Fig. 1

Fig. 2

1. Operations at Heidelberg

(cf. Fig. 1)

Markus Demleitner
msdemlei@ari.uni-heidelberg.de

(cf. Fig. 2)

- What are we?
- Publication workflow
- Keeping the stuff running

2. What are we?

GAVO's data center:

- Established 2007 to support VO publication, primarily for Germany
- Project funding from federal bodies
- Catalogs, images, spectra, services, etc: We (have to) take what we get
- Simultaneously: home of DaCHS

3. Almost everything in an RD

DaCHS has resource descriptors (RDs). These specify:

- resource metadata (what? whence? by whom?)
- table metadata (including their columns)
- import rules
- service definitions, registration
- regression tests
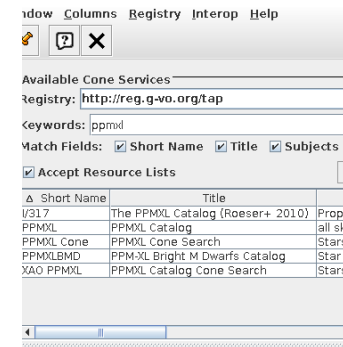
At Heidelberg: Publication $\approx$ write an RD.

Fig. 3

Fig. 4

4. Publication workflow

1. Provider upload (mail, rsync, http)
2. Subset on developer machine, RD in SVN
3. SVN checkout on server, non-linked service
4. Provider clarifications/feedback
5. Repeat until consensus
6. Publish to VO/web page

5. VO and web page

Here's a typical piece of RD:
```
<publish render="scs.xml" sets="ivo_managed"/>
<publish render="form" sets="local,ivo_managed"/>
```
The effect is a publication to a local portal page of a web interface („form"), and of both a cone search and the web-based thing to the registry.

It used to be that all the astronomers cared about was the web page.

These days, at least for catalogs, it's become more important to be "in TOPCAT".

For SSAP and SIAP, in my experience there's still some way to go until we're there.

(cf. Fig. 3)

(cf. Fig. 4)

6. Keeping the Stuff Running

Most important: **Test**.

- Before every commit to the DaCHS source: Run regression tests on almost all services (currently about 300 tests) and validate their RDs.
- After an svn update on the operational system (every couple of weeks): Run regression tests on the production service (including some extra ones requiring "big data").
- Monthly checks of the USVO validator page
- Monthly in-depth regression tests of the TAP service

Of course, most of this is automated and has very terse output on success.

7. Finally –

Main missing thing: DB failover/hot standby (DB restore for us: about 1 week). Tips welcome.

Otherwise: Nagios. It does help.

"Our part" of the publication process in version control is great for us – and the DaCHS users. In case you want to have a look: most of http://svn.ari.uni-heidelberg.de/svn/gavo/hdinputs/ is public. There's also a cross reference RD elements at http://docs.g-vo.org/DaCHS/elemref.html.