CMS Update

XRootd @ Tokyo University, 11/8/2016

Matevž Tadel, UCSD

30 kft view for 2016

- 14% of Production and Analysis jobs uses AAA
    - same percentage in core-hours
  - Failure rate for AAA is 2x higher (13% vs. 6%)
  - CPU efficiency hit:
    - 5% for production (73% vs. 68%)
    - 14% for analysis (71% vs. 57%)
- During 2016 AAA traffic grew from 25% to close to 50% of all transfers.
- In general happy … *cmsd* lockups are the only thing that is really killing us at this point.

What follows are slide selections from

1. WLCG/CHEP talk of David Lange

   about multi-threading and AAA usage

2. ------””-------------   Brian Bockelman

   about AAA usage and plans

3. CMS Data Mgmnt review talk by Siddharth Narayanan

   about monitoring / operational procedures

SELECTED SLIDES FROM:

CMS software and infrastructure performance and efficiency gains

David Lange

October 10, 2016

Pilot efficiency considerations

- **Pilot scheduling efficiency is closely monitored. Generally ~90%**
  - Job pressure variation (e.g., contracting global pool) impacts efficiency
  - Mix of single and multi threaded jobs fragments pilots
- **Going more and more multithreaded is helping too**
  - 2015+ Tier-0 and data re-reco is multithreaded
  - 2016+ DIGI-RECO MC now transitioned to multithreaded
  - 2017+ GEN-SIM will be multithreaded
  - Legacy configurations to remain single threaded (in most cases). These will naturally reduce over time as analysis is more focused on 2016+ given large data samples accumulated this year

I/O impacts job performance

- More important to monitor threaded jobs:
  - Reads from ROOT soruces are serialized, so impact of any latencies scales with N threads.
    - Recently parallelized FlushBaskets saving ~10% of total time for processing Tier-0
- Time for random (ROOT) file reads can dramatically improve if file is written in optimal way (to let you read and uncompress just what you need)
- Impact of saturating site bandwidth on CPU efficiency is big
  - Improving I/O of pileup simulation via "premixing"

Pileup simulation: Premixing now deployed in CMS production

- Premixing concept: Instead of doing pileup simulation during DIGI-RECO processing, we do it once to make a “premixed data set”.
  - Process ~0.25B heavy I/O events in a single dedicated production instead of ~6B during DIGI-RECO campaign.
  - Combine premixed data set with hard scatter events during digitization processing
- After long validation and production integration/testing period, premixing is now deployed
- Results
  - Premixing reduces the I/O from pileup in digitization step by **40x**
  - The premixed data set is big (250M, 0.5 PB)
    - We do not have a copy of it everywhere. Often use AAA to read it
    - PU scenarios needed for special studies will use old DIGI-RECO method
  - **Large increase in DIGI job efficiency (even with remote reads)**

SELECTED SLIDES FROM:

CMS Input

Brian Bockelman

How can we reduce data & storage costs?

- CMS has aggressively invested in software and computing infrastructure to be able to **stream data from remote centers**.
    - In Run2, more data-intensive workflows are “WAN-ready”.
- The large bulk of our organized processing workflows *no longer require local data*.
    - Local data provides a modest CPU efficiency boost.
    - Current reprocessing campaign: **Average transfer rates are ~.5Mbps/core**; instantaneous rates are ~35Mbps/core.
- Production can effectively utilize **large CPU resources without storage** for production campaigns.
- **Stageout not as robust** and efficient as streaming inputs. Note that Amazon EC2 has higher costs for stageout than streaming inputs.

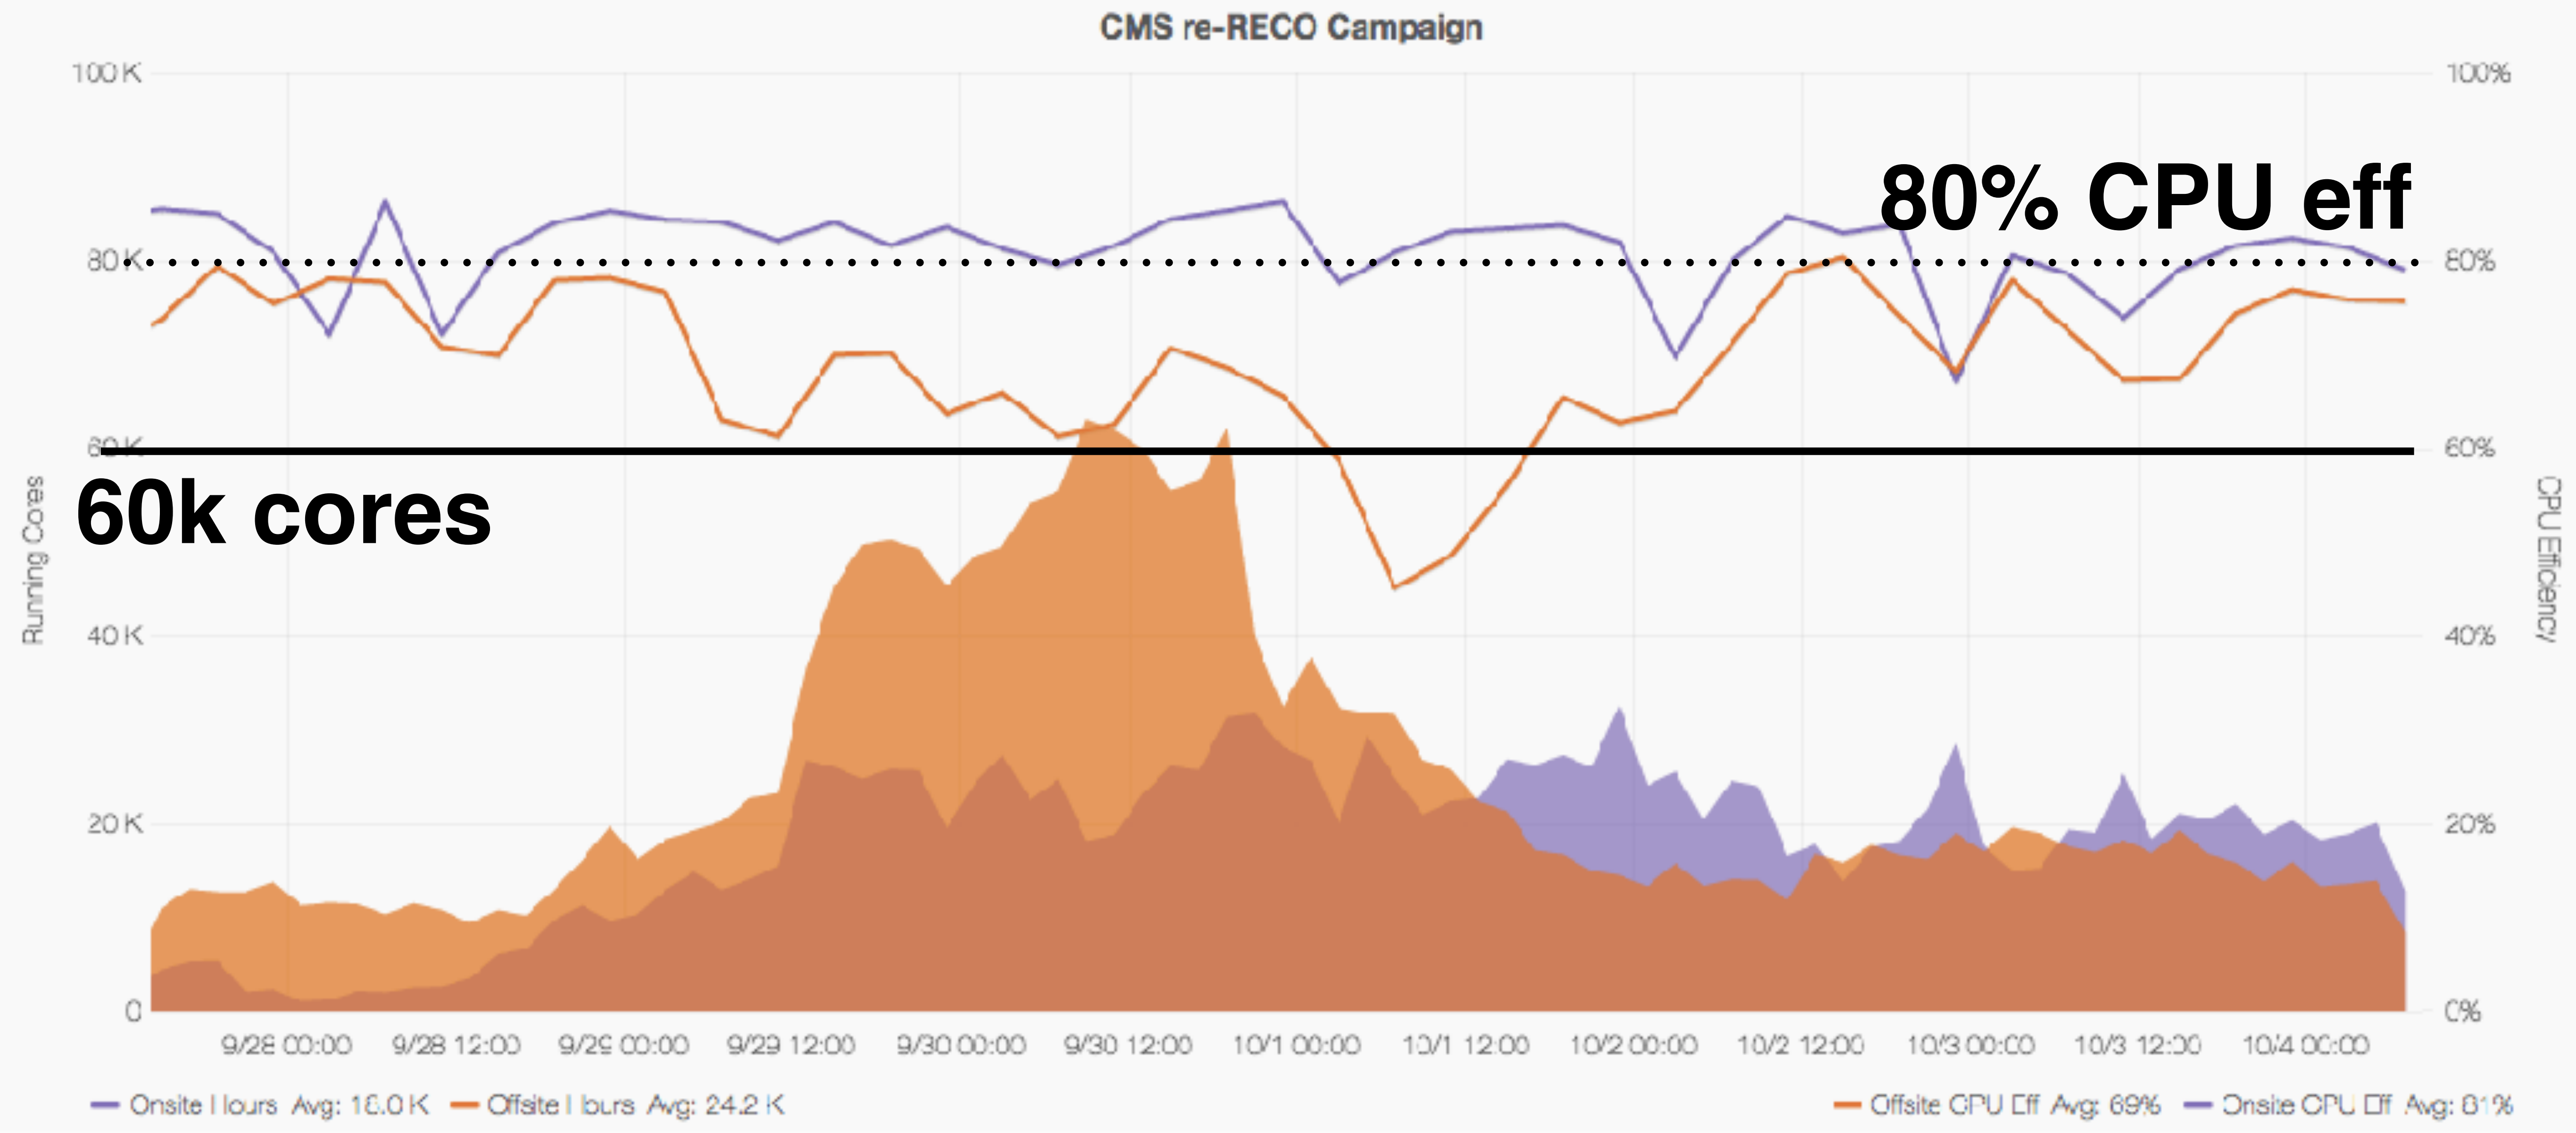

80% CPU eff

60k cores

Current CMS reprocessing campaign:

- By streaming remotely, more sites can participate.
- More core-hours spent on “offsite” jobs than onsite.
- Offsite CPU efficiency hit is noticeable but manageable.
- CPU efficiency somewhat independent of data rates.

Data Analysis

- Analysis jobs are far more difficult to characterize and outlier workflows can span multiple order-magnitude of scale.
    - A recent user submitted a workflow requiring 2.6B file-opens…
    - Many well-formed user analysis jobs can effectively stream over the WAN (there are clusters of sites that are “nearby” in RTT). **Many user analysis jobs cannot**.
- Question: can we instead **cache** the analysis datasets locally? That would potentially provide much higher bandwidth and provide protection against “clever users”.

Caching Pilot Goals

- Evaluate XRootd Caching Software at large scale:
    - Performance evaluation, i.e., simultaneous read/write vs number of clients
    - Operational cost evaluation: operate it in production for 3-6 months before data taking restarts in 2017.
    - Write interim report after 2-3 months: decide on whether to continue.
    - Write final report after 6 months: continue ops if deemed worthwhile.
- Use findings as input to future planning

Scale of Pilot

- 300+ TB in SoCal (Caltech & UCSD)
- 30-50 Gbps IO performance for simultaneous read/write with up to 20,000 clients reading.
  - 100+ disks, diverse hardware & filesystems
  - estimate to need ~20 disks to fill 10Gbps pipe
- Operations Phase:
  - Host part of CMS namespace, most likely all MINIAOD from currently most used release
  - default access for all jobs in SoCal to hit cache for this namespace

Interested in collaborations with others!!

Take-home message:

CMS can more effectively utilize CPU-only sites

CMS is working to simplify storage management, particularly at smaller T2 sites.

CMS is interested in continued improvements in the effectiveness of its use of storage

SELECTED SLIDES FROM:

AAA from the Ops Perspective

Siddharth Narayanan (for the CMS Data Transfer and AAA Teams)

Data Management Review

Redirector status

- Tracks all production (global, US, EU) and transitional redirectors
  - DNS aliases
  - Individual machines
- Failures trigger email notification ⇒ reduces shifter responsibilities

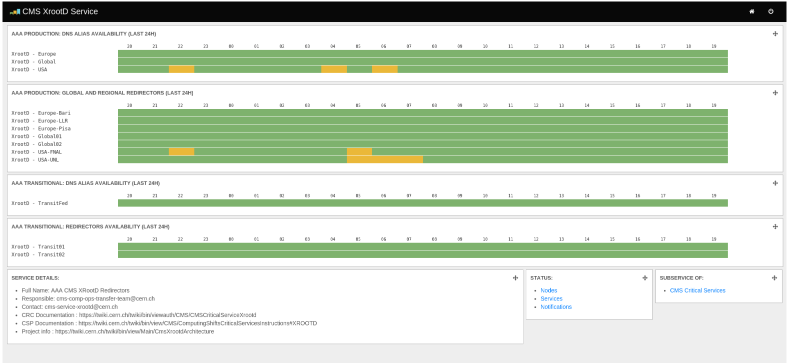

AAA PRODUCTION: DNS ALIAS AVAILABILITY (LAST 24H)

AAA PRODUCTION: GLOBAL AND REGIONAL REDIRECTORS (LAST 24H)

AAA TRANSITIONAL: DNS ALIAS AVAILABILITY (LAST 24H)

AAA TRANSITIONAL: REDIRECTORS AVAILABILITY (LAST 24H)

SERVICE DETAILS:

- Full Name: AAA CMS XRootD Redirectors
- Responsible: cms-comp-ops-transfer-team@cern.ch
- Contact: cms-service-xrootd@cern.ch
- CRC Documentation : https://twiki.cern.ch/twiki/bin/viewauth/CMS/CMSCriticalServiceXrootd
- CSP Documentation : https://twiki.cern.ch/twiki/bin/view/CMS/ComputingShiftsCriticalServicesInstructions#XROOTD
- Project info : https://twiki.cern.ch/twiki/bin/view/Main/CmsXrootdArchitecture

STATUS:

- Nodes
- Services
- Notifications

SUBSERVICE OF:

- CMS Critical Services

Site information - summary and detail

- MonALISA-based monitoring at `http://xrootd.t2.ucsd.edu/`
- Each machine reports summaries via UDP every 30 seconds
- Summaries of traffic, connections, authentications, etc
- Can also investigate specific sites and nodes: `xrootd` version, CPU usage, local configuration, etc[1]
  - Used by Operations to get detailed information to provide support and identify sites which need to update configuration

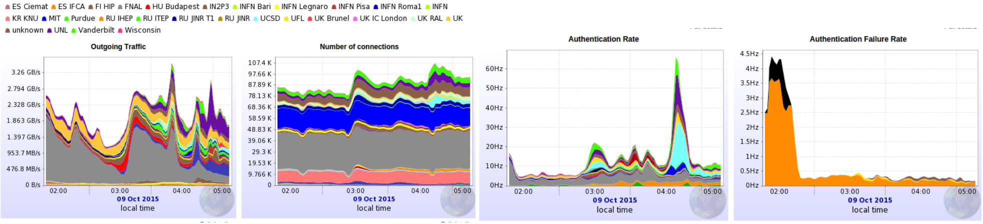

[1] `twiki.cern.ch/twiki/bin/view/CMSPublic/CompOpsAAAOperationsGuide#Filtered_queries`

Data transfer information

- Dashboard[1] provides fine-grained information on transfers
- Part of WLCG-wide transfer monitoring
- Used mostly by Operations and site admins

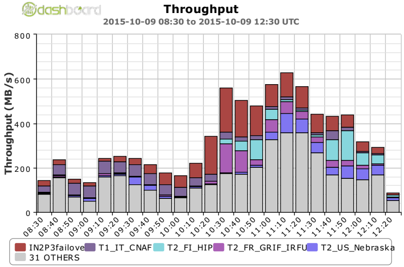

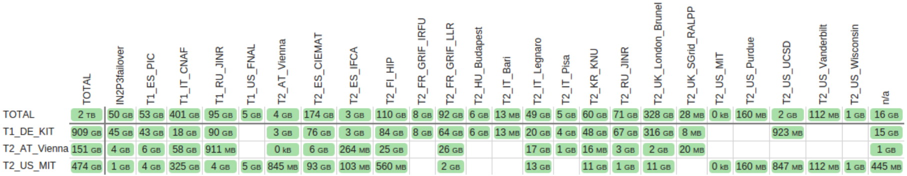

	TOTAL	IN2P3failover	T1_ES_PIC	T1_IT_CNAF	T1_RU_JINR	T1_US_FNAL	T2_AT_Vienna	T2_ES_CIEMAT	T2_ES_IFCA	T2_FI_HIP	T2_FR_GRIF_IRFU	T2_FR_GRIF_LLR
TOTAL	2 TB	50 GB	53 GB	401 GB	95 GB	5 GB	4 GB	174 GB	3 GB	110 GB	8 GB	92 GB
T1_DE_KIT	909 GB	45 GB	43 GB	18 GB	90 GB		3 GB	76 GB	3 GB	84 GB	8 GB	64 GB
T2_AT_Vienna	151 GB	4 GB	6 GB	58 GB	911 MB		0 kB	6 GB	264 MB	25 GB		26 GB
T2_US_MIT	474 GB	1 GB	4 GB	325 GB	4 GB	5 GB	845 MB	93 GB	103 MB	560 MB		2 GB

	T2_HU_Budapest	T2_IT_Bari	T2_IT_Legnaro	T2_IT_Pisa	T2_KR_KNU	T2_RU_JINR	T2_UK_London_Brunel	T2_UK_SGrid_RALPP	T2_US_MIT	T2_US_Purdue	T2_US_UCSD	T2_US_Vanderbilt	T2_US_Wisconsin	n/a
TOTAL	6 GB	13 MB	49 GB	5 GB	60 GB	71 GB	328 GB	28 MB	0 kB	160 MB	2 GB	112 MB	1 GB	16 GB
T1_DE_KIT	6 GB	13 MB	20 GB	4 GB	48 GB	67 GB	316 GB	8 MB			923 MB			15 GB
T2_AT_Vienna			17 GB	1 GB	16 MB	3 GB	2 GB	20 MB						1 GB
T2_US_MIT			13 GB		11 GB	1 GB	11 GB		0 kB	160 MB	847 MB	112 MB	1 GB	445 MB

[1] http://dashb-cms-xrootd-transfers.cern.ch/ui/

Production federation

- Available at `http://dashb-ssb.cern.ch/dashboard/request.py/siteviewhistorywithstatistics?columnid=224`
- Detailed explanation in previous talks[2]
- Site is marked as failing if 2 week average of SAM (Hammercloud) `xrootd` tests is less than 50% (80%) or if any AAA-related GGUS tickets are older than 2 weeks
- Detail view allows us to check if sites have fixed problem recently, even if criteria are failing

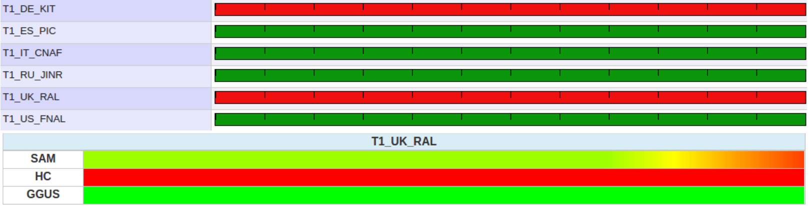

T1_DE_KIT	🟥
T1_ES_PIC	🟩
T1_IT_CNAF	🟩
T1_RU_JINR	🟩
T1_UK_RAL	🟥
T1_US_FNAL	🟩

T1_UK_RAL	
SAM	🟩
HC	🟥
GGUS	🟩

[2] `https://indico.cern.ch/event/448825/session/12/contribution/38/attachments/1169371/1687674/AAA_Production_and_Transitional_Federation_Monitoring.pdf`

Transitional federation

- For T3s and any T2s that do not pass production federation metric
- Available at `http://dashb-ssb.cern.ch/dashboard/request.py/siteviewhistorywithstatistics?columnid=219`

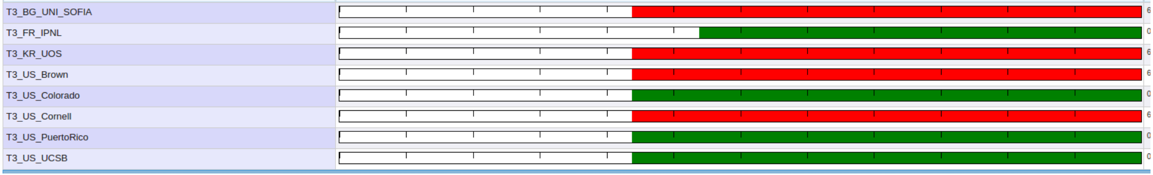

T3_BG_UNI_SOFIA	🟥	6
T3_FR_IPNL	🟩	0
T3_KR_UOS	🟥	6
T3_US_Brown	🟥	6
T3_US_Colorado	🟩	0
T3_US_Cornell	🟥	6
T3_US_PuertoRico	🟩	0
T3_US_UCSB	🟩	0

Site availability monitoring

- Results can be accessed here
- Two types of `xrootd` tests:
  - Access: access files located at the site
  - Fallback: access remote files from the site
- SAM tests used for assessing AAA production status and site readiness
  - `CMS_CRITICAL_FULL` for US and EU sites
- Because SAM test is now critical, have extra checks to lower false error rate:
  - Check global redirector is up before performing test
  - Re-try test if it fails due to a transient issue

Site scaling tests ran once per week

Operations teams

- Operations support is provided by the Transfer Team and AAA experts
- Entire TT is involved: Aram, Jorge, Engin, Sebastian, and myself
- From AAA, we have Marian, Ken, Brian, Federica, Matevz, and Carl
- TT acts as ‘front line’ for GGUS tickets
  - If the issue is global or concerns redirectors, we refer the issue to AAA experts
  - If it is site-specific:
    - We attempt to debug the issue ourselves with the site admin
    - Usually involves a change in the site’s configuration
    - If the problem is deeper, we again refer to AAA team
- TT also opens maintenance tickets for sites, for example:
  - Ask sites to change which XrdMon collector they use
  - Inform sites of changes necessary to enter AAA or production federation

CMS requests for XRootd development

1. We need a way to trace down locks that get held for excessive amounts of time.
   a. Lightweight solution for important global locks – so it can run in production;
   b. Heavyweight solution – use timed-lock on every mutex lock to support checking of all locks (requiring special build is ok)
2. Focus on code correctness … do not attempt to over optimize thread synchronization.
3. Release Caching Proxy V2 as soon as possible Needed by CMS, OSG / StashCache, PRP / LHC@UC