



CD FY08 Tactical Plan Status

FY08 Tactical Plan Status Report for dCache

Gene Oleynik

6/3/2008

FY08 Tactical Plan for dCache



- *Relevant Strategic Plan(s):*
 - FY08-09 Strategic Plan for Data Storage and Caching
 - Strategic Plan for Scientific Facilities (2007-2009)
- *Tactical Plan Leader:* Timur Perelmutov
- *Organizational Unit home:* Scientific Facilities
- *Tactical Plan Goals*
 - Satisfy the emerging needs of Fermilab, LHC experiments and other Fermilab supported scientific endeavors in storage by implementation and improvements of the dCache software system as well as its associated namespace software and the development of the future dCache features and versions of the namespace service.

Activities Summary: FTEs

Personnel Usage (FTEs)							
Tactical Plan	Level 1 Activity	Allocation		Actual YTD			
		FTE-yrs	FTE-mos	FTE-yrs (Ave/mo)	FTE-mos (7 mo)	% consumed YTD	% consumed sans other contributors
dcache	Collaboration	0.10	1.2	0.44	3.11	259.17	141.67
	Development	1.00	12	0.24	1.69	14.08	12.42
	Evaluations	0.00	0	0.96	6.73	0.00	
	Development /Resilient Dcache	0.50	6	1.03	7.2	120.00	119.00
	Operations	0.80	9.6	0.36	2.49	25.94	25.94
	Project Management	0.12	1.44	0.37	2.56	177.78	63.19
	Support	0.40	4.8	0.38	2.63	54.79	53.13
Gene, Matt	Project Leadership	0.40	4.8	0.00	0.03	0.63	
Total		3.32	39.84	3.78	26.44	66.37	45.48

Project Activity: Resilient Manager Development

dCache Resilient Manager—improvements for integrity, stability, monitoring

Original Goals Related to this Activity

- *Improve the scalability, reliability and performance of the resilient manager*
 - **Completed and deployed at US-CMS T1**

Adjusted Goals Related to this Activity

- *Same*

Project Activity: Resilient Manager Development

- *Key Milestones*

1. Improve stability, reliability, performance of resilient manager. **Done**
2. Improve and develop monitoring and administration interface. **Done**
3. Improve installation and configuration procedures. **Done:** Procedures were developed and tested on RM test stand. Migration of DB schema procedures developed. Installation of monitoring.
4. Develop documentation. **Done:** All Procedures documented in dCache book
5. Integrate Resilient Manager with SRM v2.2 spaces. **Not complete.** May not be required by customers. Estimate 2 week effort.
6. Expand functionality including installation specific customization of replication rules. **Not done.** May not be required by customers. Freezing functionality

Project Activity: dCache Development

dCache Resilient Manager—improvements for integrity, stability, monitoring

Original Goals Related to this Activity

- *Collaboratively develop and deploy a production level dCache system with the critical features agreed upon with the stakeholders (CMS, LQCD, OSG, WLCG, CEDS)*
 - **In progress**
- *Perform investigations to exploit new technology and caching systems.*
 - **In progress**

Adjusted Goals Related to this Activity

- *Same*

- *Same*

Project Activity: dCache Development

Key Milestones:

1. Commission dCache 1.8 with SRM 2.2 interface to US-CMS, OSG and WLCG (depends on both DESY and FNAL). **Done** (phase I)
2. Deploy a scalable, performant, backwards compatible namespace with native large file support. **Under investigation** through storage evaluation wg
3. Deploy dCache.org RPM at Public dCache at FNAL, Q1: **Done** with configuration assistant based on Bakken's configuration scripts
4. Improve GridFTP scaling and interoperability – **Done**. Completed bi-directional adaptor. NGDF has since taken over parts of gridFTP work.
5. Investigate other caching systems, underlying file systems and I/O scheduling systems – **In progress** – storage evaluation. Evaluated GPFS, evaluating Lustre, alternative namespaces.
6. Implement GridFTP Version 2 features including data integrity verification . **Done** Get, out completed by NGDF. Integrity verification by FNAL (Andri Baronovski)

Project Activity: dCache Development (cont.)



Key Milestones:

7. Improve Fermilab specific dCache monitoring, alarms and scans and convert those utilities into generic dCache tools maintained and distributed centrally. **In progress**
Q3: Incremental pnfs scans for stken. **Not complete** for cdfen as well
8. Improve robustness against OS/file system weaknesses and destructive usage patterns, such as implementing per node service limits. **Not completed.**
9. Develop/update procedures for cleanup/recovery from cpu, disk and other failures. **Not complete.** Replacement node installation & configuration procedures simplified
10. Collaborate in the SCIDAC2/CEDS collaboration to define and design how dCache will fit into the Data Placement Services **Not done.**

Project Activity: Collaboration

dCache Resilient Manager—improvements for integrity, stability, monitoring

Original Goals Related to this Activity

- *Collaborate with DESY in the development and software release process*
 - Ongoing

Adjusted Goals Related to this Activity

- *Same*

Project Activity: Collaboration



Key Milestones

1. Continuously improve life-cycle system for quick turnaround of bugfix releases and full-system testing of major releases. **Ongoing. Resilient Manager development process includes robust testing before release.**
2. Improve installation infrastructure (including VDT). **Completed** by VDT team with consultation.
3. Work with DESY to supply the critical features needed by our stakeholders, including per-node Service limits. **Not Complete:** Work has been done to fix issues, but per-node limits not implemented.

Service Activity: Support

Level 1 and 2 Support of Fermilab dCache systems

Original Goals Related to this Activity

- *Continue transition of routine dCache support to SSA and provide scripts and tools to ease and simplify operational maintenance*
 - **Change of Goal**
- *Provide developer-level support (debugging, installation assistance) of agreed-to production dCache systems*
 - **Ongoing**

Adjusted Goals Related to this Activity

- *Goal changed into hiring a developer devoted to support work and support development*
 - **Lost position due to budget**
- *Same*

Service Activity: Operations

Key Milestones

1. Test major releases of dCache across a multi-node test bed. **In progress**
2. Provide primary operational and configuration support, to Fermilab dCache Systems, e.g. CDF and Public dCache Systems at Fermilab, **Ongoing**
3. Provide secondary operational and configuration support for groups that have their own primary support, in both Fermilab and to Fermilab supported experiments e.g. US CMS T1, WLCG dCache Systems, etc. **Ongoing**
4. Provide Tertiary (high level expert support) to agreed to installations - e.g. OSG T2 storage **Ongoing**
5. Establish an operations bug tracking system. **Completed.**
6. Transition of routine dCache support to SSA Q2 . **Not Done** (no new hire)
7. Timely test collaborative point releases intended for production deployment at Fermilab in production Q1. **Not complete**
8. Achieve uniformity in agreed-to production dCache systems in Fermilab, including monitoring, alarms and scans Q4. **Not complete**

Tactical Plan Issues and Risks

- *Resolution of Past Action Items*
 - Review the level of testing done on v1.8 of dCache. The SRM-specific testing mentioned does not cover most past problems on local supported dCache, systems which kept developer-intensive operations at a high level. **We are in the process of doing this for the public dCache 1.8 upgrade**
- *Issues and risks*
 1. dCache maintenance must be passed on to the SSA administrators so that developers can focus on the more complex support issues and long-term development and stability work. **This is still an issue.**
 2. Failure to deliver these milestones will leave stakeholders underserved and result in great operational problems, given the absence of viable alternatives. Progress on certain milestones depends on acceptance by and contributions from DESY.

Tactical Plan Issues and Risks

- *Issues and risks*
3. dCache may not scale or operate as needed. Alternative file systems and namespaces need to be explored. So far it looks like there is no suitable alternative for HEP. **Should focus be applied to dCache?**
 4. While a certain level of release formalism has been achieved, the robust, comprehensive and continuous testing on a test-bed is still absent, which results in increased operations load with each software upgrade, long periods of manual testing of release candidates, and inability to deliver features to customers as quickly as they are wanted. However, DESY, OSG, CERN, and LBNL do testing—the latter two of the grid interface only. **With the current level of staffing and furloughs, implementation of a rigorous and formal release process has not been not realistic. dCache is however being tested at more sites as the collaboration has grown, including OSG testing.**

Tactical Plan Status Summary



- Very good progress in stabilizing resilient dCache and dCache 1.7 and 1,8 (dCache 1.7 is very stable on the CDF system)
- Expected two hires that were lost, with the following impact:
 - operations support (still handled by core developers)
 - development of operation software and tools
 - core improvements like per node resource limits
- Progress has been and is being made on major core issues (proper accounting of used space, pin manager issues, etc), but there is more to do.
- Significant effort expended in storage-evaluation for alternative file systems. Looks promising for HPC but not in the near future for HEP.