

TWiki at CERN

Past , Present and Future

Peter Jones IT-PES

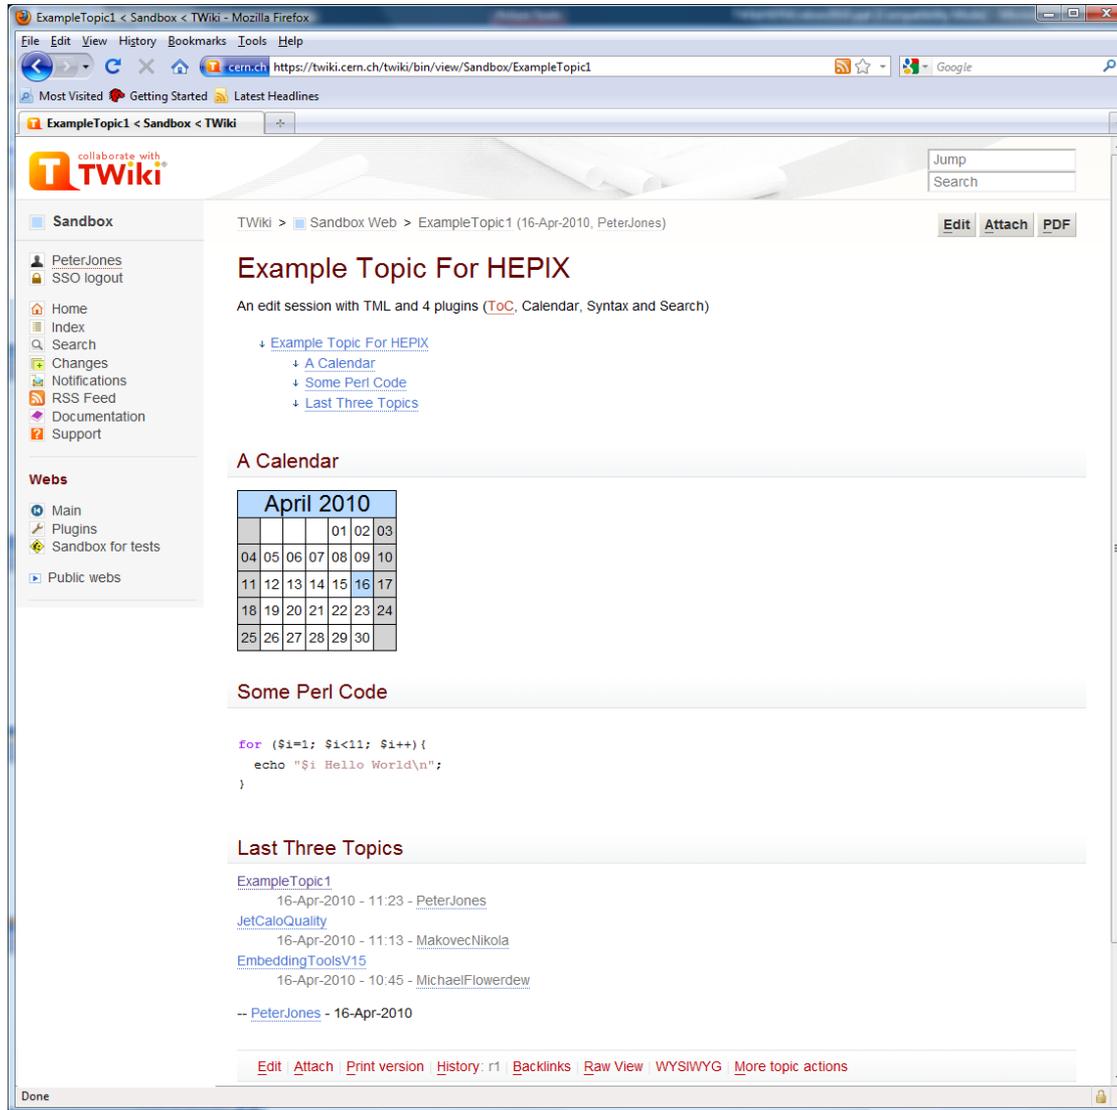
April 2010

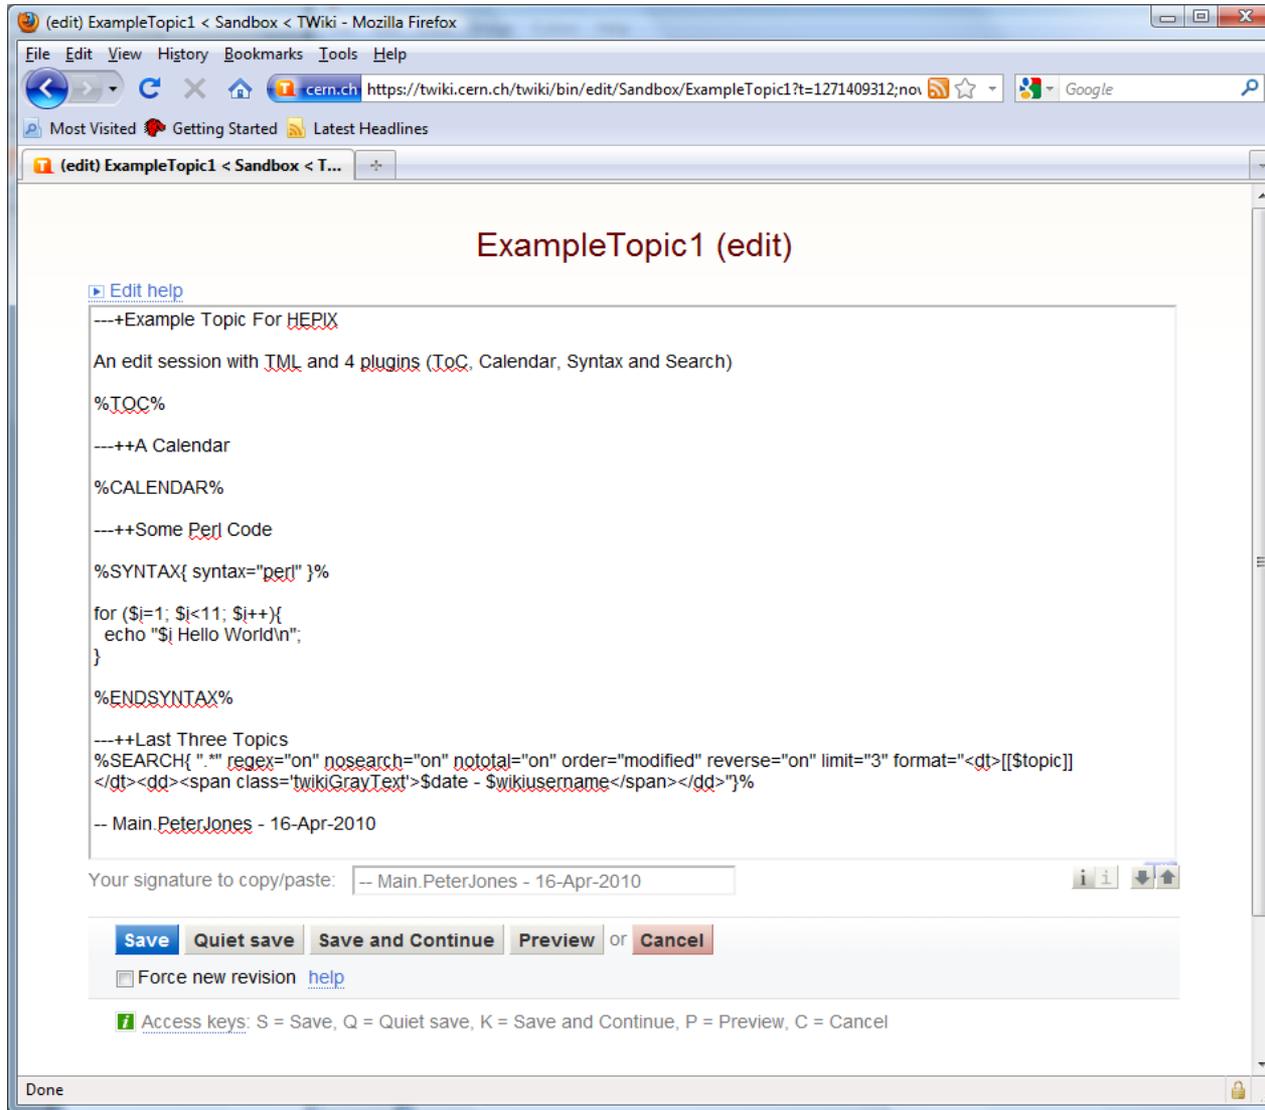
- Introduction
 - What is a wiki?
 - Why do we have TWiki?
- Overview and issues of the current service
 - Configuration
 - Performance
 - Authentication
 - Access Control
 - Authentication and Access Control Issues
 - Searching
- Future developments
- Questions

- In 2003 a group of CERN software developers requested a wiki.
- A wiki is a web page with an edit button.
- First wiki software was created by Ward Cunningham in 1995. "the simplest online database that could possibly work."
- There are many wikis available (see www.wikmatrix.org)
 - Confluence, MediaWiki, PhpWiki, TikiWiki, TWiki....
- TWiki was the chosen solution:
 - it matched the requirements and the CERN working environment (Perl, Linux, AFS and apache...)
 - CERN users from HIP were using TWiki at the time.

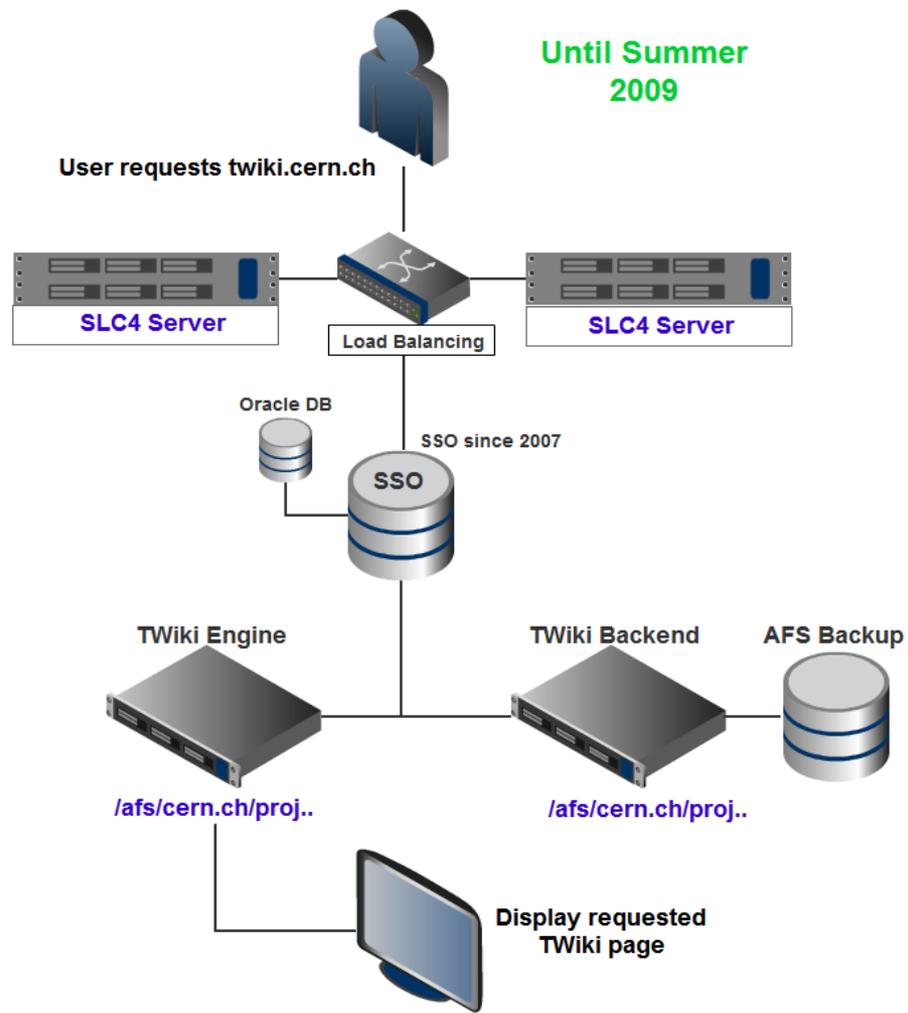
- TWiki – is a structured wiki
 - Introduced in 1998. OpenSource license (CERN input).
 - Used by: Nokia, Sony, Oracle, CISCO, EVO....
 - Perl based, plain text DB, runs on Linux and apache
 - Web and Topic hierarchy with access control
 - Revision control – all changes are tracked.
 - Extendability – over 200 plugins
 - Database like manipulation
 - Features of RSS/Atom feeds and e-mail notifications
 - Edit with HTML, TML or What You See Is What You Get (WYSIWYG) editor
 - Fast and easy to use

- TWiki has many uses:
 - Replaced many static web pages
 - Documentation, publication (PDF WorkBooks)
 - Whiteboard , Blog
 - Organiser – spreadsheet, calendars...
 - Knowledge base or FAQ
 - Software analysis and design
 - Software archive
 - Helps groups to collaborate more effectively

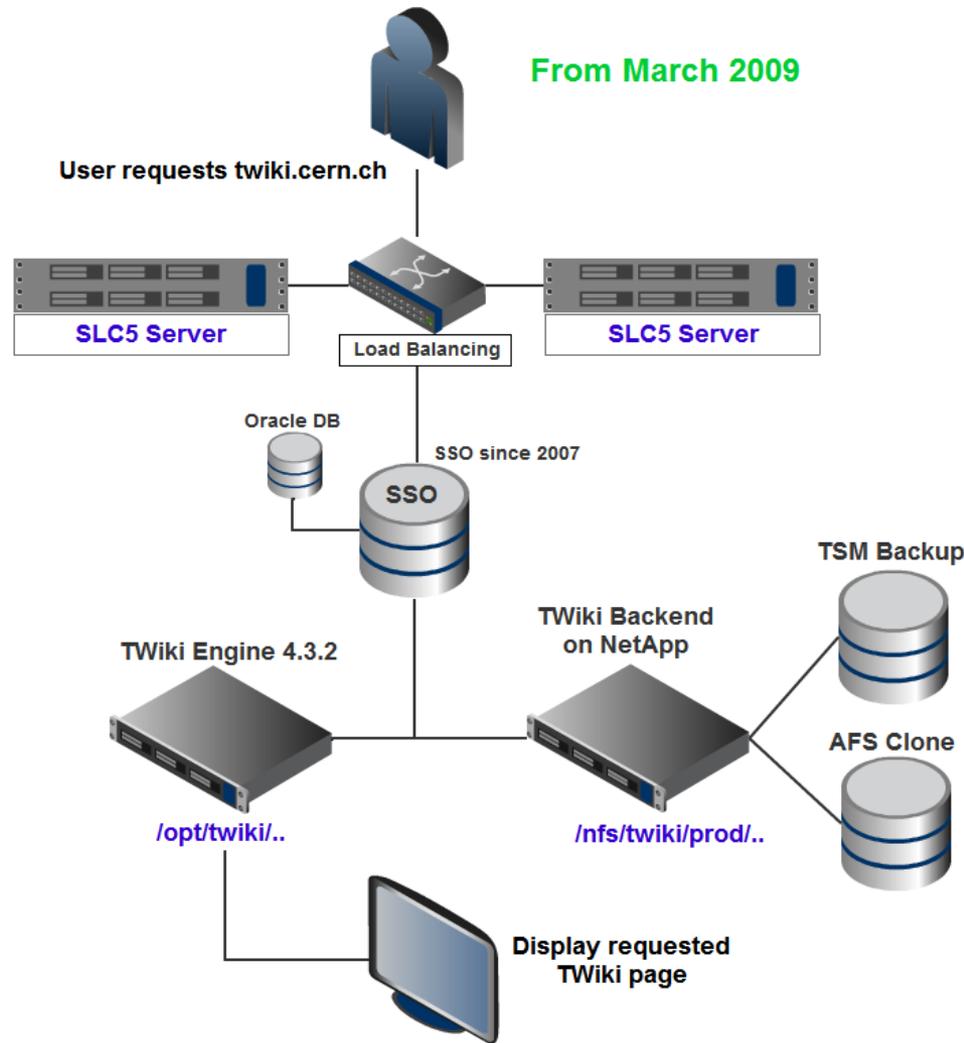




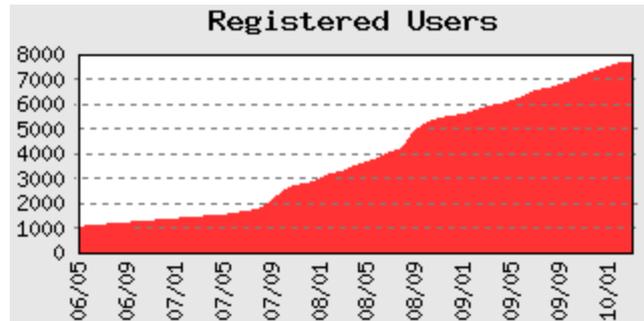
Configuration until 2009



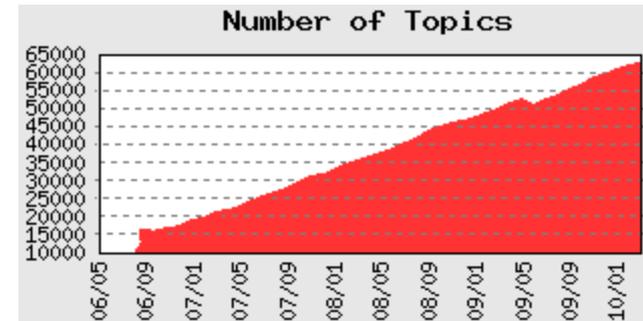
- Since 2003 TWiki has been upgraded several times
 - TWiki software Cairo -> Dakar -> Edinburgh -> Georgetown
 - To dedicated Linux servers (several upgrades)
- August 2009 installation of Georgetown (4.3.2)
 - Usability enhancements - new What You See Is What You Get editor
 - Security enhancements (reduced risk of cross site scripting)
 - SQL style query language for programmable search
 - New restore feature allows for easier rollback + over 300 bug fixes
- December 2009 upgrade of server
 - TWiki installed locally - DELL Poweredge 1950 2.33GHz, 2 CPU/8 core)
- March 2010 backend migrated from AFS->NFS
 - Service fine on AFS, however AFS has a file per directory limitation
 - Removed problem of 32K file limit (CMS web approaching this limit)
 - Removed all dependencies on AFS



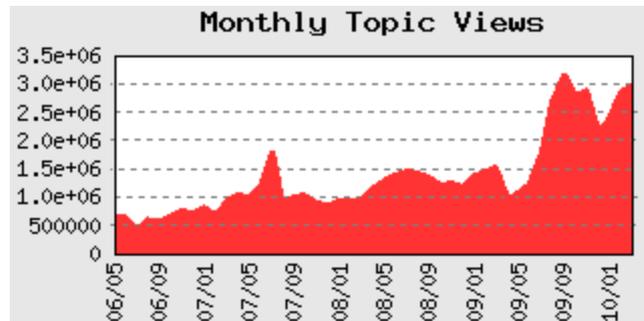
- Present Statistics:
 - 7500 registered users (does not include read-only users - TWikiGuest)
 - 190 collaboration webs
 - 60,000 topics (280,000 attachments)
 - 3,000,000 accesses/month (Mar 2010)
 - March: 50000 unique IP addresses, 10000 at CERN
 - 50,000 monthly updates
- Main users:
 - CMS, ~ 15,000 topics, 85000 attachments
 - Atlas, ~ 11000 topics, 60000 attachments
 - Others include Alice, LHCb, IT groups, Grid....



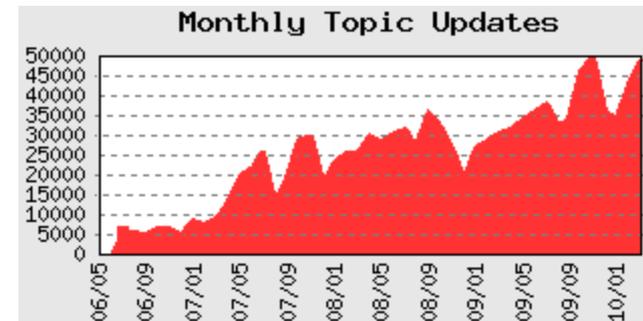
7563 users – March 2010



61743 topics – March 2010

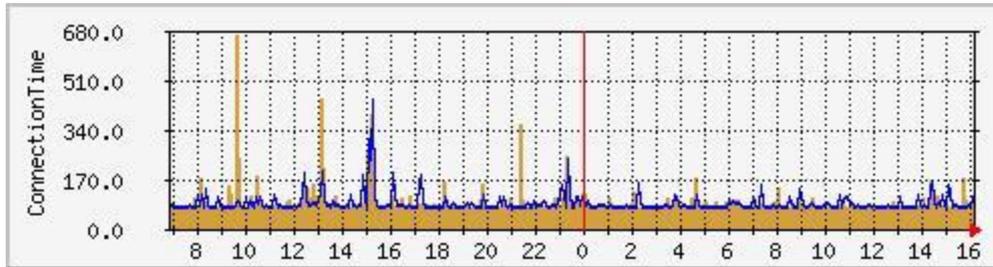


2971857 accesses – March 2010



48736 updated – March 2010

- It currently takes ~ 0.7 sec to load a page.



- Issues with slow performance addressed
 - Upgraded servers (DELL Poweredge 1950 2.33GHz, 2 CPU/8 core)
 - Upgraded Software (including Plugin updates)
 - Looked at any AFS bottlenecks (afs-cache, SSD)
 - Tuning: prepare some pages in advance (menus, feeds)
 - TWiki gardening is encouraged

- SSO since October 2007 – (initially we used Kerberos)
- TWiki Registration and membership of ‘CERN Users’ group necessary:
 - for editing files
 - uploading attachments
 - authorized views
 - topic management
- All registered users have a WikiName (for signature and access control)
- TWikiName is matched to UPN email address (from SSO)
 - PeterJones = peter.l.jones@cern.ch
- All users are called TWikiGuest until logged in
- Light weight registered users are allowed in but are vetted

- Set using the TWikiName(s)
- Defined by the Admin for the site level
- Set on the web level by the coordinator
 - Ex. Set DENYWEBVIEW = list of Users and/or Groups
- Set by the user for each topic
 - Ex. Set ALLOWTOPICVIEW = list of Users and/or Groups
- TWiki groups can also be used:
 - A list of registered TWiki users
 - Created by the TWiki user
 - Also created by a script based on mailing lists or egroups
- Since Nov 2009 direct use of egroups
 - Ex. Set ALLOWTOPICVIEW = cms-web-access

- TWiki is written in OO Perl and is extendable.
- A new user mapping can inherit existing modules

```
#!/usr/bin/perl
package TWiki::Users::ADFSUserMapping;
use base 'TWiki::Users::TWikiUserMapping';
```
- Collaboration requested Access Control based on egroups
- Apache returns information about SSO session
 - \$ENV{HTTP_ADFS_GROUP}
- Technical Student wrote new code to handle egroups
- New access control handles TWiki and egroup entries
 - Ex - Set ALLOWTOPICVIEW = PeterJones, hepix-users-egroup
- In production since November 2009 as requested by LHC
- Work still ongoing...
- <https://twiki.cern.ch/twiki/pub/Main/EgroupTWikiIntegration/EgroupsAndTWiki.pdf>

- Trouble-shooting for performance
 - High loads due to erratic TWiki use (feeds, search, browser issues, recursive calls...)
 - Some groups request static exports of data in case of downtimes.
- SSO has advantages (one login) but there are some disadvantages
 - Modification to the code and user registration
 - User mapping uses email address, which can change
 - Some users can have more than 1 SSO account
 - Load balancing has been problematic and is currently disabled
- Web management
 - Change of policy with some web or topic protection
 - Large TWikiGroups need maintaining
 - Web name changes. Causes problems with backlinks
- Search
 - TWiki WebSearch is slow on large webs and the results are not always satisfactory
 - People want Google search but this does not work on protected data
 - CERN search is the answer and will soon be ready on protected webs.

- Re-introduce loadbalancing
- Upgrade to new servers
 - 4 Dell M610 blades – 48Gb ram
 - At least one server to be on uninterruptable power
- Complete SSO integration
 - Complete egroup integration (for sub groups and PDF)
 - Solve user mapping issues (changes of email address)
- Automize TWiki installation (Quattor) – (bare SLC5 not enough)
- Install and test Helsinki version (v5) and upgrade extensions
- Gardening – work with collaborations to maintain their pages
- Improve search feature
- Keep an eye on developments and evaluate other wiki options

- Oct -27 2008 TWiki.net took over governance of the TWiki project. (LHC started Oct 2008)
- Core contributors disagreed with the new governance model.
- TWiki is now led by TWiki.net instead of the Open Source community. However it is still Open Source
- Foswiki available since early 2009.
 - Based on TWiki core.
 - Very different look and feel to TWiki by default.
 - Latest release is 1.09
- Stability is most important for our user community so we remain with TWiki for now.

- Total effort ~ 1FTE
 - IT-PES group
 - Including: Peter Jones, Nils Høimyr, Alexander Bernegger, Juan Manuel Guijarro
 - IT-DB group
 - Including: Giacomo Tenaglia, Artur Wiecek

(The above were all in the same group until end 2009)

- TWiki installed at CERN following user requests
- TWiki quickly grew in popularity
- Complements other IT web services
- The service has adapted to user requirements
- Regular software and hardware upgrade needed to satisfy user demand
- Performance –maintain quick page load
- Continue with SSO integration
- Improve search
- Continue to be part of the TWiki.org project
- Wikis evolve and other options are monitored

- CERN
 - <https://twiki.cern.ch/>
- TWiki
 - <http://www.twiki.org/>
 - <http://www.twiki.net/>
- Other
 - <http://www.wikimatrix.org/>
 - <http://www.foswiki.org/>

Thanks For Listening....

Peter Jones IT-PES

Peter.L.Jones@cern.ch