Criteria taken into account:

- AAA-related ticket in GGUS open for longer than two weeks.
- SAM xrootd access test < 50% for two weeks.
- NO (HC) test for T3s

SAM tests for AAA

- Site Availability Monitoring (SAM) tests include two tests supporting AAA
 - XrootD **access**
 - XrootD **fallback**
- **Access** tests checks whether site files accessible through AAA
 - Checks VOMS proxy, CMSSW, TFC configuration
- **Fallback** tests checks whether site can access remote files via AAA
- AAA transfer team uses SAM access test to assess CMS T1 & T2 sites
- SAM **access** test provides one criterion for remaining in production or transitional federation
- **Both tests** are not only used for the AAA federation testing but are part of the site evaluation (i.e. site readiness)

Performance assessment

- To evaluate the potential of data federation, CMS needs to understand the current performance of each site
 - how are the sites performing? Is their performance and quality of service sufficient?
- The “**File opening and reading scale tests**” measures ability of CMS sites to handle predicted peak load for AAA
- Tests allow to discover unoptimized sites and suggest improvements

Performance assessment

- To evaluate the potential of data federation, CMS needs to understand the current performance of each site
 - how are the sites performing? Is their performance and quality of service sufficient?
- The “**File opening and reading scale tests**” measures ability of CMS sites to handle predicted peak load for AAA
- Tests allow to discover unoptimized sites and suggest improvements

File opening and reading scale tests

- Tests emulate CMS jobs running at CMS sites choosing the site through regional redirectors
- two measurement are performed:
 - Rate to open files at a site via regional redirector
 - Rate to reading data from files at a site opened via regional redirector

File opening and reading scale tests

- CMS target for tests are:
 - **File-opening test:** access total rate of 100 Hz at a site
 - tests run up to 100 jobs simultaneously, that open files at rate of 2Hz each
 - **File-reading test:** 600 jobs reading average rate of 2.5MB every 10s at a site → reading total rate of 150MB/s
 - Test runs up to 800 jobs simultaneously, each one reading data blocks of 2.5MB from a file
- Target numbers comes from internal CMS analysis based on historical figures

File opening and reading scale tests

- CMS target for tests are:
 - **File-opening test:** access total rate of 100 Hz at a site
 - tests run up to 100 jobs simultaneously, that open files at rate of 2Hz each
 - **File-reading test:** 600 jobs reading average rate of 2.5MB every 10s at a site → reading total rate of 150MB/s
 - Test runs up to 800 jobs simultaneously, each one reading data blocks of 2.5MB from a file
- Target numbers comes from internal CMS analysis based on historical figures

How the tests are run

- Sites provide a special path to allow redirector to match only the site to test – TFC trick as plugin for the xrootd at a site
- the list of input files is obtained via PhEDEx
- tests run from a “controlled” condor pool in Wisconsin (no grid jobs); also having resources at CERN owned by CMS

Status of the tests

- Two regions (sub-federations) to test: EU and US
- Tests usually run on a subset of CMS sites
 - first phase of work is the check of site setup
 - Correctness of TFC and special path
 - 34 non-US sites + 9 US sites are tested once a week via EU DNS alias redirector (Bari, Pisa and Paris)

Status of the tests

- Two regions (sub-federations) to test: EU and US
- Tests usually run on a subset of CMS sites
 - first phase of work is the check of site setup
 - Correctness of TFC and special path
 - 34 non-US sites + 9 US sites are tested once a week via EU DNS alias redirector (Bari, Pisa and Paris)

Various storage backend of tested sites

- xrootd protocol is the common access interface

	non-US sites	US sites
→	17 dCache (4 tier1)	1 dCache (tier 1)
	2 Hadoop	6 Hadoop/BeStMan
	16 DPM (~6 in 2016!)	1 Lustre/BeStMan
	6 StoRM (1 tier 1)	1 LStore/BeStMan
	1 Castor	

Summary of the tests - [results](#)

Summary of AAA opening and reading tests

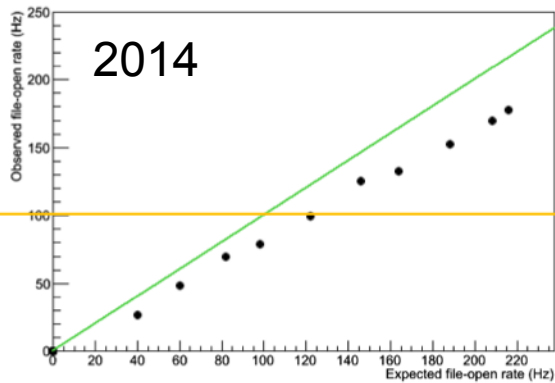
SITES	castor	dcache	dpm	storm	hadoop/BeStMan	LStore/BeStMan	Lustre/BeStMan
TEST STATUS	OPENING TEST						READING TEST
FAILED	completely failed test, no plots produced						completely failed test, no plots produced
PROBLEM	the opening rate is lower than 10 Hz						the reading rate is lower than 150 MB/s and the number of simultaneous clients is lower than 600
WARNING	the number of simultaneous clients is lower than 90						the reading rate is lower than 150 MB/s even if the number of simultaneous clients reaches 600
OK	the opening rate reaches 10 Hz and the number of simultaneous clients is bigger than 90						the reading rate reaches 150 MB/s

T1_DE_KIT-xrootd-cms.infn.it_26_05_18	OPENING	READING
T1_ES_PIC-xrootd-cms.infn.it_25_05_18	OPENING	READING
T1_FR_CCIN2P3-xrootd-cms.infn.it_26_05_18	OPENING	READING
T1_IT_CNAF-xrootd-cms.infn.it_26_05_18	OPENING	READING
T1_RU_JINR-xrootd-cms.infn.it_26_05_18	OPENING	READING
T1_UK_RAL-xrootd-cms.infn.it_25_05_18	OPENING	READING
T1_US_FNAL_Disk-cmsxrootd.fnal.gov_27_05_18	OPENING	READING
T2_AT_Vienna-xrootd-cms.infn.it_26_05_18	OPENING	READING
T2_BE_IIHE-xrootd-cms.infn.it_28_05_18	OPENING	READING
T2_BE_UCL-xrootd-cms.infn.it_25_05_18	OPENING	READING
T2_BR_SPRACE-xrootd-cms.infn.it_28_05_18	OPENING	READING
T2_CN_Beijing-xrootd-cms.infn.it_28_05_18	OPENING	READING
T2_DE_DESY-xrootd-cms.infn.it_28_05_18	OPENING	READING
T2_DE_RWTH-xrootd-cms.infn.it_28_05_18	OPENING	READING


File-opening tests of DPM sites

- Tests run 100 jobs simultaneously, with opening file rate of 2Hz each → plots shows attempted file open rate vs. observed rate. Ideal is observed=attempted (green line). If the average time to open a file is longer than 0.5s, performance will be below the green line
 - *performance depends on the storage backend, configuration and hardware. If a site is supporting multi-VO, a slow rate can be due to contention with other VO's*


T2_AT_Vienna



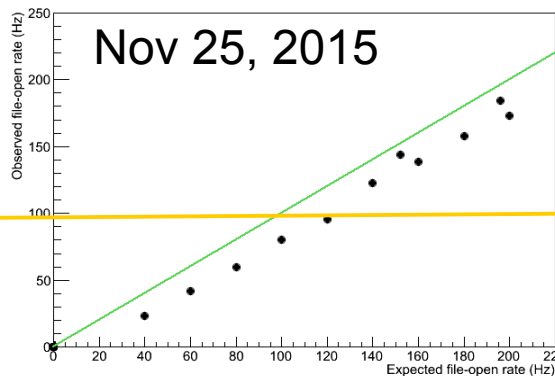
T2_IN_TIFR



T2_UA_KIPT




T2_AT_Vienna

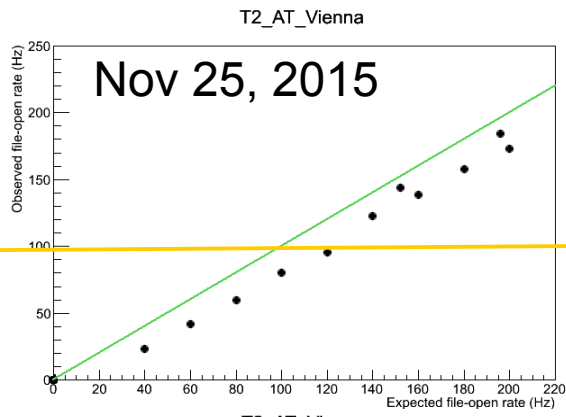


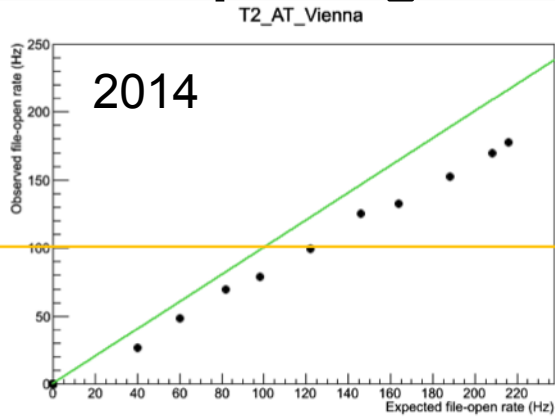


complete failed test for T2_IN_TIFR, thus no plots

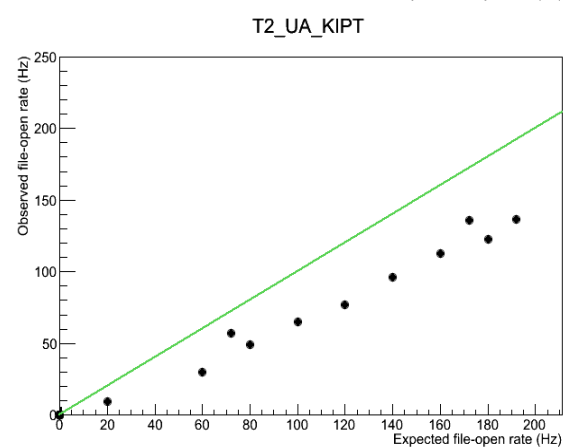
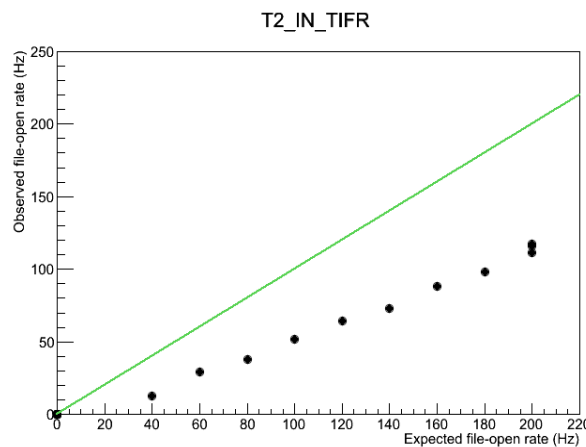
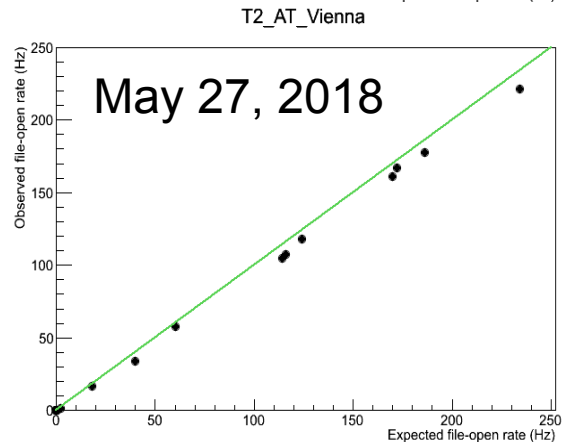

T2_UA_KIPT



File-opening tests of DPM sites





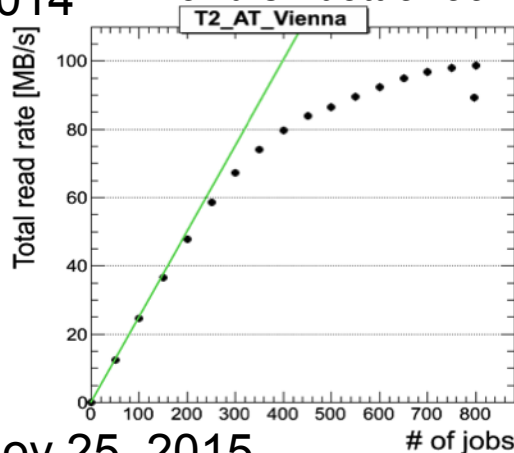
complete failed test
for T2_IN_TIFR,
thus no plots



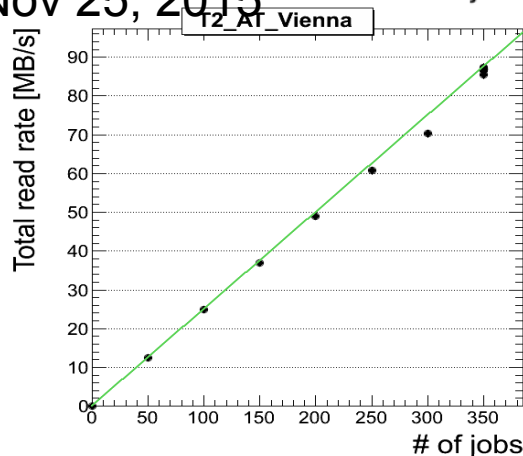
File-reading tests of DPM sites

- Tests run up to 800 simultaneously jobs reading block of 2.5MB every 10s from a file (an input file each job)
 - **plots show the total rate as function of number of jobs, the expected rate should follow the green line**
 - *if a site doesn't reach the target, the reason could be a network or filesystem or disk bottleneck. Only the site monitoring of nodes' load can say more.*


2014




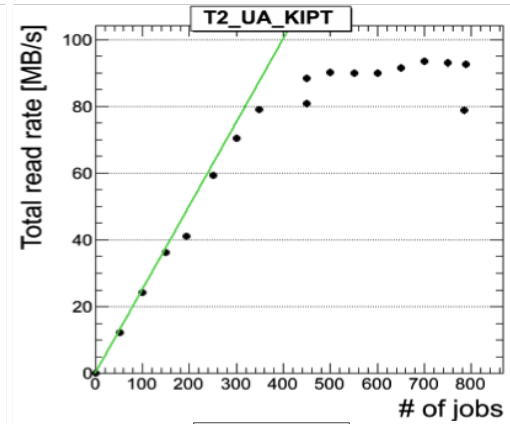
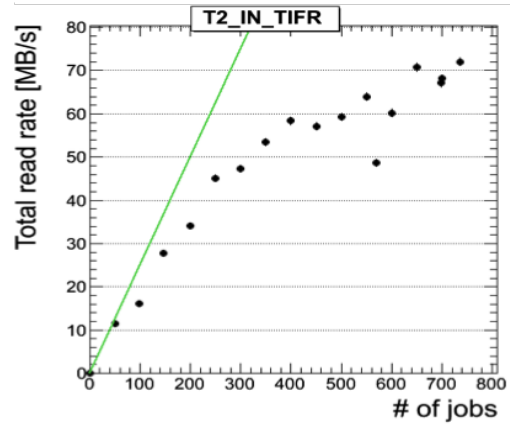

Nov 25, 2015



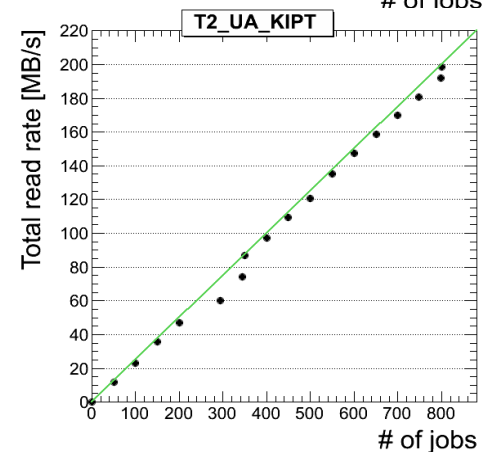
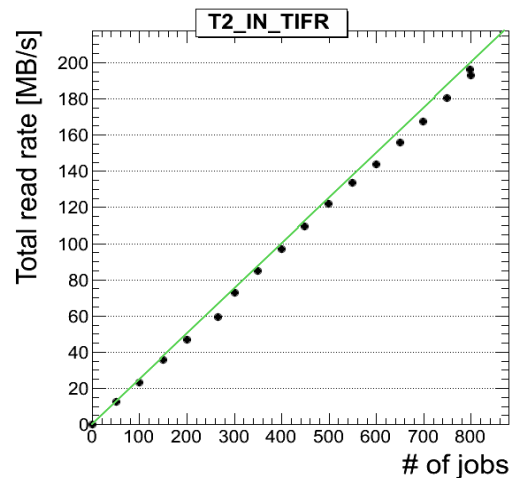


complete failed test for T2_IN_TIFR, thus no plots




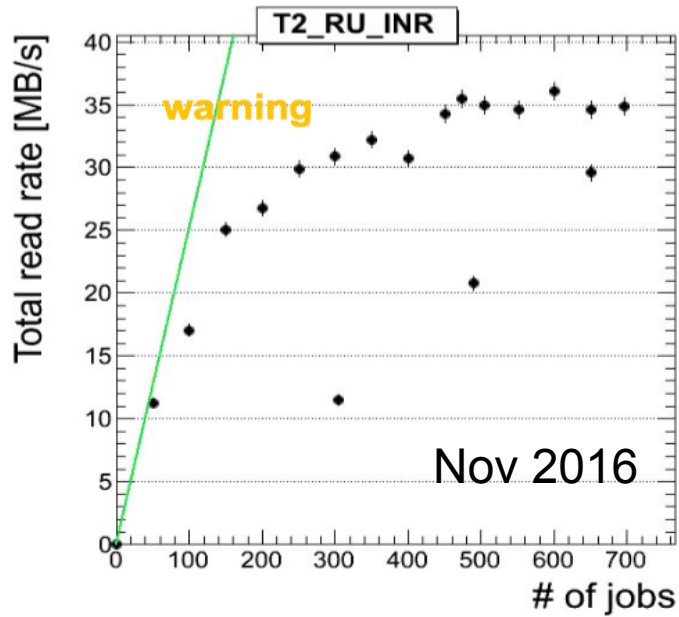
File-reading tests of DPM sites



complete failed test
for T2_IN_TIFR,
thus no plots



File-reading tests of DPM sites



Operations and Monitoring – SSB view



dashb-ssb.cern.ch/dashboard/request.py/siteview?view=Site%20Readiness#currentView=test&highlight=true

Help Metric Login Register for site notifications Site Status for the CMS sites, v0.17.0 r0_rc6.wlcgmon6

dashboard Index Expanded Table

Show 200 entries

Print view: test Search...

Site Name	Space Check	phedex-node in site-local-config	xrootd-redirector	CMSSW Generic File Monitoring	Under investigation	AAA Federations	AAA Transitional Federation	AAA Production Federation	AAA-XrootD Read Tests	AAA-XrootD Open Tests	Phedex version
T2_UA_KIPT-cms.infn.it									OK_18_02_18	OK_18_02_18	
T2_RU_JINR-cms.infn.it									OK_18_02_18	PROBLEM_10_02_18	
T2_IN_TIFR	2017-03-26	yes	yes	yes	mark	prod		on			PHEDEX_4_2_1
T2_US_Caitech-cmsxrootd.fnal.gov									OK_18_02_18	OK_18_02_18	
T2_RU ITEP	2015-09-06	yes	no	yes	mark	trans	on				PHEDEX_4_2_1
T2_AT_Vienna-cms.infn.it									OK_17_02_18	OK_17_02_18	
T2_BE_IHE-cms.infn.it									OK_18_02_18	OK_18_02_18	
T2_ES_CIEMAT-cms.infn.it									OK_18_02_18	PROBLEM_10_02_18	
T2_FR_GRIF_LLR-cms.infn.it									OK_17_02_18	OK_17_02_18	
T2_FR_IPHC-cms.infn.it									OK_17_02_18	OK_17_02_18	
T2_UK_SGrid_Bristol-cms.infn.it									OK_16_02_18	OK_16_02_18	
T2_FL_HIP	2016-10-03	yes	yes	yes	mark	prod		on			PHEDEX_4_2_1
T2_RU_JINR	2018-02-26	yes	yes	yes	mark	prod		on			PHEDEX_4_2_1
T2_UK_London_Brunel	2015-06-15	yes	yes	yes	mark	prod		on			PHEDEX_4_2_1
T2_PT_NCG_Lisbon-cms.infn.it									WARNING_16_02_18	PROBLEM_10_02_18	

Found a bug?

Status of deployed versions to date

Site_Name	SB_Version	Xrootd_Version	AAA Federation
T2_AT_Vienna	1.9.0	v4.7.1	Production Federation
T2_FR_GRIF_IRFU	1.9.0		Production Federation
T2_FR_GRIF_LLQ	1.9.0	v4.8.2	Production Federation
T2_FR_IPHC	1.9.0	v4.6.1	Production Federation
T2_GR_loannina	1.9.0	v4.8.2	Production Federation
T2_HU_Budapest	1.9.0	v4.8.1	Production Federation
T2_IN_TIFR	1.9.0	v4.8.2	Production Federation
T2_PL_Swierk	1.9.0	v4.8.3	Production Federation
T2_PL_Warsaw	1.9.0	v4.7.1	Production Federation
T2_RU_INR	1.9.0	v4.7.1	Production Federation
T2_TR_METU	1.8.11	v4.8.1	Production Federation
T2_UA_KIPT	1.8.11	v4.4.1	Production Federation
T2_UK_London_Brunel	1.9.0	v4.8.1	Production Federation
T2_PK_NCP	1.9.0	v4.8.3	Transitional Federation
T3_FR_IPNL	1.8.10	v4.2.3	Transitional Federation
T3_TW_NTU_HEP	1.9.0	v4.8.3	Transitional Federation

Status of the GGUS tickets since 1/2017

Site_Name	SB_Version	Xrootd_Version	AAA Federation
T2_AT_Vienna	1.9.0	v4.7.1	Production Federation
T2_FR_GRIF_IRFU	1.9.0		Production Federation
T2_FR_GRIF_LLR	1.9.0	v4.8.2	Production Federation
T2_FR_IPHC	1.9.0	v4.6.1	Production Federation
T2_GR_loannina	1.9.0	v4.8.2	Production Federation
T2_HU_Budapest	1.9.0	v4.8.1	Production Federation
T2_IN_TIFR	1.9.0	v4.8.2	Production Federation
T2_PL_Swierk	1.9.0	v4.8.3	Production Federation
T2_PL_Warsaw	1.9.0	v4.7.1	Production Federation
T2_RU_INR	1.9.0	v4.7.1	Production Federation
T2_TR_METU	1.8.11	v4.8.1	Production Federation
T2_UA_KIPT	1.8.11	v4.4.1	Production Federation
T2_UK_London_Brunel	1.9.0	v4.8.1	Production Federation
T2_PK_NCP	1.9.0	v4.8.3	Transitional Federation
T3_FR_IPNL	1.8.10	v4.2.3	Transitional Federation
T3_TW_NTU_HEP	1.9.0	v4.8.3	Transitional Federation

Not many tickets overall
for DMP sites
1.5 year period of
time!!!

Count of CMS Site/Notified Site	
Row Labels	Total
T2_AT_Vienna	2
T2_FR_GRIF_IRFU	2
T2_FR_IPHC	3
T2_GR_loannina	1
T2_HU_Budapest	3
T2_IN_TIFR	3
T2_PL_Swierk	2
T2_PL_Warsaw	5
T2_RU_INR	3
T2_RU_PNPI	1
T2_TR_METU	8
T2_TW_NCHC	3
T2_UK_London_Brunel	3
T2_UK_London_Brunel	1
T3_FR_IPNL	1
Grand Total	41

What information can help to debug?

- Info about DPM site hardware and configuration to be compared in order to suggest optimal HW vs SW setup.
- Needed feedback from sites.
- Is it a good idea to collect hardware and configuration info in a webpage for operators?
- DPM parameter setup
 - Help from developers about the tuning of DPM cluster.
 - Is the <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Admin/TuningHints> page updated?
- Is the available documentation about how to join the AAA federation clear enough?

Summary

- We do a lot in terms of monitoring! And supporting sites!
- A complete “debug” system allows AAA to evaluate the reliability of a site to be included in the production federation or to move-maintain it in the transitional one.
- The correct evaluation of test results and the debug of problems require the collaboration of site-manager and backend developers.
 - Would like to increase sites’ connectivity to AAA to guarantee success going forward!
- As this becomes routine of operations simpler debug of AAA failure should be the new target for smarter maintenance of the complete system
- CMS is continuously exploring the current performance of remote sites joined AAA federation
- With the collaboration of site managers, storage backend developers and the AAA team a lot can be done, thank you!

Support

- **Users:**
 - CompOps Transfer Team using GGUS and support unit: CMS AAA - WAN Access
 - Hypernews list: hn-cms-wanaccess@cern.ch
- **Regional redirector admins:**
 - Global: cms-service-xrootd-global@cern.ch
 - US: cms-service-xrootd-us@cern.ch
 - EU: cms-service-xrootd-eu@cern.ch

Backup slides



Debugging and cross-check

- Debug of failed results is really time consuming, sometimes error messages don't identify in a clear way the reason of problems (site setup, redirector problem, old hardware, etc...)
- To exclude temporary or more general problems a cross check is necessary with older AAA results and with results of other tests run on site
 - **Site readiness:** example of 100% failure when site was in downtime

		T2_BE_IIHE																					
Site Readiness Status:		R	R	R	R	R	R	R	SD	SD	SD	SD	R	R	R	R							
Daily Metric:		O	O	O	O	O	O	O	O	O	O	SD	SD	SD	SD	O	O	O	O				
Maintenance:		Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	SD	SD	SD	SD	Up	Up	Up	Up				
HammerCloud:		na	100%	100%	100%	100%	100%	100%	100%	100%	100%	na	na	na	100%	100%	100%	100%	100%				
SAM Availability:		100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	25%	0%	0%	100%	100%	100%	100%	100%				
Good T2 links from T1s:		14/15	14/15	14/15	14/15	14/15	14/15	14/15	14/15	14/15	14/15	0/0	0/0	0/0	14/15	14/15	14/15	14/15	14/15				
Good T2 links to T1s:		14/15	14/15	14/15	14/15	14/15	14/15	14/15	14/15	14/15	14/15	0/0	0/0	0/0	14/15	14/15	14/15	14/15	14/15				
Active T2 links from T1s:		13	13	13	13	13	13	13	13	13	13	13	13	13	13	13	13	13	13				
Active T2 links to T1s:		13	13	13	13	13	13	13	13	13	13	13	13	13	13	13	13	13	13				
Waiting Room:		out	out	out	out	out	out	out	out	out	out	out	out	out	out	out	out	out	out				
		21	22	23	24	25	26	27	28	29	30	01	02	03	04	05	06	07	08	09	10	11	12

Report made on 2015-10-12 04:35:01 (UTC)

- **SAM tests:** to check the xrootd access and fallback results that should be coherent with AAA scale tests
- **HC tests:** to check the capacity of a site to provide files for fallback solution (8028 error code)
 - every 2 days HC submits real analysis jobs running on all possible sites and reading input files available only on one specific SE
- if the reason of problem is still not clear the help from site manager is needed → GGUS ticket

How to control federated access?

- **We use CMS site name information from SiteDB**
- **Production federation:**
 - make list of allowed sites join the federation (access list) and distribute across regional level of redirectors
 - mapping of domains for allowed sites, e.g. US region, makes config:


```
cms.allow host *.fnal.gov
cms.allow host *.mit.edu ...
```
 - the list should consist of sites which meet criteria of periodic scale testing
- **Transitional federation:**
 - no access list needed, let anyone join (by default help T3 sites)
 - use as temporary room
 - for the new sites joining (later) production till they pass criteria
 - for the sites which were disallowed in production federation based on bad results of the tests and other metric defined from SAM and HC tests