Subversion server

- can also be used for DSCM later on
- Server: zone on Sun Fire X4150,
> 8 cores (3.16GHz), 300 GB storage (ZFS).
- Second zone serves as backup and failover system
- URL: svn.services.openoffice.org
- Will host also OpenGrok index for OOo

Subversion

- ## Subversion 1.5
  - > merge tracking!

- ## Tentative repository structure:
  - > …/trunk
    …/tags   DEV300_m18, OpenOffice_3_0_0 …
    …/branches   DEV300, DEV301 …
    …/cws   swbugfix03, sb87, ….
    …/dist/…   anyone can get space here for patches etc
    …/contrib   is this needed?
    …/patches   globally useful patches

- ## path based access needed? Hopefully not!

Access

- access methods are
  - > svn+ssh: read/write access with ssh keys
  - > svn: read-only
  - > http: read-only access via web server

Modified CWS-tooling

- cws create [-m milest.] <mws> <cws>
  - > creates branch and registers with EIS
- cws fetch [-c CWS] [-m milest] [-s]
  - > checks out sources in a configurable way
  - > makes use of pre-checked out source tree
  - > modules are really only top level directories now
  - > includes cwsclone functionality
- cws rebase [-C | -m milestone]
  - > thin wrapper around subversion

Modified CWS-tooling (continued)

- cws analyze
- cws integrate
- cws query
- cws task
- cws eisclone
- cws ...

What should be migrated?

- Full migration, including all CWS branches
  - > not really possible, migration will take forever
- Full migr. minus "finished" CWS branches
- Trunk migration only.
- No history migration at all.
  - > probably to radical
- http://wiki.services.openoffice.org/wiki/Scm_migration_scope

Full mig. minus "finished" CWSs

- Pro
  - > Almost everything of relevance there
  - > Historic releases states can be extracted
  - > Current work checked out from new SCM
  - > Can be done, this has been tried out
  - > CVS server can be switched off (or r/o)
- Con
  - > The resulting repository will still be huge
  - > Number of revisions estimated > 500000
  - > The conversion will take about a week
  - > Quiet period for everyone
  - > Accuracy? Legacy workspaces need to be converted

Trunk only (that's what we did)

- Pro
  - > Still almost everything of relevance there
  - > Conversion is easy, has been tried out
  - > Repository is lean and mean
  - > Accuracy is no concern, history is for reference only
  - > No quiet period except for integrations
  - > No hassle with workspace migration
- Con
  - > CVS server remains online
  - > One time patch effort when integrating CWS
  - > Real ancient history might browsing CVS
  - > Historic state constructed from CVS