



# **Storage at Caltech**

## **September 22, 2010**

**Michael Thomas**



# Aggregate resources



## 3 CMS clusters, each running Hadoop for local storage

Cluster	Raw TB	Usable TB	Used TB
Tier2	1054 <sup>1</sup>	527	818 (77%)
Tier3-A	15.25 <sup>2</sup>	7.62	9.2 (60%)
Tier3-B	43.25	21.62	39.57 (91%)

<sup>1</sup>+120TB pending

<sup>2</sup>+60TB pending



# Services



**Bestman SRM – Gateway mode, shared with FrontierDB**

**PhEDEx – shared with FrontierDB**

**Namenode – Metadata Service shared with condor collector/negotiator**

**Secondary NN – Filesystem checkpoint merges, on dedicated server**

## **54 GridFTP servers**

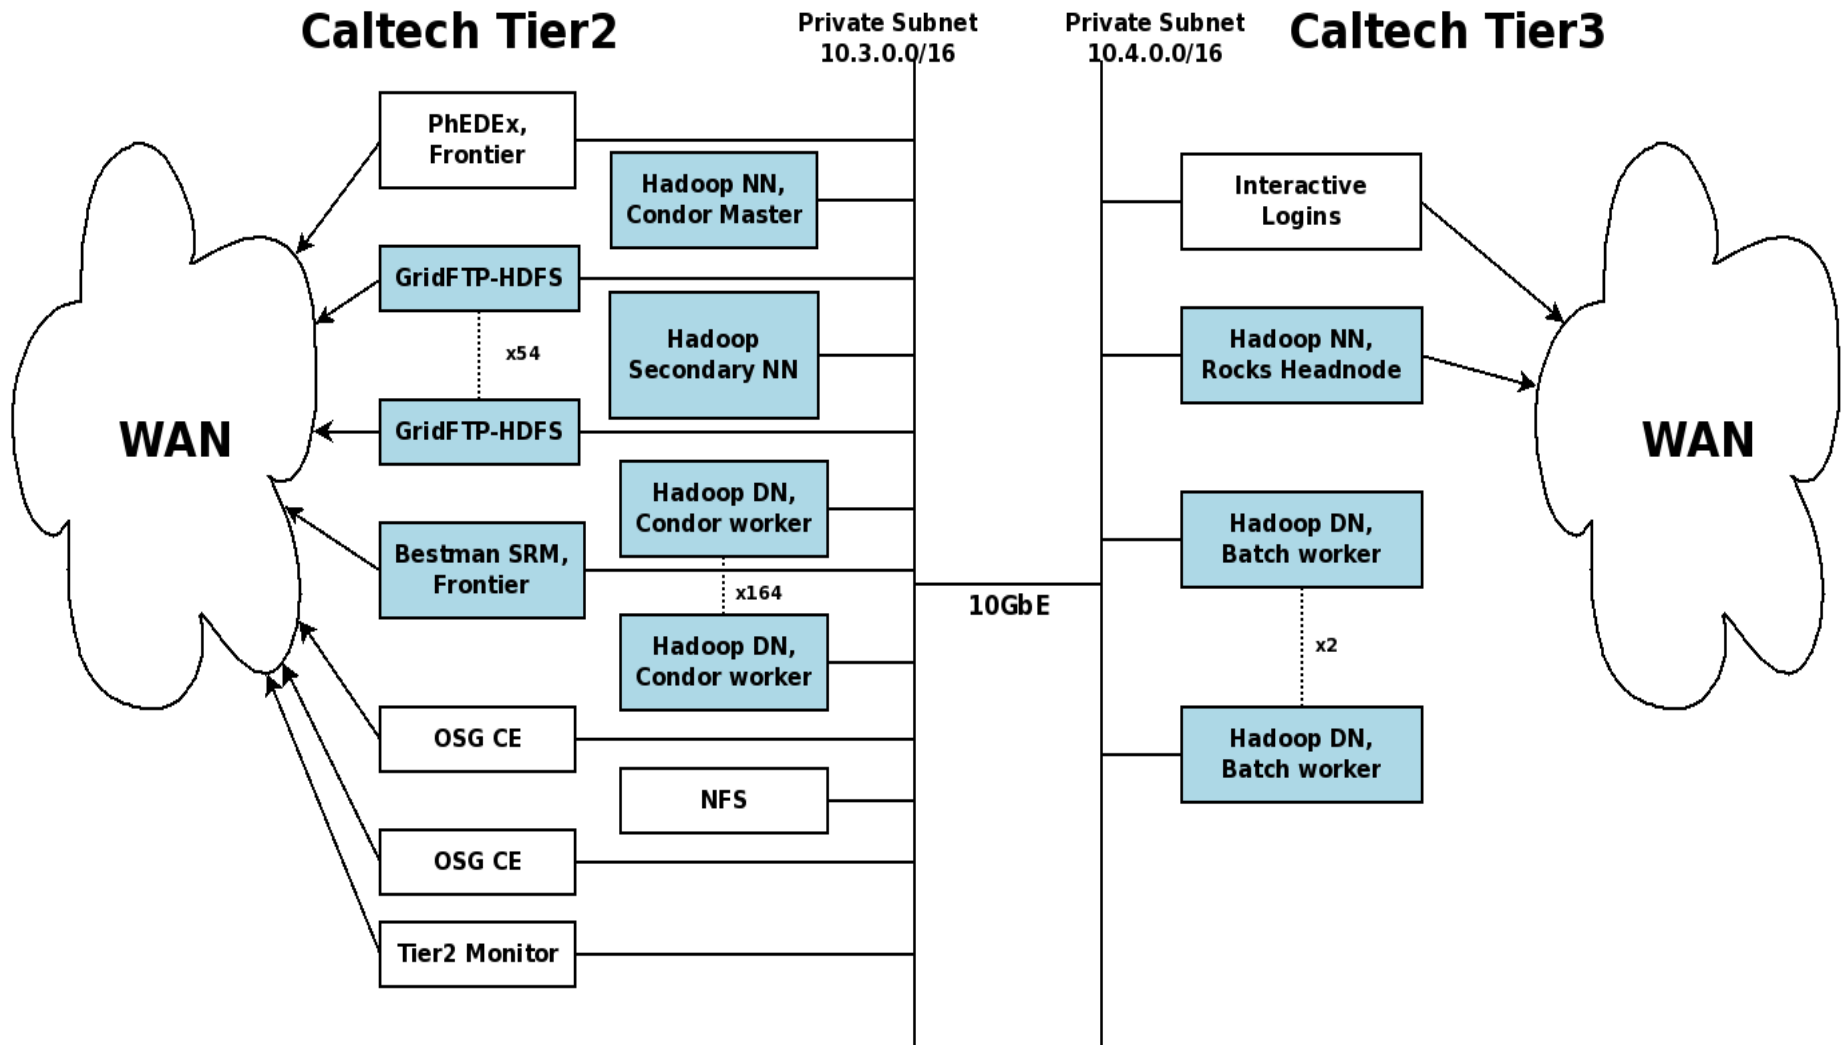
- \* 4 x 10GbE dedicated
- \* 8 x 1GbE dedicated
- \* 42 x 1GbE shared worker nodes

## **164 datanodes**

- \* 16 dedicated disk servers (including 2 Sun Thumpers)
- \* 148 shared with worker nodes
- \* All disks 1TB or larger



# Cluster diagram

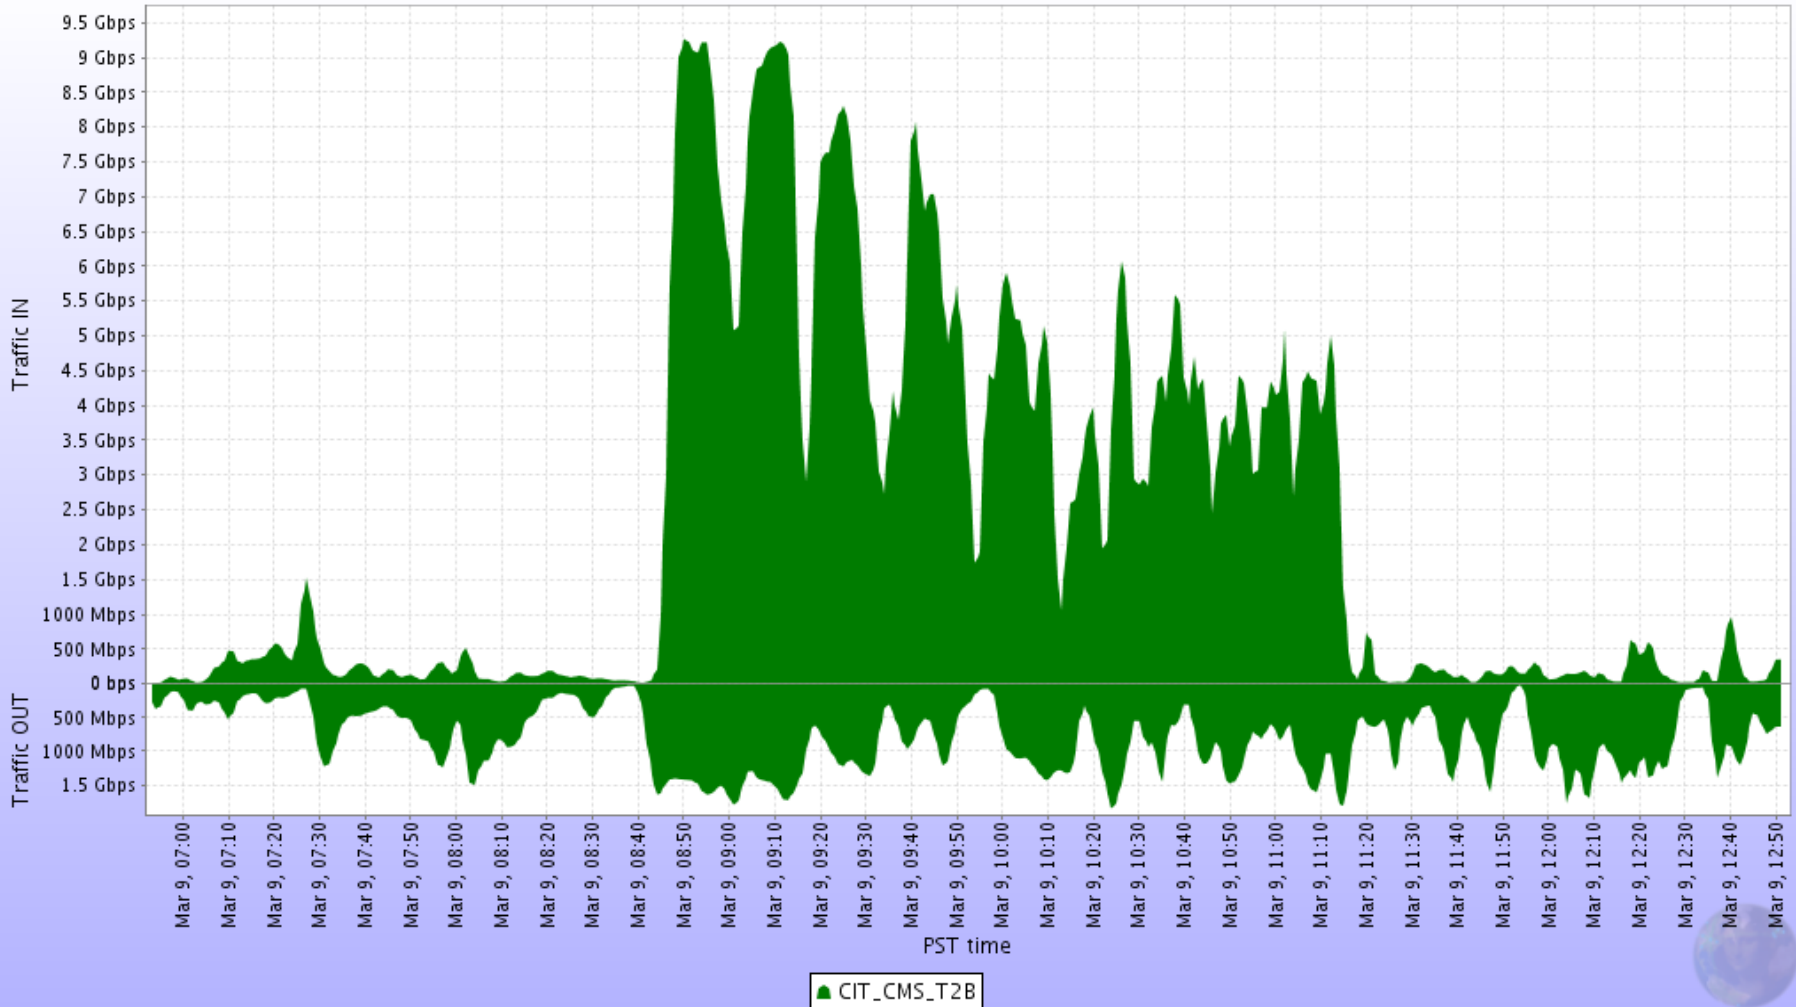




# Performance



### External traffic





# Daily Operations - Automated

**Cron – Balance datanode block distribution once/day**

**Cron – Force replication across entire FS once/day**

**Cron – Incremental replication across partial FS 3/day**

**Cron – cmsSync to detect lost files twice/day**

**Nagios – Wait for alerts**

- ★ SNN not processing checkpoints
- ★ Namenode not running
- ★ gridftp services down

**Email – Wait for user complaints**



# Daily Operations - Manual



## Approve/Reject PhEDEx requests

## Reset datanodes

- \* Our kernels like to panic 2-3 nodes/day, depending on temp

## Fix misreplicated blocks (due to datanode resets)

## Check storage monitoring pages

- \* Namenode status page
- \* SAM test results
- \* Aggregate WAN traffic
- \* PhEDEx errors



# Daily Operations - Manual




The Hadoop Chronicle - Mozilla Firefox 3.5 Beta 4

File Edit View History Bookmarks Tools Help

caltech.edu https://cms.hep.caltech.edu/hadoop/ Google

Most Visited EVO docs Caltech T2 Fedora

The Hadoop Chronicle



Selected or last chronicle

2009\_06\_26\_08:30

=====

The Hadoop Chronicle | 46 % | Fri Jun.26.2009 08:30

=====

Global storage

-----

Configured Capacity: 191813069336576 (174.45 TB)  
 Present Capacity: 191707600577536 (174.36 TB)  
 DFS Remaining: 102718547960182 (93.42 TB)  
 DFS Used: 88989052617354 (80.94 TB)  
 DFS Used%: 46.42%

-----

Path	Size(GB)	#Files	#Dirs
/store/PhEDEx_LoadTest07	667	262	668
/store/data	24337	33474	462
/store/mc	2353	2146	15
/store/unmerged	438	3416	172
/store/user	8461	25234	433

-----

User area

-----

Path	Size(GB)	#Files	#Dirs
/store/user/burt	0	2	1
/store/user/chorbo	902	5341	127
/store/user/dkcira	0	17	13
/store/user/dorian	41	286	1
/store/user/hpi	2	6	21
/store/user/ligioi	0	2	1
/store/user/litvin	0	3	8
/store/user/oatramen	3178	1036	77
/store/user/ssekmen	0	4	4
/store/user/test	0	2	5
/store/user/tucker	0	17	6
/store/user/uscms0377	10	7	1
/store/user/uscms0755	614	2754	3
/store/user/vlitvin	3709	15754	153
/store/user/wart	1	3	1

-----

System health

-----

Total size: 38929932017766 B (Total open files size: 3087007744 B)  
 Total dirs: 1765  
 Total files: 64551 (Files currently being written: 2)  
 Total blocks (validated): 358503 (avg. block size 108590254 B) (Total open file blocks (not validated): 23)  
 Minimally replicated blocks: 358503 (100.0 %)  
 Over-replicated blocks: 63 (0.017573075 %)  
 Under-replicated blocks: 0 (0.0 %)  
 Mis-replicated blocks: 0 (0.0 %)  
 Default replication factor: 2  
 Average block replication: 2.349721  
 Corrupt blocks: 0  
 Missing replicas: 0 (0.0 %)

All Chronicles

- 2009\_06\_26\_08:30
- 2009\_06\_25\_19:42
- 2009\_06\_25\_08:30
- 2009\_06\_24\_19:42
- 2009\_06\_24\_08:30
- 2009\_06\_23\_19:42
- 2009\_06\_23\_08:30
- 2009\_06\_22\_19:42
- 2009\_06\_22\_08:30
- 2009\_06\_21\_19:42
- 2009\_06\_21\_08:30
- 2009\_06\_20\_19:42
- 2009\_06\_20\_08:30
- 2009\_06\_19\_19:42
- 2009\_06\_19\_08:30
- 2009\_06\_18\_19:42
- 2009\_06\_18\_08:30
- 2009\_06\_17\_19:42
- 2009\_06\_17\_08:30
- 2009\_06\_16\_19:42
- 2009\_06\_16\_08:30
- 2009\_06\_15\_19:42
- 2009\_06\_15\_08:30
- 2009\_06\_14\_19:42
- 2009\_06\_14\_08:30
- 2009\_06\_13\_19:42
- 2009\_06\_13\_08:30
- 2009\_06\_12\_19:42
- 2009\_06\_12\_08:30
- 2009\_06\_11\_19:42
- 2009\_06\_11\_08:30
- 2009\_06\_10\_19:42
- 2009\_06\_10\_08:30
- 2009\_06\_09\_19:42
- 2009\_06\_09\_08:30
- 2009\_06\_08\_19:42
- 2009\_06\_08\_08:30
- 2009\_06\_07\_19:42
- 2009\_06\_07\_08:30
- 2009\_06\_06\_19:42
- 2009\_06\_06\_08:30
- 2009\_06\_05\_19:42
- 2009\_06\_05\_08:30
- 2009\_06\_04\_19:42
- 2009\_06\_04\_08:30
- 2009\_06\_03\_19:42
- 2009\_06\_03\_08:30
- 2009\_06\_02\_19:42
- 2009\_06\_02\_08:30
- 2009\_06\_01\_19:42
- 2009\_06\_01\_08:30
- 2009\_05\_31\_19:42
- 2009\_05\_31\_18:55
- 2009\_05\_31\_18:33
- 2009\_05\_31\_17:52

Michael Thoma

cms.hep.caltech.edu FoxyProxy: Patterns





# Daily Operations - Manual



t2-headnode.ultrallight.org

File

Main

Log file:

Server	Job PID	userid	User Name	Remote host	Start Date	End Date	Size	Rate	# buffers	direction	file
compute-16-6	9134	uscms1264	Wei Li	cmsfts1.fnal.gov	Tue Sep 21 03:35:36 PM PDT 2010	Tue Sep 21 03:47:05 PM PDT 2010	2736269993	3.79MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
compute-10-9	17712	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 03:35:39 PM PDT 2010		0	0	1/1		/store/PhEDEx_LoadTest07/LoadTest07_Debug_RU_IHEP/US_Cattech/10/LoadTest07_L
compute-16-10	19812	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 03:35:40 PM PDT 2010	Tue Sep 21 04:03:46 PM PDT 2010	2684354560	1.52MB/s	1/1	in	/store/PhEDEx_LoadTest07/LoadTest07_Debug_PL_Warsaw/US_Cattech/7/LoadTest07_
compute-10-13	3452	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 03:35:40 PM PDT 2010	Tue Sep 21 03:46:54 PM PDT 2010	2684354560	3.8MB/s	1/1	in	/store/PhEDEx_LoadTest07/LoadTest07_Debug_RU_JINR/US_Cattech/12/LoadTest07_J
compute-10-15	2018	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 03:35:40 PM PDT 2010		0	0	1/1		/store/PhEDEx_LoadTest07/LoadTest07_Debug_UK_SGrid_RAL/PP/US_Cattech/11/Load
compute-10-15	2026	cmsprod006	stuart wakefield	lcgfts02.gridpp.rl.ac.uk	Tue Sep 21 03:35:41 PM PDT 2010	Tue Sep 21 03:57:24 PM PDT 2010	2761391150	2.02MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
compute-10-18	5122	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 03:35:44 PM PDT 2010	Tue Sep 21 03:43:39 PM PDT 2010	2684354560	5.39MB/s	1/1	in	/store/PhEDEx_LoadTest07/LoadTest07_Debug_CN_Beijing/US_Cattech/12/LoadTest07_
compute-10-33	3497	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 03:35:46 PM PDT 2010		0	0	1/1		/store/PhEDEx_LoadTest07/LoadTest07_Debug_IN_TIFR/US_Cattech/6/LoadTest07_IN_
compute-10-12	10925	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 03:35:55 PM PDT 2010	Tue Sep 21 03:55:50 PM PDT 2010	2684354560	2.14MB/s	1/1	in	/store/PhEDEx_LoadTest07/LoadTest07_Debug_RU_PNPI/US_Cattech/3/LoadTest07_P
compute-16-8	16306	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 03:36:02 PM PDT 2010		0	0	1/1		/store/PhEDEx_LoadTest07/LoadTest07_Debug_IT_Rome/US_Cattech/17/LoadTest07_R
compute-16-9	11271	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 03:36:57 PM PDT 2010		0	0	1/1		/store/PhEDEx_LoadTest07/LoadTest07_Debug_RU_SINP/US_Cattech/4/LoadTest07_SI
compute-16-3	14317	uscms1413	Francesco Pandolfi	fts01-sc.cr.cfnaf.infn.it	Tue Sep 21 03:40:42 PM PDT 2010	Tue Sep 21 04:00:35 PM PDT 2010	2736560411	2.19MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
compute-10-31	8595	uscms1108	Massimo Biasotto	fts03-sc.cr.cfnaf.infn.it	Tue Sep 21 03:41:39 PM PDT 2010	Tue Sep 21 03:45:43 PM PDT 2010	2732349350	10.68MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
gridftp-11-18	19240	uscms1093	Marco Dias	spstrj01.sprace.org.br	Tue Sep 21 03:42:51 PM PDT 2010	Tue Sep 21 03:52:08 PM PDT 2010	2765457002	4.73MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
compute-10-31	9200	uscms1093	Marco Dias	spstrj03.sprace.org.br	Tue Sep 21 03:43:11 PM PDT 2010	Tue Sep 21 03:52:54 PM PDT 2010	2756889671	4.51MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
gridftp-14-3	22690	phedex	Eduardo Azevedo Revoredo	cmsfts1.fnal.gov	Tue Sep 21 03:48:28 PM PDT 2010	Tue Sep 21 03:52:20 PM PDT 2010	2742659189	11.27MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
compute-10-34	18541	uscms2014	Aaron Nowack	polgrid118.in2p3.fr	Tue Sep 21 03:48:35 PM PDT 2010		0	0	1/1		/store/user/nowack1/mono_skim/full_run/mono_photon_skim_799_1_r6H.root
compute-10-29	6089	uscms358	Jon A Bakken	cmsssrv116.fnal.gov	Tue Sep 21 03:48:37 PM PDT 2010		0	0	1/1		/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
compute-16-5	21856	phedex	Eduardo Azevedo Revoredo	cmsfts1.fnal.gov	Tue Sep 21 03:48:46 PM PDT 2010	Tue Sep 21 03:53:19 PM PDT 2010	2737722083	9.56MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
compute-16-14	5185	uscms358	Jon A Bakken	cmsssrv116.fnal.gov	Tue Sep 21 03:51:15 PM PDT 2010	Tue Sep 21 03:56:9285	34.68MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech	
gridftp-11-19	12176	uscms2014	Aaron Nowack	polgrid85.in2p3.fr	Tue Sep 21 03:51:18 PM PDT 2010	Tue Sep 21 03:59:26 PM PDT 2010	4800007939	9.54MB/s	1/1	in	/store/user/nowack1/mono_skim/full_run/mono_photon_skim_1093_1_dKS.root
gridftp-11-22	18700	uscms358	Jon A Bakken	cmsssrv116.fnal.gov	Tue Sep 21 03:51:36 PM PDT 2010		0	0	1/1		/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
gridftp-11-23	27896	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 03:51:52 PM PDT 2010		0	0	1/1		/store/PhEDEx_LoadTest07/LoadTest07_Debug_UK_London_Brunel/US_Cattech/9/Load
gridftp-14-2	21314	phedex	Dorian Kcira	cmsfts1.fnal.gov		Tue Sep 21 03:52:09 PM PDT 2010	0	0	1/1	in	/store/PhEDEx_LoadTest07/LoadTest07_Debug_UA_KIPT/US_Cattech/1/LoadTest07_KI
gridftp-11-23	28000	phedex	Dorian Kcira	cmsfts1.fnal.gov		Tue Sep 21 03:52:09 PM PDT 2010	0	0	1/1	in	/store/PhEDEx_LoadTest07/LoadTest07_Debug_UA_KIPT/US_Cattech/1/LoadTest07_KI
compute-16-2	14430	phedex	Dorian Kcira	cmsfts1.fnal.gov		Tue Sep 21 03:52:11 PM PDT 2010	0	0	1/1	in	/store/PhEDEx_LoadTest07/LoadTest07_Debug_UA_KIPT/US_Cattech/1/LoadTest07_KI
gridftp-14-3	23075	uscms2014	Aaron Nowack	cabinet-6-6-23.i2.ucsd.edu	Tue Sep 21 03:53:24 PM PDT 2010	Tue Sep 21 03:57:05 PM PDT 2010	5050538738	21.79MB/s	1/1	in	/store/user/nowack1/mono_skim/full_run/mono_photon_skim_1928_1_3B9.root
gridftp-14-2	21472	phedex	James R Letts	cmsfts1.fnal.gov	Tue Sep 21 03:55:17 PM PDT 2010	Tue Sep 21 03:57:07 PM PDT 2010	2729735588	23.67MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
compute-10-28	6392	uscms2014	Aaron Nowack	polgrid85.in2p3.fr	Tue Sep 21 03:56:31 PM PDT 2010	Tue Sep 21 04:02:03 PM PDT 2010	4745325867	13.63MB/s	1/1	in	/store/user/nowack1/mono_skim/full_run/mono_photon_skim_1096_1_UVa.root
compute-16-10	3671	cmsprod006	stuart wakefield	lcgfts02.gridpp.rl.ac.uk	Tue Sep 21 03:57:56 PM PDT 2010	Tue Sep 21 04:04:21 PM PDT 2010	2736124784	6.78MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
gridftp-14-4	1283	phedex	Preston M. Smith	dcache-11.rcac.purdue.edu	Tue Sep 21 04:00:46 PM PDT 2010	Tue Sep 21 04:01:53 PM PDT 2010	2762698031	39.32MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
compute-10-19	27519	phedex	Preston M. Smith	dcache-09.rcac.purdue.edu	Tue Sep 21 04:00:50 PM PDT 2010	Tue Sep 21 04:02:57 PM PDT 2010	2758922597	20.72MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
compute-10-33	24490	uscms0570	zhang xiaomel	ccicgfts05.in2p3.fr	Tue Sep 21 04:00:52 PM PDT 2010	Tue Sep 21 04:03:19 PM PDT 2010	2763714494	17.93MB/s	1/1	out	/mnt/hadoop/store/PhEDEx_LoadTest07/LoadTest07_Prod_Cattech/LoadTest07_Cattech
compute-10-33	3786	uscms2014	Aaron Nowack	cabinet-6-6-31.i2.ucsd.edu	Tue Sep 21 04:04:18 PM PDT 2010	Tue Sep 21 04:06:26 PM PDT 2010	5368699570	40.0MB/s	1/1	in	/store/user/nowack1/mono_skim/full_run/mono_photon_skim_2088_0_QZ.root
gridftp-11-22	19339	uscms358	Andrea Sciacca	vocms36.cern.ch	Tue Sep 21 04:06:25 PM PDT 2010	Tue Sep 21 04:06:26 PM PDT 2010	41472	40.5KB/s	1/1	in	/store/annerged/SAM/testSRM/SAM-cit-se.ultrallight.org/cg-uti/testfile-cp-notoken-20
compute-10-27	13912	cmsprod002	Andrea Sciacca	vocms36.cern.ch	Tue Sep 21 04:07:16 PM PDT 2010	Tue Sep 21 04:07:16 PM PDT 2010	41472		1/1	out	/mnt/hadoop/store/annerged/SAM/testSRM/SAM-cit-se.ultrallight.org/cg-uti/testfile-cj
compute-10-21	22413	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 04:07:46 PM PDT 2010	Tue Sep 21 04:08:35 PM PDT 2010	2684354560	52.24MB/s	1/1	in	/store/PhEDEx_LoadTest07/LoadTest07_Debug_US_Florida/US_Cattech/1361/LoadTest
gridftp-11-19	12949	phedex	Dorian Kcira	fta03.pic.es	Tue Sep 21 04:07:54 PM PDT 2010		0	46/85			/store/PhEDEx_LoadTest07/LoadTest07_Debug_PIC/Cattech/130/LoadTest07_PIC_B2_U
compute-10-17	1331	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 04:07:53 PM PDT 2010		0	1/1			/store/PhEDEx_LoadTest07/LoadTest07_Debug_US_Florida/US_Cattech/1361/LoadTest
compute-16-1	18877	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 04:07:51 PM PDT 2010	Tue Sep 21 04:08:58 PM PDT 2010	2684354560	38.21MB/s	1/1	in	/store/PhEDEx_LoadTest07/LoadTest07_Debug_UCSD/Cattech/1180/LoadTest07_UCSD
gridftp-14-3	24019	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 04:07:51 PM PDT 2010		0	1/1			/store/PhEDEx_LoadTest07/LoadTest07_Debug_UCSD/Cattech/1180/LoadTest07_UCSD
gridftp-14-4	1752	phedex	Dorian Kcira	fts01-sc.cr.cfnaf.infn.it	Tue Sep 21 04:07:54 PM PDT 2010		0	308/308			/store/PhEDEx_LoadTest07/LoadTest07_Debug_IT_CNAAF/US_Cattech/122/LoadTest07_I
gridftp-11-24	3295	phedex	Dorian Kcira	fts01-sc.cr.cfnaf.infn.it	Tue Sep 21 04:07:54 PM PDT 2010		0	88/277			/store/PhEDEx_LoadTest07/LoadTest07_Debug_IT_CNAAF/US_Cattech/122/LoadTest07_I
compute-10-13	20144	cmsprod002	Andrea Sciacca	vocms36.cern.ch	Tue Sep 21 04:07:55 PM PDT 2010	Tue Sep 21 04:07:55 PM PDT 2010	41472		1/1	in	/store/annerged/SAM/testSRM/SAM-cit-se.ultrallight.org/cg-uti/testfile-gt-mg-gt-notok
compute-16-1	18945	phedex	Dorian Kcira	lcgfts02.gridpp.rl.ac.uk	Tue Sep 21 04:07:56 PM PDT 2010		0	7/7			/store/PhEDEx_LoadTest07/LoadTest07_Debug_UK_RAL/US_Cattech/38/LoadTest07_R
compute-10-14	29408	phedex	Dorian Kcira	fta03.pic.es	Tue Sep 21 04:07:57 PM PDT 2010		0	29/83			/store/PhEDEx_LoadTest07/LoadTest07_Debug_PIC/Cattech/129/LoadTest07_PIC_F5_H
compute-10-15	15356	phedex	Dorian Kcira	lcgfts02.gridpp.rl.ac.uk	Tue Sep 21 04:07:56 PM PDT 2010		0	7/7			/store/PhEDEx_LoadTest07/LoadTest07_Debug_UK_RAL/US_Cattech/38/LoadTest07_R
compute-16-1	19000	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 04:08:08 PM PDT 2010		0	1/1			/store/PhEDEx_LoadTest07/LoadTest07_Debug_UA_KIPT/US_Cattech/1/LoadTest07_KI
compute-10-11	4967	phedex	Dorian Kcira	cmsfts1.fnal.gov	Tue Sep 21 04:08:09 PM PDT 2010		0	1/1			/store/PhEDEx_LoadTest07/LoadTest07_Debug_UA_KIPT/US_Cattech/1/LoadTest07_KI
gridftp-11-22	19495	cmsprod002	Andrea Sciacca	vocms36.cern.ch	Tue Sep 21 04:08:39 PM PDT 2010	Tue Sep 21 04:08:39 PM PDT 2010	41472		1/1	in	/store/annerged/SAM/testSRM/SAM-cit-se.ultrallight.org/cg-uti/testfile-gt-notoken-201
compute-16-13	28799	phedex	Dorian Kcira	cmsfts1.fnal.gov		Tue Sep 21 04:09:10 PM PDT 2010	0	1/1			/store/PhEDEx_LoadTest07/LoadTest07_Debug_UA_KIPT/US_Cattech/1/LoadTest07_KI

gridftp-14-4: [1752] Tue Sep 21 16:09:24 2010 :: Initializing buffer number 311.



# Site-Specific Features



**HTTP + mod\_ssl + fuse + gums-client == Secure read-only browser access for CMS**

**Xrootd on 12 gridftp servers**

**Automount of fuse\_dfs**

**Export Tier2 HDFS to Tier3**



# Storage R&D



**Build RPMS and maintain Yum repository for the OSG,  
with Abhishek**

**Maintain koji+mock build system for simplified  
package/yum repo management**

**Working on Hadoop 0.20 stack of packages**



# **FDT+HDFS+PhEDEx**



**Development of a HDFS adapter to FDT to buffer and write non-sequential data to HDFS**

**Integration of PhEDEx and FDT**

**Conveniently packaged as RPMs**

**Deployed at Caltech and Nebraska for Debug transfers**



# Future Plans

## **Add second SRM server (round-robin)**

- \* Upgrade to bestman2

## **Continue upgrading disks**

- \* Also adding 121TB (raw) from new worker nodes
- \* 1TB -> 3TB when available

## **Production deployment of FDT/PhEDEx**

## **Fix those darned fuse\_dfs memory leaks!**

## **Better storage accounting (by VO group)**

## **Upgrade to Hadoop 0.20**