

TWiki a collaboration tool for the LHC


Peter L. Jones
CERN
IT/PES
Geneva 23, Switzerland
+41 22 7674901
Peter.L.Jones@cern.ch

Nils Høimyr
CERN
IT/PES
Geneva 23, Switzerland
+41 22 7672945
Nils.Hoimyr@cern.ch

ABSTRACT

At the European Laboratory for High Energy Physics, CERN[1], the Large Hadron Collider (LHC)[2] accelerator is colliding beams of protons at energies of 3.5 TeV, recreating conditions close to those at the origin of the Universe. The four main LHC experiments, Alice, Atlas, CMS and LHCb are complex detectors with millions of output channels. These experiment detectors, “large as cathedrals”, have been designed, built and are now operated by collaborations of physicists from universities and research institutes spread across the world.

Wikis are a perfect match to the collaborative nature of CERN experiments and since TWiki[3] was installed at CERN in 2003 it has grown in popularity and the statistics from April 2011 show nearly 10000 registered editors and about 110000 topics (Figure 1). Since the start-up of the LHC more and more users are accessing TWiki requiring better server performance as well as finer control for read and write access and more features. This paper discusses the evolution of the use of TWiki at CERN.



Categories and Subject Descriptors

H.4 [Informations Systems Applications]:

General Terms

Documentation, Human Factors,

Keywords

CERN, LHC, TWiki.

Copyright is held by the author/owner(s).

WikiSym '11, Oct 03-05 2011, Mountain View, CA, USA
ACM 978-1-4503-0909-7/11/10.

1. INTRODUCTION

The World Wide Web was invented at CERN to improve sharing of documents among collaborating physicists on diverse computer systems. Wiki pages, where the user has the possibility to edit the page directly from the browser, is very much in the original democratic spirit of the web. For scientific organizations like CERN, this allows collaborations where users are in different buildings, cities or even countries a convenient way of working together on projects and documents.

TWiki, one of many available Wikis, is a structured one targeting mainly corporate intranets. TWiki was originally chosen because of its project oriented focus and simple setup on computing resources that were readily available at CERN. TWiki runs on Linux machines with the Apache web server and is written in Perl. To keep track of changes to web pages, an RCS (revision control system) back end is used. There are a number of Perl “plugins” available for TWiki that provide additional functionality. The TWiki service quickly became popular at CERN following its introduction and as figure 1 shows the number of topics has grown at a steady rate over the past years.

2. USE CASES

The Atlas[4] collaboration consists of 3000 physicists from more than 174 institutes in 38 countries on 5 continents. These collaborators need efficient means of communicating information. To this end, Atlas has enthusiastically embraced TWiki since 2004 and now has over 14000 web pages some of which are world readable containing technical information about Atlas as well as protected ones for physics preparations and results. New pages are created at a rate of 150/month and averaging over 10,000 updates a month. Atlas creates workbooks with TWiki and the application's working environment allows their users to contribute to the development and maintenance of the documents.

CMS[5] also uses TWiki for creating software guides and workbooks and makes use of the PDF creation feature that allows a one click creation of an entire book. They also benefit from the dynamic web page creation features and implement virtual blackboards that can be written on by users from all around the world.

Both the Atlas and CMS experiment collaborations employ specific review processes to ensure that Wiki content is maintained and current. Certification features are used to ensure document integrity and authors are reminded by email to check the validity of information in their pages, in case it has not been updated by one of their peer researchers.




Figure 2. The Compact Muon Solenoid detector

3. PROBLEMS AND SOLUTIONS

TWiki was initially used at CERN as an Intranet application by a small group of Linux based software developers for their documentation. Once authenticated, using their CERN account, collaboration was classic Wiki free-form, within project areas called “webs”. As the usage of TWiki took off, notably in the Atlas detector collaboration, users requested that TWiki content be made public on the Internet for the benefit of the wider Atlas and High Energy Physics user community. Over the years the service has had to adapt to: increasing demands on the hardware; requests for additional user functionality; requests for more flexibility with access control and the need to integrate the service with the rest of the CERN computing environment.

Performance and search issues

The increasing number of people using TWiki had an effect on the system CPU load and resulted in slower page response times especially for searches. New hardware with more CPU power and memory corrected these performance issues.

The TWiki data back-end was initially on AFS (Andrew File System) that has a limit on the number of files in a particular directory. TWiki now runs on a dedicated cluster of powerful machines and the TWiki data storage was migrated to NFS 4 on NetApp servers.

In order to overcome issues with load-peaks caused by use of the internal TWiki search (based on *grep*), the FAST search engine (Now Microsoft Enterprise Search) in use at CERN for organization wide web-search, was adapted for use within TWiki. This allows for efficient search on protected data, yet retaining full access control to the content.

Data protection versus collaboration

TWiki culture lends to open free-form editing and most pages are world-readable and editable by CERN authenticated users. The experiments benefited from this ease of use and openness. Now integrated with the CERN Single Sign On authentication system external users can edit TWiki documents, which has led to the need for a user vetting procedure.

Since the LHC accelerator has started to provide physics results, researchers felt that some of the data should be protected and made available only to their users. Various levels of access control have been requested from the many small and large projects using TWiki.

TWiki has a system of access control based on specifying which users or groups of users can view and edit documents but this is difficult to manage for large groups. To group users CERN uses a concept called 'e-groups' for many other applications and these are managed by the experiments themselves and by CERN's HR department. By integrating e-groups into TWiki, the experiments can now easily manage the access control. Users are frequently joining or leaving experiments or moving from one group to another and so these e-groups are updated regularly to ensure that users have the correct rights when accessing a document.

Life-cycle of documentation

With regards to the life-cycle of Wiki pages and Webs, we have seen that the large LHC experiment collaborations keep their information alive and have active maintainers for house-keeping. For smaller projects and groups, the trend has been to use a TWiki web for the duration of the project or until the group is re-organized. New groups or projects have often started from scratch with a new TWiki web, instead of updating and evolving the pages of their predecessors. This split in terms of methodology is interesting, and mirrors the work habits of CERN departmental groups versus that of global collaborations with many visitors and short term contributors.

4. THE FUTURE

TWiki is today a key web collaboration tool at CERN. However, as the IT landscape and collaboration needs evolve, upgrades or extended Wiki services may be needed. CERN currently deploys Drupal as a corporate web content management system. TWiki pages will co-exist with the Drupal setup and also feed contents into Drupal sites. As TWiki is a structured Wiki and with the back-end based on simple flat files it would be easy to import into another database or Wiki application if required.

5. CONCLUSIONS

TWiki has proven to be a simple and easy tool to use for the creation and development of collaborative documents. The Wiki approach is a perfect match to the collaborative nature of the experiment communities at CERN. New members of the LHC experiments can easily access information and do not need to follow a course in order to contribute.

6. ACKNOWLEDGMENTS

Our thanks to local user contacts handing distributed support of the application, notably for the Atlas and CMS collaborations. We would also like to thank our IT colleagues for the services we rely on.

7. REFERENCES

- [1] CERN. The European Organization for Nuclear Research. DOI=<http://public.web.cern.ch/public/>
- [2] The Large Hardron Collider. DOI=<http://public.web.cern.ch/public/en/LHC/LHC-en.html>.
- [3] TWiki. An open source enterprise Wiki. DOI=<http://www.twiki.org>
- [4] The Atlas experiment. DOI=<https://twiki.cern.ch/twiki/bin/view/AtlasPublic>.
- [5] The CMS experiment. DOI=<https://twiki.cern.ch/twiki/bin/view/CMSPublic>.