Documentation and Knowledge Retention in LCLS

15 March 2018

Oliver Hickman

Background:

On October 4th 2017 LCLS held an offsite meeting titled “LCLS Operations Improvement Workshop” where staff were assigned into three working groups: *Operational Delivery Structure*, *Enabling Tools and Techniques*, and *Process control*. Within the Enabling Tools and Techniques group, several key areas were identified for substantive process creation or improvement, among them “document control”.

Document control is one element of the much larger need for knowledge retention in LCLS.

Knowledge includes the information in people’s heads, written procedures, information contained in engineering drawings, comments in code and the code itself, and even the information contained in the very construction of our instruments.

As LCLS grows and increases the throughput of experiments, it becomes critical that we can learn from our operations and iterate to keep LCLS competitive. It is also important that we find ways to operate more efficiently, including handing off some duties of performing experiments to users. To facilitate both of these items LCLS needs a knowledge retention system, and because of the complexity of the organization, the knowledge retention system must be one that fits this organization. This report and the investigation that informed it seeks to highlight best path forward for knowledge retention in LCLS as directed by the charge in Appendix A.

Method:

Between 11 December 2017 and 25 January 2018 Bob Sublett and I interviewed 23 individuals in LCLS; interviewees included both individual contributors and supervisors from groups throughout LCLS.

All interviewees were informed that their comments would be kept confidential and that findings from the interviews would be shared with LCLS management in aggregate and permission would be acquired before attributing quotes.

Interviewees were invited to interpret terms such as “group” in a way that they felt would give the most meaningful answers (as an example for myself “group” could mean “area managers”, “MEC”, or “LCLS Operations” and is not limited to just these).

We began asking the questions listed in Appendix 2, after about 18 interviews when the themes we saw were very apparent, we instead began asking about the onboarding process and how staff learns about what knowledge is and where knowledge is kept.

Findings:

While conducting these interviews strong themes quickly emerged about the state of knowledge retention in LCLS. In general, staff is interested in retaining knowledge both to be able to use it to iterate and improve design and operations more quickly and also to reduce the workload associated with being the only person who can perform a given task. Groups and individuals are hungry to find a way to store and share knowledge and numerous attempts to find tools have been made. These individual efforts have led to a very fractured knowledge environment where staff may not even know where to begin to look for information. Appendix 3 contains a list of all the locations which were identified where knowledge is stored in LCLS, and it is expected that if more people were interviewed that this list would grow.

Across all staff interviewed we saw that the status of knowledge does not seem to change what is retained (status means how important the knowledge is seen to be for operating the facility) and that verbal communication is the most common method of knowledge transfer in LCLS.

With one exception on-the-job training is an informal process and new staff learn by being paired with a more experienced staff member. Most interviewees volunteered that knowledge siloed with one staff member is lost when that staff member leaves LCLS. Breaking the chain of siloed information is difficult: without proper documentation, the one expert will frequently find that it is easier to perform a task than to delegate that task.

“I don't want to be the only one who knows how to fix this thing and I don't want to have to do it every time.”

Ultimately staff did not feel that knowledge retention was a core component of their job. Several management factors appear to lead to this belief, among them are that management at all levels does not emphasize knowledge retention. This is seen where documentation is not included as a deliverable in projects and project planning does not include documentation (Large projects like L2SI being the exceptions), and where P-A’s are frozen before documentation is complete.

“We're trying to write procedures and checklists, but we don't have a good place to put them.”

We asked specifically about what information was not being retained and what would give the most value if it were retained.

Among the knowledge that we heard was not being retained:

- As-built drawings
- Changes to configurations (motor channels, code changes, etc.)
- Configuration change logs
- Deployment notes
- Design notes

- Experiment configurations
- Installation notes
- Lessons learned
- Machine status for experiments
- Problem reports
- Redlines (reviewed engineering drawings)
- Service history for equipment
- Software change logs

Two outliers in the field of knowledge retention are

- Mechanical Engineering where *Teamcenter* (a project lifecycle management tool by Siemens) is used to store notes and calculations related to a project, though capturing data in this is not universal in that group and the information stored in Teamcenter is not accessible outside the group.
- CDS where *Confluence* (a team collaboration software by Atlassian) is a heavily used tool by some members of the group. While some in this group contribute frequently to this knowledge base, the search function is poor enough that it is often difficult for the uninitiated to find the information they are looking for, and effectively impossible for staff outside the group to find information. Several staff said that they had essentially given up on trying to find things in Confluence because accurate information was hard to find and it was mixed with out of date information, and despite the existence of plug-ins to improve the search functionality, these have not been installed.

We also encountered other factors which aggravate the shortcomings in knowledge retention. Participants in projects may change; this is particularly problematic in the design phase where the information about why design decisions were made is then lost. Not having standardized equipment and code bases across hutches effectively increases the amount of knowledge required to run LCLS, meaning that when people move between hutches the knowledge they hold from another hutch has a diminished utility in the new hutch. The merging of science groups in SRD means that people more frequently move between hutches. The merging of the groups also aids in distributing knowledge as staff gain familiarity with new areas.

As is widely observed in LCLS the pace of work is very high and often doesn't leave time to record lessons learned during experiments. While some SRD groups have begun distributing end-of-shift emails, this serves primarily to spread the current operational knowledge of a hutch and to a lesser extent to broaden the visibility of problems which may be building. The culture of reflecting on an experiment in its entirety is slowly building with the use of tools like the experiment post mortem, but not all relevant information is captured there.

"it all falls apart when you...put that shutter in."

“I think it’s just the pace at which we do things doesn't allow for thoughtful reflection.”

Additionally, as LCLS matures into a steady state facility the role of staff changes. Staff who once felt personally invested in each experiment now see the pipeline of experiments as too large to invest in each individually; this again diminishes the feeling of the need to retain knowledge. This change in mentality has not been countered by a refinement or restatement of our mission.

Do the minimum for the experiment and then move on...that's been our operational mode for the past two or three years.

Outside of SRD and operations the pace of development of our facility is also very rapid. Engineers, designers, technicians, associates, etc. all reported feeling like they worked too fast and were very quickly moved to new jobs as the old job was nearing completion, rarely reconvening for a design postmortem after an instrument is commissioned.

“there is no feedback...unless it’s a total failure”

To this point it would be easy to conclude that everything uncovered in this enquiry was negative, however most staff feel confident in their ability to find the information that they need most of the time. Most staff also recognizes the need for an improved knowledge retention system and many had even spent time thinking about what a solution might look like.

Recommendations:

While LCLS exists to create new knowledge and in part measures this success in the number of proposals submitted, user groups awarded time, and published papers, we don’t have a mechanism or dominant culture of creating, retaining and sharing knowledge within our organization. As other enquiries into LCLS operations have concluded in other areas, our failing in this area is masked by the tremendous commitment of our staff and their incredible *firefighting* abilities.

“when you fight fires you tend to spend a lot of time investigating and there is a lot of knowledge gained there, but we don't retain that knowledge.”

Overall LCLS senior management needs to emphasize the importance of knowledge retention as a part of operations and of projects, and reinforce this emphasis with setting workloads, project planning, budgeting, and providing tools to facilitate knowledge retention. The specific recommendations are as follows.

General:

- Standardize knowledge retention systems across hutches and departments - this reduces the number of locations where someone might search for information and increases the likelihood of finding information in a search. Options include CDMS, Confluence, Wiki media software.
- Standardize calendaring system to single application or to interoperable applications - this would relieve the need for staff who use the LCLS Gateway to maintain private Gmail accounts which then get used for other LCLS business and thus become shadow repositories of knowledge.
- Establish on-boarding routine for staff at all levels to disseminate locations of knowledge in LCLS.
- Appoint or hire a documentation librarian with the authority to sign off or re-assign incomplete documentation, who can unify templates across the facility and support staff in retaining knowledge. Without a single gatekeeper who oversees the entire knowledge retention system it is likely that different groups would establish different protocols, templates, and styles making accessing knowledge from different groups more difficult.
- Develop library of “SMART” goals related to documentation to give staff a vocabulary to discuss documentation and tools to measure its completeness. See Appendix 4 for a discussion on measurable aspects of documentation.

There is ‘not retaining’, and not ‘retaining well’.

Projects:

- Include documentation in project planning and list it as a deliverable.
- Budget for documentation and ensure that P-A’s stay active until documentation is complete.
- Include transition to operation plan in design phase
- Include installation instructions in the design process
- Develop engineering, design, and controls specifications (similar to vacuum specifications)
- Develop feedback system where problems and changes encountered during installation, commissioning, and operation are communicated to the designer.

Operations:

- Include knowledge retention as a step in a work ticket
- Create a universally accessible repository for documentation
- Formalize location for operations feedback which may be used to influence future engineering decisions
- Mandate use of ServiceNow for equipment maintenance records

Other recommendation(s):

- Expand job descriptions (formerly R2A2s) to explicitly state who is responsible for knowledge retention and documentation.

In the execution of these recommendations we identified several traps: currently the use of non-enterprise Google drives breaks all central control of permissions and staff who are unaware of the implications can grant permission to group folders to users who should not get such broad visibility. This highlights the need to move our systems away from those that depend on private accounts at services outside of LCLS and the need for centrally managed permissions. We also may lose access to non-enterprise knowledge repositories when staff leave the organization.

Requirements:

A widely deployed knowledge retention system would need to account for some enterprise needs and some management issues. Among those are:

- Management must both fully support and navigate rolling out a system that some staff may not want to use and which will not provide tangible benefits for a period of time until a critical mass of knowledge is contained in it.
- Staggered approach for approving and releasing documents or updating knowledge based on the importance and visibility
    - formal documents for safety with approval workflow (BLAs with rigid and robust approval process)
    - Operation manuals which may have initial release but lower bar for updates (beamline alignment procedures)
    - Working knowledge, (log books, with no formal approval workflow)
- The needs of exchanging information and developing requirements may be different than the needs for archiving the information and systems for communication need to be accounted for even if a very powerful tool for the ultimate storage of knowledge is employed.
- Guidance must be given as to what must be retained. At the extreme end is recording too much where the cost of retaining isn't offset by the gains of having increased knowledge. Staff must feel confident that they are meeting the spirit of the effort when they make decisions about what to retain. For example:
    - Alignment procedures which may reasonably be expected to be repeated should be documented in detail after performing a few times.
    - Experiment setups may be recorded with photographs and block diagrams that allow an expert in the field to reconstruct the experiment in the future if necessary.

    - Any deviation from released drawings must be recorded at a minimum, drawings updated where appropriate.

Final thoughts:

Knowledge retention is a critical component of our organization and it is impossible to develop a working knowledge retention system without impacting other areas of LCLS. Among the management areas that impact knowledge retention are:

- The lack of a chief of staff means that as staff leave or change roles some tasks get dropped and a continuity of that task and its associated knowledge is not accounted for. Self appointed group librarians may not have their duties reassigned after they leave their position.
- Budgeting will be a challenge because we may not have historical information about what documentation cost for past projects where it was not included in planning or where P-A's were prematurely frozen. Historically these costs have been bourn by operations.
- Project planning is often not sufficient to account for the complexity of a project. Staff do not have training in project planning and do not have guidelines to guide the development of a project or to specify what must be included in a project. Documentation is often omitted in project planning. LCLS should develop project planning guidelines that mimic those of the DOE.
- Service history on equipment deployed in LCLS does not exist in any searchable system meaning that the histories of equipment is stored in the heads of individuals or may be exhumed through searching emails, but in both cases if staff leaves then the service history they contain is lost to LCLS.

- Equipment and code (python) is not standardized throughout LCLS. Though standardization is a monumental task, the lack of standardization means that the total amount of knowledge to operate LCLS is increased.

Despite the monumental task of corralling knowledge in LCLS and bringing all parties in to the same system, I believe that it is agreed by an overwhelming number of staff that improved knowledge retention and/or improved searching of existing knowledge would aid our operations. Most of the staff I interviewed were very comfortable talking about the problems they see and were interested in being part of a solution. It is my belief that when a solution is implemented, gaining buy-in from staff will be relatively easy to get.

"People are willing to work smarter, we just don't have a unified way of doing it."

Appendices:

Appendix 1: Charge

1. Document Control Database (lead: Oliver Hickman)
    a. Identify which groups have failed at preserving, managing, or effectively transferring knowledge
    b. Establish possible solution(s) to these cases
    c. Make a recommendation

Appendix 2: Questions

LOCATION:
Where does your group store knowledge?
What system/location do you store knowledge?

PLANNING:
Does your group include documentation in project planning?
Does your group budget for documentation/knowledge retention?
How much time per week/day/project do you (people in your group) spend documenting?

KNOWLEDGE:
What kinds of things does your group retain? Data, instructions, policies, etc.?
What kinds of things are noticeably not retaining?
What kinds of things do you think would add the most value if they were retained?

CULTURE:
What guidelines are you given/do you give your group for storing knowledge?
What are best practices of your group that could be adopted by LCLS?
What does your group need to work on to improve documentation? What would be better guidelines?
Do you store knowledge for yourself that isn't shared with the group? How, why?

REFERENCE:
How often do you query your group's knowledge base?
How readily can you find the information people in your group save? (Specifically, information that you yourself did not put into the system)

How often do you query the knowledge bases of other groups?
How readily can you find the information retained by other groups?

PROBLEMS:
In your opinion what is the biggest impediment to storing the knowledge you think should be stored?
How often do you find yourself reinventing the wheel? Creating solutions to problems that have already been solved but the solution has been lost.

OTHER:
What did we miss that you like / think LCLS/your group does well?
What did we miss that you dislike / think LCLS / your group does poorly?
Who else do you think has strong, well-formed opinions that we should hear?

Appendix 3: List of knowledge retention tools in use

1. Asana
2. Binders (with dead trees inside)
3. Confluence
4. CDMS
5. DAQ
6. DropBox
7. Evernote
8. e-log
9. Email
10. End of Run Summaries
11. FAMIS
12. GitHub (https://github.com/slaclab)
13. Google Calendars (personal accounts and non-enterprise shared accounts)
14. Google Drive (both personal and Stanford issued (Stanford Webdocs))
15. Insite for Design
16. Jira
17. MS Exchange calendars
18. Neocaptar
19. Networked drives (V:\)
20. One Drive
21. Postmortems
22. “The Questionnaire”
23. SEDA
24. SharePoint
25. Slack
26. Subversion
27. Team Drive
28. TeamCenter - for mechanical engineering drawings and knowledge

29. Track (ticketing system)
30. Trello
31. Websites - for outward knowledge dissemination
32. Wikis

Appendix 4: Developing SMART Goals for Documentation

Each of the components of a SMART goal is important for the goal as a whole, however when discussing human knowledge and documentation finding the metric that is measurable is difficult. Overly simplistic metrics for documentation and knowledge retention are at best unhelpful and at worst destructive of the goal itself. Other metrics like user feedback are difficult to measure in a statistically valid way in such a small environment as one organization. Below are a few examples of measurable elements of documentation which may be useful in goal setting.

- Were all fields in documentation template completed, or explanation provided for empty fields.
- Does documentation contain step-by-step list of operations to perform task?
- Are engineering decisions supported by calculations contained in documentation?
- Adoption of procedure by one or more other groups.
- Is documentation free of factual errors?
- Is equipment unambiguously enumerated in documentation?
- Delta in hours billed for completing task with and without documentation.
- Does the documentation adhere to a technical style guide?
- Is documentation free of grammatical errors?
- Does documentation contain photographs and diagrams describing equipment?
- Is documentation structured to be discoverable with a search tool?
- Is documentation correctly tagged?
- Are documents made obsolete by new documentation marked as such or rescinded?
- Can staff execute procedure without additional assistance?

Appendix 5: Other Findings

In addition to what was reported on above, some staff felt that the interviews were a safe space to share other concerns and observations about LCLS. I would be doing a disservice to LCLS and to our staff if I did not share some of the other issues we encountered.

One interviewee highlighted what was seen in LCLS as an overly white-male dominated environment that makes women and minorities feel like outsiders. As a follow up to this I met with Karen Van Der Pyl to pass on the concerns and I followed up with the staff member passing on the resources that Karen provided me. A partial solution might be to treat implicit

bias mitigation training materials the same way we treat other supervisor mandated training: with a course assigned on an STA.

Morale issues were frequently brought up in the interviews, several staff reported not feeling personal ownership of experiments, and contrasted this with the feeling they had early in the operation of LCLS. Others felt that their contribution was not recognized sufficiently.

"*now it's the user's experiment...I [used to feel] like every experiment was my experiment, but now I feel like every experiment is a job."*

We heard that the existence of tools such as SEDA for finding engineering documentation are seen as invaluable in their concept, but the execution is lacking. The search function in SEDA is poor at best and does not allow staff to find the resources they already know exist, much less discover new drawings.

Additionally with the move to SEDA old features, such as the ability to link to the current release of a document were lost and so searches must be repeated to ensure that you are referencing the most current drawing.

"I like the SEDA engine but the search is horrible."

While there was certainly no consensus on the solution, many people felt that the dizzying pace at which we do science and implement upgrades was impressive, but also destructive of our mission.

We have a culture here of expediency, the pressure is always to get things done and get things done now, the next experiment is coming...damn the torpedoes, full speed ahead!

It's very much related to the pace you acquire knowledge, and here we do it very fast.

Out in industry it's slower, so even old documents have value because there are incremental improvements.

for us after three years we may find a new way of doing things and the old way is [totally obsolete]

What kinds of things does your group retain?

Water, he he he. I'm sorry...I shouldn't be allowed out in public.