

GAVO DaCHS operator's guide

Author: Markus Demleitner
Email: gavo@ari.uni-heidelberg.de
Date: 2019-05-28

Contents

Starting and stopping the server	3
Publication	4
Registry Matters	5
Defining Basic Metadata	5
Creating an Organisation Record	7
Registering DaCHS-external Services	8
Registering Web Interfaces (And More)	9
Simple OAI operation	10
Making the VO see your Registry	11
Ensure There's a Backup of your Registry State	11
Adapting DaCHS for Your Site	12
Customisation Hooks	13
Operator CSS	13
XSL configuration	13
Userconfig RD	14

Simple Web Resources	15
Templates	15
Overridden System RDs	15
Other documents	16
The Vanity Map	16
Enabling HTTPS	17
Letsencrypt	17
HTTP preferred	19
Configuration Settings	19
Walkthrough	20
The general Section	20
The web Section	20
The db Section	21
The profiles Section	22
Reference	22
Section [general]	22
Section [adql]	24
Section [async]	24
Section [db]	25
Section [ivoa]	25
Magic Section [profiles]	26
Section [ui]	26
Section [web]	26
Managing Runtime Resources	28

Admin Interfaces	29
Admin Web Interfaces	29
robots.txt	29
Admin CLI Interfaces	30
Housekeeping	30
Two-Server Operation	31
Upgrading	32
Upgrading DaCHS	32
Upgrading Installations from Debian Package	32
Upgrading Installations from SVN	33
Upgrading Postgres	33

This document details the configuration and operation of the GAVO DaCHS server. For information on installing the software, please refer to the [installation guide](#), to learn how to import data, see the [tutorial](#). For an overview of the available documentation, see [DaCHS documentation](#)

Starting and stopping the server

The `gavo serve` subcommand is used to control the server. `gavo serve start` starts the server, changes the user to what is specified in the `[web]` user config item if it has the privileges to do so (that's "gavo" by default; you will already have created that user if you followed the installation instructions) and detaches from the terminal.

Analogously, `gavo serve stop` stops the server. To reload some of the server configuration (e.g., the resource descriptors, the vanity map, and the `/etc/gavo.rc` and `~/.gavorc` files), run `gavo serve reload`. This does not reload database profiles, and not all configuration items are applied (e.g., changes to the bind address and port only take effect after a restart). If you remove a configuration item entirely, their built-in defaults do not get restored on reload either.

Finally, `gavo serve restart` restarts the server. The start, stop, reload, and restart operations generally should be run as root; you can run them as the server user (by default, gavo), too, as long as the server doesn't try to bind to a privileged port (lower than 1025).

All this can and should be packed into a startup script or the equivalent entity for the init system of your choice. Our Debian package provides both a [System V-style init script](#) and a [systemd unit](#). They would typically be installed to `/etc/init.d/dachs` or `/etc/systemd/system/dachs` (but this might, of course, be different if you're running non-Debian systems).

For development work or to see what is going on, you can run `gavo serve debug`; this does not detach and does not change users, and it also gives a lot more tracebacks in the logs.

Publication

To "publish" a resource – which means include it either on your site's home page or in what you report to the VO registry –, add a `publish` element to a `service` (for a normal CatalogService publication), `data`, or `table` (these are data publications, typically for the TAP service) elements. The `publish` element lets you specify the sets the resources shall be published to. Unless you have specific applications, only two sets are relevant: `ivo_managed` for publishing to the VO (see [Registry Matters](#)), and `local` for publishing to your data center's service roster. Other sets can be introduced and used for, e.g., specific sub-rosters.

In services, the `publish` element needs, in addition, a `render` attribute, giving a comma-separated list of renderers the publication is for. The various renderers are translated into capability element in the VO resource records. For example, a typical pattern could be:

```
<publish render="scs.xml" sets="ivo_managed"/>
<publish render="form" sets="local,ivo_managed"/>
```

This generates one capability each for the simple cone search and a browser-based interface; the browser-based interface is, in addition, listed in the local service roster.

When you publish tables (or collection of tables via a `data` element), the notion of renderers makes no sense. Instead, you would have to define services that serve that data, except that when you publish tables that have `adql="True"`, the local TAP service is automatically considered to be a service for that data. Otherwise (or in addition), add service references in `service` children.

So, to publish an ADQL-queriable table to the VO for querying via TAP, just write:

```
<publish/>
```

within the table element. A table containing, e.g., data that's queried in a SIAP service in a different RD, would require something like:

```
<publish service="other/rd#siapsvc/>
```

(use `service` elements if there are multiple services).

Once you have done this, run `dachs pub <rdid>`. This causes all publishable items in the RD to be published. It also unpublishes everything that was published through the RD before and is no longer published. If the `[web]serverURL` config item on the machine running `gavo pub` is pointing to the actual running server and `[web]adminpasswd` matches, the server will automatically be made aware of these changes. Otherwise, you need to prod the server as discussed in [Managing Runtime Resources](#).

Registry Matters

As explained in the [tutorial chapter on the registry](#), you must provide enough data to allow the VO to tell who you are before the VO registry can usefully include your registry records.

Once that's done, you can add the publishing registry included in DaCHS to the list of registries harvested by VO full registries.

This chapter tries to guide you through this process.

Defining Basic Metadata

The first step is to define your authority (i.e., something like `org.g-vo.dc`) in your config (`/etc/gavo.rc`), in the `[ivoa]authority` item. Please make sure that the authority is not already taken by someone else; you have probably fulfilled your due diligence if you've run an [authority query against the registry](#) and did not find a match. Using your DNS name is not a bad idea. Please don't repeat our (the GAVO DC's) mistake and invert the sequence of particles in your DNS name. Also, this is just the name, there is no `ivo://` or other decoration.

Then, add metadata about yourself in `$GAVO_ROOT/etc/defaultmeta.txt`; these provide some values that stand in in registry records you generate whenever their values are not overridden in actual records, which obviously is particularly pertinent to things like the publisher, which very typically will be you for all the resources. It is a file in the [meta stream format](#); basically it's lines of `<key>:
<value>`.

- publisher – A short, human-readable name for you
- publisherID – An IVOA id for yourself; leave this empty, or create an `ivo://<authority>/org` resource in your `userconfig.rd` (see [Creating an organisation record](#)).

- `contact.name` – A human-readable name for some entity people should write to. This is not necessarily different from `publisher`, but ideally people can write "Dear <contact.name>" in their mails.
- `contact.address` – A contact address for surface mail
- `contact.email` – An email address. It will be published on web pages, so there probably should be some kind of spam filter in front of it.
- `contact.telephone` – A telephone number people can call if things really look bad.
- `creator.name` – A name to use when you give no creator in your resource descriptors. Could be some error sentinel ("we forgot to give credit, please complain") or just `contact.name` if you produce resources yourself.
- `creator.logo` – A URL for a logo to use when none is given in the resource metadata. Use a small PNG here; if you've put in a `logo_medium.png` (see [Simple Web Resources](#)), just use `<your-DaCHS-url>/favicon.png`, which is a suitably scaled version of that provided by DaCHS.
- `site.description` – A description of your site (i.e., "data center")
Example: The GAVO data center provides V0 publication services to all interested parties on behalf of the German Astrophysical Virtual Observatory. (use backslashes at the end of the lines to break long lines).

Then, fill out the metadata for the system registry resources in your `userconfig.RD`. See [Userconfig RD](#) below if you are not sure what we are talking about.

The registry is configured in the `registry-interfacerecords` stream (which you can copy from `//userconfig` if it's not yet in your `etc/userconfig.rd`).

In authority, change in particular

- `creationDate` – A UTC datetime (with trailing Z); technically, it should be the date the resource record is created, but realistically, just use "now" at the time you're writing the `defaultmeta.txt`. Example: `2007-12-19T12:00:00Z`.
- `title` – A human-readable descriptor of what the authority corresponds to.
Example: `The Utopia Observatory Data Center`
- `description` – A sentence or two on what the authority you are using means. This could be the same as `site.description` if all you're claiming authority for is that; if you're claiming authority for your institute or organisation, this obviously should be different. Example: `The Data Center at the Observatory of Utopia manages lots of substantial data sets`

created by the bright scientists all over the world (use backslashes at the end of the lines to break long lines).

- shortName – a short (about 16 chars) identifier for your authority. Example: GAVO DC.
- referenceURL – A URL at which people can learn more about your data center. Example: `http://www.g-vo.org`.
- managingOrg – A human-readable name for the organisation running this authority. The default is what you gave as publisher, which should be ok most of the time. If you make a registry record for yourself, you can also give `managingOrg.ivo-id`.

After you've specified all that, you're ready to define your first resources, viz, your registry itself, the authority, and the organisation that's managing it. These are predefined using the data you just filled in in the `//services` RD. To publish them, you say:

```
dachs pub //services
```

Depending on other services you're running, you may want to publish the following additional build-in resources:

```
# If you're publishing spectra, images, and the like, in particular
# if you produce pubDIDs
dachs pub //products
# If you run a TAP service
dachs pub //tap
# If you run a TAP service and want to point people to the built-in
# web interface
dachs pub //adql
# If you have an obscure service with images and want to also
# offer the material via SIAv2
dachs pub //siap2
```

Creating an Organisation Record

If you want to fill out the `publisherID` meta item – and there's no strong reason to do so at this point, you have to create a registry record for you, i.e., the organisation that runs the data center and hence the registry.

To do that, put something like:

```

<resRec id="manager">
  <meta>
    resType: organization
    creationDate: 2007-12-19T12:00:00
    title: GAVO Heidelberg Data Center
    subject: Organization
    referenceURL: http://www.ari.uni-heidelberg.de
    identifier: ivo://org.gavo.dc/org
    sets: ivo_managed
  </meta>
  <meta name="description">
    Operating at the Astronomisches Rechen-Institut (ARI) (part of
    Centre for Astronomy of Heidelberg University) on behalf of
    the German VO organisation GAVO, the GAVO Heidelberg Data Center
    is a one-stop shop for astronomical data publication projects
    of (almost) any description. We are also active within the
    IVOA in standards development and Registry operations.
  </meta>
</resRec>

```

into your `userconfig.rd`'s `registry-interfacerecords` and then say `dachs pub //services`.

Of course, you'll have to change things to match your situation; in particular, make sure identifier simply is `ivo://<your-authority>/org` – and that is what you'll use as `publisherID`.

Registering DaCHS-external Services

The registry interface of DaCHS can be used to register entities external to DaCHS; actually, you're already doing this when you're claiming an authority.

To register a non-service "resource", you can fill out a `resRec` RD element. You could reserve an RD (say, `$GAVO_ROOT/inputs/ext.rd` to collect such external registrations, or you could put them alongside internal services into their respective RDs. You will then usually just use the `resRec`'s `id` attribute to determine the IVORN of resource record. It will then be `ivo://<your authority>/<rd id>/<id of resRec>`.

In all likelihood, however, you will want to register services. To do that, use a normal service definition with with a `nullCore`. You probably need to manually give an `accessURL`. The most common case is that of a service with a `WebBrowser` capability. These result from `external` or `static` renderers. Thus, the pattern here usually is:

```

<service id="myservice" allowed="external">
  <nullCore/>
  <meta>

```



```

    shortName: My external service
    description: This service does wonderful things, even though\
        it's not based on GAVO's DaCHS software.
</meta>
<publish render="external" sets="ivo_managed">
    <meta name="accessURL">http://wherever.else/svc</meta>
</publish>
</service>

```

Of course, you will normally need to add further metadata as discussed above. dachs pub should complain if there's metadata missing, though.

The "services" can be fairly funky, actually; here's how GAVO registers their ADQL reference card:

```

<service id="adqlref" allowed="external">
  <nullCore/>
  <meta>
    shortName: GAVO ADQL ref
    creationDate: 2012-11-05T14:24:00Z
    title: The GAVO ADQL reference card
    subject:Virtual Observatory
    subject:Standards
    subject:ADQL
    description: GAVO's ADQL reference card briefly gives an overview \
of the SQL dialect used in the VO. It is available as a PDF\
file and as Scribus source under the CC-BY license.
    referenceURL:http://www.g-vo.org/pmwiki/About/ADQLReference
  </meta>
  <publish render="external" sets="ivo_managed,local">
    <meta name="accessURL">http://docs.g-vo.org/adqlref/adqlref.pdf</meta>
  </publish>
</service>

```

It is likely that if you register external services, you'll want to manage authorities other than [ivoa]authority as used by DaCHS. If you do, just add authority record(s) as before in the registry-interfacerecords STREAM in your [userconfig RD](#). And do not forget to add lines like:

```

<meta name="managedAuthority">edu.euro-vo.org</meta>

```

within the <service id="registry" in the user config.

Registering Web Interfaces (And More)

A typical situation is that you have a standard service (SSA, SCS, SIAP, etc) and a form-based custom service on the same data. Since the form-based service

caters to humans, it can require quite different input parameters (and thus usually cores) and output tables, and so you'll usually have a different service on it.

If you want to publish both services to the VO, you could add `publish` elements with `sets="ivo_managed"` to both `service` elements – but that would yield two resource records (which you then should link via `relatedTo` metas). At least when the form interface doesn't add significant functionality, this would usually seem overkill – e.g., your service would show up twice in resource listings.

Therefore, it is typically preferable to add the web interface as a capability to the resource record of the standard service. To let you do that, the `publish` element takes an optional `service` attribute containing the id of a service that should be used to fill the capability's metadata.

Here's an example:

```
<service id="web" defaultRenderer="form">
  <meta name="title">Form-based service</meta>
  <!-- add this service to the local roster -->
  <publish render="form" sets="local"/>
  ...
</service>

<service id="ssa" allowed="form,ssap.xml">
  <publish render="ssap.xml" sets="ivo_managed"/>
  <!-- now make a WebBrowser capability on this service in the IVOA
  published resource record, based on the service with the id web -->
  <publish render="form" sets="ivo_managed" service="web"/>
  ...
</service>
```

To publish

Simple OAI operation

If you want to check what you have published, see the `/oai.xml` on your server, e.g., <http://localhost:8080/oai.xml>. This is a plain OAI-PMH interface with some style sheets (if you want to customize them, copy them to `rootDir/web/xsl/`). The default style sheets add a link to "All identifiers defined here". Follow it to a list of all records you currently publish.

The OAI endpoint can also be used to help you in debugging validity problems with your registry content. To XSD-validate your registry without bothering the RofR (see above), you can do the following:

```
curl <your oai.xml url>?verb=ListRecords&metadataPrefix=ivo_vor | \
xmlstarlet fo > toval.xml
gavo admin xsdValidate toval.xml
```

This may result in a few error messages; if you don't understand them, it's a good idea to just go to the respective line in toval.xml and give it a long, hard look.

Making the VO see your Registry

The VO registry is a distributed system. There still is some sort of root, the [Registry of Registries](http://rofr.ivoa.net/registry) or RofR. Once your system provides sufficient metadata, go to <http://rofr.ivoa.net/regvalidate/regvalidate.html> and enter your registry endpoint (i.e., your installation's root URL with /oai.xml appended).

GAVO DaCHS is lenient with missing metadata and will deliver invalid VOResource for records missing some. It is not unlikely that your registry will not validate on the first attempt. Reading the error messages should give you a hint what's wrong. You can also use the `gavo val` command on the RDs that generate invalid records to figure out what's wrong.

Once your registry passes the validation test, you can add it to the RofR, and the full registries will start to harvest your registry (after a while).

Ensure There's a Backup of your Registry State

From your file system, there's no way to figure out what you have published when. To ensure that your registry stays consistent even when your database disappears, you should make sure that either you have a safe backup of your entire database (and no, backing up the files doesn't count, as that probably won't let you restore the thing) or, simpler, just add:

```
[db]
dumpSystemTables: True
```

to your `/etc/gavo.rc`. With that, DaCHS will every midnight dump the contents of the tables in `//services` and `//users` (which contain everything permanently operator-changable that doesn't go through imp) to `$stateDir/systemtables.dump`. You can restore this file using `dachs dump load` in case of a disaster (in particular if you back it up off-site once more).

Note that this will currently contain cleartext passwords of your site users (the stuff you added through `dachs adm adduser`). This doesn't hurt much over what's the case anyway, as the passwords are unencrypted in the database. But it would still be bad if anyone used valuable passwords with DaCHS. If you foresee that, implement password hashing and salting for DaCHS' user table or find someone to do that before going ahead.

Adapting DaCHS for Your Site

As delivered, the web interface of DaCHS will make it seem you're running a copy of the GAVO data center, with some metadata defused such that you are not actually disturbing our operation if you accidentally activate your registry interface. You should thus first customize the items given in `etc/defaultmeta.txt` (as discussed in [Registry Matters](#)).

The next adaptations are done through the configuration (as discussed in [Configuration Settings](#), i.e., usually in `/etc/gavo.rc`). The most relevant item here is `[web]sitename`, which should contain a terse identifier for the site (like "GAVO Data Center"). It is shown in titles and top headlines in many places. If you plan to use DaCHS' embargo feature together with user authorisation, you must also set `[web]realm` to some characteristic string. You could use the site name here; some user agents use it to display a prompt like "Credentials for <realm>" or similar.

Unless you plan to publish a sizeable number of services, you may want to override the root page. Essentially, you can just dump an XHTML page to `web/templates/root.html`, and DaCHS will use it as its root page. However, you may still want to consult the guide on [HTML templating in DaCHS](#).

If you want, you can set `[web]favicon` to either a `webDir`-relative path or a full URL to a [favicon](#).

It is also advisable to configure `[general]maintainerAddress` to a mail address of a person who will read problem reports. DaCHS doesn't send many of those yet, but it's still valuable if the software can cry for help if necessary. Sending mail only works if the local machine can actually send mail. If there is no MTA on your machine yet, we recommend `nullmailer` as a lightweight and easy-to-configure `sendmail` stand-in. If you use something else, you may need to adapt `[general]sendmail`.

For the rest, you can customize almost everything by overriding built-in resources. There are five major entities that you can override:

- [customisation hooks](#)
- [userconfig RD](#)
- [Simple Web Resources](#)
- [Templates](#)
- [Overridden System RDs](#)

If you find you need to override anything but the logo, please talk to us first – we'd in general prefer to provide customisation hooks. Overridden distribution files are always a liability on upgrades.

Customisation Hooks

Operator CSS

To override css rules we distribute or add new rules, avoid changing gavo_dc.css as described in [Simple Web Resources](#), as that will be a liability when upgrading. Instead, drop a CSS file somewhere (recommended location: `$GAVO_ROOT/web/nv_static/user.css`) and add a configuration item in `[web]operatorCSS`. With the recommended location, this would work out to be:

```
[web]
operatorCSS: /static/user.css
```

in `/etc/gavo.rc`.

This can also be an external URL, but we recommend against that, as that would force a browser to open one external connection per web page delivered.

By far the most common complaint is that we are limiting the width of `p` and `li` elements to 40em. We believe that text lines longer than about 80 characters are hard to read and should be avoided. On pages with tables where users might actually want to run browsers filling the entire screen, this choice cannot be made through a sensible choice of the width of the user agent window on the user side but requires CSS intervention.

Having said that, if you really think you want window-filling text lines, just put:

```
p, li {
    max-width: none;
}
```

into your operator CSS.

XSL configuration

DaCHS employs client-side XSLT for some purposes -- for instance, to show OAI-PMH (registry) responses in web browsers, to allow perusing datalink results in the browser, and to allow web browsers some rudimentary interaction with UWS applications like TAP.

The default XSLT contains references to the GAVO data center; to change these (or something else), override the xsl config stylesheet, which is expected at `/static/xsl/dachs-xsl-config.xsl`. The recommended way to go about this is:

```
cd /var/gavo # or wherever your DaCHS root is
cd web/nv_static
mkdir -p xsl
cd xsl
gavo admin dumpDF web/xsl/dachs-xsl-config.xsl > dachs-xsl-config.xsl
```

Then edit `dachs-xsl-config.xsl`. Note that you have to restart the server once to make it notice the override.

Userconfig RD

Fairly new in DaCHS is an RD exclusively for configuration. This is a place in which you can put streams that fill certain hooks; we expect to move more configuration into `userconfig`.

DaCHS has a builtin RD `//userconfig` that is updated as you update DaCHS. It always contains fallbacks for everything that can be in `userconfig` used by the core code. To override something, pull the elements in questions in your own `userconfig` RD and edit it there.

Your own `userconfig` RD is expected in `$GAVO_DIR/etc/userconfig.rd`. If it's not there yet, there's nothing wrong with starting with the distributed one:

```
cd 'gavo config configDir'
gavo admin dumpDF //userconfig > userconfig.rd
```

Once it's already there, use `dumpDF //userconfig` and, say, `less` to pick out the templates for whatever elements you need to copy. Currently, `userconfig` is already used in configuring the registry interface, extending the built-in obscure schema, and providing SIAPv2 metadata, and its use is growing.

Changes to `userconfig.rd` are picked up by DaCHS but will usually not be visible in the RDs they end up in. This is because DaCHS does not track which RDs make use of `userconfig`, so these will typically need to be reloaded manually. For instance, if you changed TAP examples, you'd need to run:

```
gavo serve exp //tap
```

to make your change show up in the web interface. Although usually not necessary, you can reload `userconfig` itself using:

```
gavo serve exp %
```

Note that for `gavo serve exp` to work, you need `[web]adminpasswd` set in your `/etc/gavo.rc`.

Simple Web Resources

For items coming from `static` (e.g., images, css, javascript), this overriding works by dropping same-named files in `$GAVO_ROOT/web/nv_static`.

Thus, you should put a PNG of your logo into `$GAVO_ROOT/web/nv_static/img/logo_medium.png`; scale it to about 250 pixels width or so (it will typically be used at 100 pt in the CSS).

Other files you may want to override in this way include

- `css/gavo_dc.css` – the central CSS; you could use this for skinning. However, you are *strongly* advised not to replace this file since some of the css is almost necessary for the web interface to do useful things, and the upstream CSS gets changed fairly regularly. Instead, use [Operator CSS](#).
- `help.shtml` – the help file. Unfortunately, we blurb quite a lot about GAVO in there right now. We'll think of something more parametrisable, but meanwhile you may want to have your own version
- `js/gavo.js` – could be the place for additional javascript; but frankly, if you want custom javascript, write to us and we'll think of a sane mechanism.
- `xsl/oai.xsl`, `xsl/uws-joblist-to-html.xsl`, `xsl/uws-job-to-html.xsl`, and `vosi.xsl` – XSLT stylesheet files. If you override these to customize them, please let us know. We'd try to put out generic stylesheets that are customisable without having to muck around in stuff that's basically functionality.

Templates

There is now a document on [HTML templating in DaCHS](#)

Overridden System RDs

You can copy system RDs from `gavo/resources/inputs/__system__` in the distribution to `$GAVO_ROOT/inputs/__system__` (adapt if you have played tricks with `inputsDir`) and edit them there. Again, if you feel you need to do that, contact us first, maybe we can work something out; it's a liability for upgrades.

Other documents

The default help file and the default sidebar link to a privacy policy that you should put down in `$GAVO_ROOT/web/nv_static/doc/privpol.shtml`. The document must be well-formed XHTML. Also, files with an extension `shtml` will be interpreted as templates over the service `//services/root`, which means that you can use the usual render functions and data items; the same goes for `disclaimer.html` (referenced from the standard sidebar) and, if you offer SOAP services, `soaplocal.html`. See the respective pages in the GAVO DC (<http://dc.g-ov.org/static/doc/...>) for ideas as to what to include.

The Vanity Map

DaCHS' URL scheme leads to somewhat clunky URLs that, in particular, reflect the file system underneath. While this doesn't matter to the VO registry, it is possibly unwelcome when publishing URLs outside of the VO. To overcome it, you can define "vanity names", single path elements that are mapped to paths.

These mappings are read from the file `$GAVO_ROOT/etc/vanitynames.txt`. The file contains lines of the format:

```
<target> <key> [<option>]
```

Target is a path that must *not* include `nevowRoot` and must *not* start with a slash (unless you're going for very special effects).

Key normally is a single path element (i.e., a string without a slash). If this path element is found in the first segment, it is replaced with the segments in target.

`<option>` can only be `!redirect` or empty right now.

If it is `!redirect`, `<key>` may be a path fragment (as opposed to a single path element); leading and trailing slashes are ignored. If the *enire* query path matches this key, a redirect to this key is generated. This is intended to let you shut down services and introduce replacements. If the incoming URL contains a query, it will be appended to the replacement URL. Thus, even stored queries or forms can potentially work across such a redirect.

You can also (ab)use the `redirect` option to give vanity names, but since the target will show up in the browser address line, normal maps are highly preferred. The only time normal maps don't work for this is when the resource directory is identical to the vanity name (you'll get an endless loop then), so you should avoid that situation.

Empty lines and #-on-a-line-comments are allowed in the input.

As an example, here's the vanity map that DaCHS had builtin as of version 0.6:

```
__system__/products/p/get getproduct
__system__/services/registry/pubreg.xml oai.xml
__system__/services/overview/external odoc
__system__/dc_tables/show/tablenote tablenote
__system__/dc_tables/show/tableinfo tableinfo
__system__/services/overview/admin seffe
__system__/tap/run/tap tap
__system__/adql/query/form adql !redirect
```

Note again that <key> must be a *single* path element only.

Enabling HTTPS

DaCHS can natively speak HTTPS; you have to let it claim port 443 on all the network interfaces you bind it to, though.

If, on the other hand, you use an external HTTPS termination (usually a reverse proxy), your best shot is probably to configure the reverse proxy's HTTPS URL as serverURL. You cannot currently have parallel HTTP in such a configuration, which is unfortunate (see [HTTP preferred](#)). If you really need external HTTPS termination, please let us know and we'll try to think of something.

To make DaCHS (attempt to) bind to an HTTPS port, all you need to do is give it a secret key and a certificate (which really is a public key with a CA's signature); these need to come in PEM format concatenated a file called \$GAVO_DIR/hazmat/bundle.pem. Just so the whole crypto thing doesn't become a total carnival, you should have the hazmat directory owned by root and with permissions 700. DaCHS will only read it when starting up using `gavo serve` and before dropping privileges. That way, if there's a minor security issue in DaCHS, you at least don't necessarily lose your private key.

On the other hand, you need to become root to update the certificate in this setup, which probably again is a security hazard, in particular when it's automated. If you want to do better, you probably can. Whatever you do, just make sure the operational DaCHS server (i.e., the `gavo` user) can't read the private key.

Letsencrypt

To get a certificate that is (hopefully) widely accepted, we strongly recommend you use letsencrypt. DaCHS has built-in support to update these certificates in time.

Here's how to go about it:

```

# DaCHS internally calls a tool called acme-tiny; we much rather trust
# it than any crypto code we'd hack together
$ sudo apt install acme-tiny
# now create the directory with the key material; we'll do all of this
# as root; if you don't like this, at least make it a user distinct
# from gavo and gavoadmin
$ sudo bash
$ cd 'gavo config rootDir'
$ mkdir hazmat
$ chmod 700 hazmat
$ cd hazmat
# generate an account key; this is what identifies you against
# letsencrypt
$ openssl genrsa 4096 > account.key
# Generate your secret key; these are the holy bits if you believe
# in https.
$ openssl genrsa 4096 > server.key
# now tell dachs to have the certificate signed; warning: this will
# restart the server
$ dachs serve updateCertificate
# Done -- hit ^D to exit the root shell.

```

Letsencrypt certificates are only good for three months. It's a good idea to renew them every two months. And that won't fly if you don't automate it. Automating it means your server will restart without manual intervention. That's bad because it might be down for a significant time (namely if someone is doing a large download at that time). As said below: HTTPS is an operational liability.

Anyway, add the following to root's crontab (if you installed DaCHS from source, you'll probably have to add `/usr/local/bin/dachs` here):

```
21 4 1 1,3,5,7,9,11 * dachs serve updateCertificate
```

-- adapt the time (here, 4:21 local time) to avoid times when it's likely you'll have many users. Also, make sure you receive mails from cron from the machine that does this, because there's a lot that can go wrong here, and DaCHS will just write error messages to stderr. That's probably not a big additional liability, as you should teach the service mails also for DaCHS itself.

Finally, to register your https endpoints, add:

```
registerAlternative: True
```

to the `ivoa` section in your `/etc/gavo.rc`. To propagate the information, restart the server and then say `dachs pub -a` to make the registries re-harvest your site.

If your machine is reachable through different host names (or box in Heidelberg, for instance, is both `dc.g-vo.org` and `dc.zah.uni-heidelberg.de`), tell DaCHS by listing all the non-primary names (i.e., anything that's not in `serverURL`) in a comma-separated list in `[web] alternateHostnames`.

HTTP preferred

We strongly recommend to have an unencrypted HTTP endpoint, and have that as the primary interface even when you support HTTPS in all but very rare cases.

Here's the reasoning: HTTPS is largely ineffective against most privacy breaches: Almost all nation-states can issue certificates that almost all user systems will accept; companies or institutions wanting to eavesdrop on their employees can force http proxies and install their own CA certificates on their employee's machines; and of course most privacy violations are side effects of platforms executing loads of Javascript sailing into the browsers with valid certificates.

On the other hand, networks become a lot more brittle with HTTPS: On the server side, the certificates need to be managed and regularly refreshed. And that's even before operators employ secure key management (which would probably require extra hardware or at least the use of passphrases).

Worse, the client side needs to keep their CA bundles (that's essentially lists of private keys they trust) up to date. With browsers that often carry these along and are updated quite regularly on most boxes, that's marginally manageable. For the VO, clients as a rule are not browsers but tools like TOPCAT or Aladin that may use a CA bundle coming with the Java VM, which often is not as well maintained. And once people start using curl or pyVO, the operating system's CA bundle is used, which on many systems is a mess. The net effect is that a given service may appear to work in the browser and in a script, but not from TOPCAT. Oh my. I mention in passing that WebSAMP on https pages may very well be impossible, and it certainly doesn't work now (which is why DaCHS swallows the SAMP button when it knows it's delivering via HTTPS).

Given the miniscule benefits and the serious operational implications, you should provide HTTP endpoints and register your those. If HTTPS is available, DaCHS will tell clients via the Registry's mirrorURL feature. If clients think HTTPS is worth it, they in this way can learn about support for it and use that as they see fit.

Configuration Settings

Many aspects of the data center can be configured using INI-style configuration files. DaCHS tries to obtain them from a global location (`/etc/gavo.rc` or whatever is in the `GAVOSETTINGS` environment variable) and a user-specific file (`~/gavo.rc` or whatever is in the `GAVOCUSTOM` variable). The server should probably be configured in the global location exclusively, since otherwise it will behave differently depending on which user starts the server.

This section starts with a walkthrough through the more relevant settings, section by section; below, there is a reference of all supported configuration items.

Walkthrough

The general Section

This mainly sets paths. The most important is `rootDir`, a directory most other paths are relative to. This is the one you'll most likely want to change. If you, e.g., wanted to have a private DaCHS tree, you could put:

```
[general]
rootDir: /home/user/gavo
```

into the personal configuration file (which DaCHS searches in `~/.gavorc`) by default; this would then override the analogous specification in `/etc/gavorc`.

The other paths in this section are interpreted relative to `rootDir` unless they start with a slash.

You may want to set `tempDir` and `cacheDir` to a directory local to your machine if `rootDir` is mounted via a network. Also note that we do no synchronisation for writing to the log (and never will -- we will provide syslog based logging if necessary), so you may want to tweak `logDir` too to keep actions from separate users separate.

The web Section

You typically want to adapt several settings here. First `bindAddress` gives the IP address of the interface DaCHS will accept requests from. By default, that's localhost, meaning that your server will only talk to the machine it runs on. Once you want to serve other people, you will need to change this. For most systems, binding to all interfaces is what you want; keep `bindAddress` empty to accomplish that.

You may also want to change `serverPort`. That is the TCP port DaCHS listens to. The default, 8080, is what's commonly used in test setups. On machines dedicated to DaCHS, you would set it to 80, the standard HTTP port; this will of course fail if there's already another web server running.

DaCHS frequently needs to produce full URLs to itself. To do that, it uses `serverURL`. While we could potentially infer that from `bindAddress` and `serverPort`, today's web setups are frequently too complicated to make that work. So, adapt `serverURL`, too, to the base URL of your server, without any trailing slash. A complete setup for a public server would thus look like this:

```
[web]
bindAddress:
serverPort: 80
serverURL: http://mydc.myvo.org
```

Note that `serverURL` *must* include the port if it is not 80; for https, DaCHS does not support nonstandard URLs. If you actually kept the default and just put the machine on the net, your web section would need to include something like:

```
[web]
bindAddress:
serverURL: http://your.machine.example.org:8080
```

– the empty `bindAddress` is necessary so DaCHS doesn't just bind to the loopback address, the `serverURL` because DaCHS has no way of knowing the preferred name of the machine it's running under; it *could* add the port, which it knows, but doing that would, e.g., make the lives of people operating behind reverse proxies hard.

While you are at it, set `sitename` to a short string describing your server (this is currently only used in the registry interface).

You will probably also want to set `adminpasswd`. If set, you can log in on your server as user `gavoadmin` with this password. `Gavoadmin` basically may do everything (access protected resources, clear caches, etc). The password is given in clear text; doing some kind of encryption would only make sense if you were prepared to enter some kind of passphrase every time you start the server. As in other places, DaCHS assumes the machine it runs on is trusted.

The db Section

In the `db` section, some global properties of the database access layer are defined. Currently, the most relevant one is `profilePath`. This is a colon-separated list of `rootDir`-relative paths in which DaCHS looks for database profiles (expansion of home directories is supported). The first match in any of these directories wins. This is useful when you have a test setup and a production setup -- just say `include dsn` in the common profiles (by default in `configDir`) and have separate `dsn` files in the `~/gavo` directories of the accounts feeding the test and production databases.

You probably do not want to mess with any settings ending in `Roles`. These are for rather exotic setups where DaCHS needs to accommodate other software.

The profiles Section

The profile section maps profile names to file names. These file names are relative to any of the directories in `db.profilePath`. Usually, you should keep whatever `gavo init` has come up with and hence not change anything here.

The profiles contain a specification of the access to the database in (unfortunately yet another, but simple) language. Each line in such a profile is either a comment (starting with `#`), an assignment (with `=`) or an instruction (consisting of a command and arguments, separated by whitespace).

Keywords available for assignment are

- `host` -- the host the database resides on. Leave empty for a Unix socket connection.
- `port` -- the port the database listens on. Leave empty for default 5432.
- `database` -- the database your tables live in.
- `user` -- the user through which the db is accessed.
- `password` -- the password of user.

There's just one command available, viz.,

- `include` -- read assignments and instructions from the profile given in the argument

`gavo init` creates four profile files, `dsn`, `feed`, `trustedquery`, and `untrustedquery`. These are referred to in the default profiles section, and are basically required by the python code.

Reference

You can get an up-to-date version of this by running `gavo config`.

Section [general]

Paths and other general settings.

- `cacheDir`: path relative to `rootDir`; defaults to `'cache'` -- Path to the DC's persistent scratch space

- `configDir`: path relative to `rootDir`; defaults to `'etc'` -- Path to the DC's non-ini configuration (e.g., DB profiles)
- `defaultProfileName`: string; defaults to `"` -- Deprecated and ignored.
- `gavoGroup`: string; defaults to `'gavo'` -- Name of the unix group that administers the DC
- `group`: string; defaults to `'gavo'` -- Name of the group that may write into the log directory
- `inputsDir`: path relative to `rootDir`; defaults to `'inputs'` -- Path to the DC's data holdings
- `logDir`: path relative to `rootDir`; defaults to `'logs'` -- Path to the DC's logs (should be local)
- `logLevel`: value from the list `info, debug, warning, error`; defaults to `'info'` -- How much should be logged?
- `maintainerAddress`: string; defaults to `"` -- An e-mail address to send reports and warnings to; this could be the same as `contact.email`; in practice, it is shown in more technical circumstances, so it's advisable to have a narrower distribution here.
- `operator`: string; defaults to `"` -- Deprecated and ignored. Use `contact.email` in `defaultmeta.txt` instead.
- `platform`: string; defaults to `"` -- Platform string (can be empty if `inputsDir` is only accessed by identical machines)
- `rootDir`: string; defaults to `'/var/gavo'` -- Path to the root of the DC file (all other paths may be relative to this)
- `sendmail`: string; defaults to `'sendmail -t'` -- Command that reads a mail from `stdin`, taking the recipient address from the mail header, and transfers the mail (this is for sending mails to the administrator). This command is processed by a shell (generally running as the server user), so you can do tricks if necessary.
- `stateDir`: path relative to `rootDir`; defaults to `'state'` -- Path to the DC's state information (last imported,...)
- `tempDir`: path relative to `rootDir`; defaults to `'tmp'` -- Path to the DC's scratch space (should be local)
- `uwsWD`: path relative to `rootDir`; defaults to `'state/uwsjobs'` -- Directory to keep uws jobs in. This may need lots of space if your users do large queries

- webDir: path relative to rootDir; defaults to 'web' -- Path to the DC's web related data (docs, css, js, templates...)
- xsdclasspath: shell-type path; defaults to 'None' -- Classpath necessary to validate XSD using an xsdval java class. You want GAVO's VO schemata collection for this. Deprecated, we're now using libxml2 for validation.

Section [adql]

(ignored, only left for backward compatibility)

- webDefaultLimit: integer; defaults to '2000' -- (ignored, only present for backwards compatibility; use [async]defaultMAXREC instead.

Section [async]

Settings concerning TAP, UWS, and friends

- csvDialect: string; defaults to 'excel' -- CSV dialect as defined by the python csv module used when writing CSV files.
- defaultExecTime: integer; defaults to '3600' -- Default timeout for UWS jobs, in seconds
- defaultExecTimeSync: integer; defaults to '60' -- Default timeout for synchronous UWS jobs, in seconds
- defaultLifetime: integer; defaults to '172800' -- Default time to destruction for UWS jobs, in seconds
- defaultMAXREC: integer; defaults to '20000' -- Default match limit for ADQL queries via the UWS/TAP
- hardMAXREC: integer; defaults to '20000000' -- Hard match limit (i.e., users cannot raise MAXREC or TOP beyond that) for ADQL queries via the UWS/TAP
- maxSlowPollWait: integer; defaults to '300' -- Maximal time a UWS 1.1-WAIT request will delay the response. This should be smaller than what you have as timeout on outgoing connections.
- maxTAPRunning: integer; defaults to '2' -- Maximum number of TAP jobs running at a time
- maxUserUWSRunningDefault: integer; defaults to '2' -- Maximum number of user UWS jobs running at a time

Section [db]

Settings concerning database access.

- `adqlProfiles`: set of strings; defaults to 'untrustedquery' -- Name(s) of profiles that get access to tables opened for ADQL
- `defaultLimit`: integer; defaults to '100' -- Default match limit for DB queries
- `dumpSystemTables`: boolean; defaults to 'False' -- Dump the tables from `//users` and `//system` to `stateDir/system_tables` once a day?
- `indexWorkMem`: integer; defaults to '2000' -- Megabytes of memory to give to postgres while making indices. Set to roughly half your RAM when you have big tables.
- `interface`: string; defaults to 'psycpg2' -- Don't change
- `maintainers`: set of strings; defaults to 'admin' -- Name(s) of profiles that should have full access to gavo imp-created tables by default
- `managedExtensions`: list of strings; defaults to 'pg_sphere' -- Name(s) of postgres extensions gavo upgrade -e should watch
- `msgEncoding`: string; defaults to 'utf-8' -- Encoding of the messages coming from the database
- `profilePath`: shell-type path; defaults to '~/.gavo:\$configDir' -- Path for locating DB profiles
- `queryProfiles`: set of strings; defaults to 'trustedquery' -- Name(s) of profiles that should be able to read gavo imp-created tables by default

Section [ivoa]

The interface to the Greater VO.

- `authority`: string; defaults to 'x-unregistred' -- The authority id for this DC; this has *no* leading `ivo://`
- `dalDefaultLimit`: integer; defaults to '10000' -- Default match limit on SCS/SSAP/SIAP queries
- `dalHardLimit`: integer; defaults to '1000000' -- Hard match limit on SCS/SSAP/SIAP queries (be careful: due to the way these protocols work, the results cannot be streamed, and the results have to be kept in memory; 1e7 rows requiring 1k of memory each add up to 10 Gigs...)

- `oaipmhPageSize`: integer; defaults to '500' -- Default number of records per page in the OAI-PMH interface
- `registerAlternative`: boolean; defaults to 'False' -- Give access URLs for the alternative protocol (https when serverURL is http and vice versa) as a mirrorURL? If you're listening to HTTPS, this is probably a good idea.
- `sdmVersion`: value from the list 1, 2; defaults to '1' -- Obsolete (SDM version 2 is shelved). Don't use.
- `VOSITableDetail`: value from the list max, min; defaults to 'max' -- Default level of detail to return on the VOSI endpoint (change to min when you have more than 100 or tables).
- `votDefaultEncoding`: value from the list binary, td; defaults to 'binary' -- Default 'encoding' for VOTables in many places (like the DAL responses; this can be user-overridden using the `_TDENC` local HTTP parameter).

Magic Section [profiles]

Ignored and deprecated, only here for backward compatibility. The items in this section are all of type profile name. You can add keys as required.

Section [ui]

Settings concerning the local user interface

- `outputEncoding`: string; defaults to 'iso-8859-1' -- Encoding for system messages. This should match what your terminal emulator is set to

Section [web]

Settings related to serving content to the web.

- `adminpasswd`: string; defaults to "" -- Password for online administration, leave empty to disable
- `adsMirror`: string; defaults to '<http://ads.g-vo.org>' -- Root URL of ADS mirror to be used (without a trailing slash)
- `alternateHostnames`: list of strings; defaults to "" -- A comma-separated list of hostnames this server is also known under. Only set this if you're running https. With this, you can handle a situation where your data center can be reached as both `example.org` and `www.example.org`.

- `bindAddress`: string; defaults to '127.0.0.1' -- Interface to bind to
- `corsOriginPat`: string; defaults to "" -- A regular expression for URLs from which to authorise cross-origin requests. This is matched, i.e., the RE must account for the whole URL including the schema. Example: `https?://example.com/apps/.*`
- `enableTests`: boolean; defaults to 'False' -- Enable test pages (don't if you don't know why)
- `favicon`: path relative to `webDir`; defaults to 'None' -- Webdir-relative path to a favicon
- `graphicMimes`: list of strings; defaults to 'image/fits,image/jpeg,application/x-votable+xml;content=datalink' -- MIME types considered as graphics (for SIAP, mostly)
- `jsSource`: boolean; defaults to 'False' -- If True, Javascript will not be minified on delivery (this is for debugging)
- `logFormat`: value from the list default, combined; defaults to 'default' -- Log format to use. Default doesn't log IPs, user agents, or referrers and thus should be ok in terms of not processing personal data, which in turn means you probably don't have to declare anything in EU jurisdictions.
- `maxPreviewWidth`: integer; defaults to '300' -- Ignored, only present for backward compatibility
- `maxUploadSize`: integer; defaults to '20000000' -- Maximal size of file uploads in bytes.
- `nevowRoot`: path fragment; defaults to '/' -- Path fragment to the server's root for operation off the server's root; this must end with a slash (setting this will currently break essentially the entire web interface. If you must use it, contact the authors and we will fix things.)
- `operatorCSS`: string; defaults to "" -- URL of an operator-specific CSS. This is included as the last item and can therefore override rules in the distributed CSS.
- `preloadRDs`: list of strings; defaults to "" -- RD ids to preload at the server start (this is mainly for RDs that have execute children that should run regularly).
- `previewCache`: path relative to `webDir`; defaults to 'previewcache' -- Webdir-relative directory to store cached previews in
- `realm`: string; defaults to 'X-Unconfigured' -- Authentication realm to be used (currently, only one, server-wide, is supported)

- `serverFDLimit`: integer; defaults to '4000' -- A hard limit of the number of file handles DaCHS should try to set. This will only take effect if DaCHS is started as root. Otherwise, DaCHS will just adjust its soft limit to the external limit irrespective of this setting.
- `serverPort`: integer; defaults to '8080' -- Port to bind the http port of the server to; https, if present at all, is always on 443.
- `serverURL`: string; defaults to '<http://localhost:8080>' -- URL fragment used to qualify relative URLs where necessary. Note that this must contain the port the server is accessible under from the outside if that is not 80; nonstandard ports are not supported for https. If you offer both http and https, use your preferred protocol here; for robustness, we recommend preferring http.
- `sitename`: string; defaults to 'Unnamed data center' -- A short name for your site
- `sqlTimeout`: integer; defaults to '15' -- Default timeout for db queries via the web
- `templateDir`: path relative to `webDir`; defaults to 'templates' -- `webDir`-relative location of global nevow templates
- `user`: string; defaults to 'gavo' -- Run server as this user.
- `voplotCodeBase`: URL fragment relative to the server's root; defaults to 'None' -- Deprecated and ignored.
- `voplotUserman`: URL fragment relative to the server's root; defaults to 'Deprecated and ignored' -- URL to the documentation of VOPlot

Managing Runtime Resources

DaCHS caches quite a lot of information rather aggressively, which means that editing information on disk may not immediately influence the behaviour of the server. This is particularly true for the default meta (`etc/defaultmeta.txt`), the vanity name translations (`etc/vanitynames.txt`), and the database profiles. Most of this can be reloaded on `gavo serve reload`, but certain settings (like `serverPort` and `bindAddress`) only take effect on a restart.

The resource descriptors are special. The server should pick up edits on RDs automatically, with the following exceptions:

- Built-in `__system__` RDs are not controlled. The main reason here is that these may not actually have disk files behind them, depending on the installation they may come from an archive.

- Only the file the RD was loaded from is checked. This means that if you override a built-in `__system__` RD with your own version in `inputs/__system__`, DaCHS will not automatically pick that up
- If you access an RD with no corresponding file and create that file afterwards, that change will also not be picked up automatically.

`gavo serve reload` will reload even those RDs. To selectively invalidate RDs that fall under these categories when you don't want to reload or restart the server, use the administration panel for the RD through the webserver; see [Admin Web Interfaces](#)

Admin Interfaces

Admin Web Interfaces

Some operation on the data center can be done from its web interface. To use these features, you have to set the `[web] adminpasswd` configuration item. You can then use the "Log in" link in the side bar using `gavoadmin` as the user name.

If you are logged in as `gavoadmin`, you should see an "Admin me"-link in the side bar of services. The page behind that link lets you block all services on the respective RD – where blocking means all requests are rejected until the RD is reloaded – and reload the RD. This is the recommended way to notify DaCHS that an RD has changed and needs re-reading.

In the form, you can also set scheduled down times. This is for VOSI, an interface clients could use to figure out whether a service can reasonable be expected to work. Since there don't seem to be clients exploiting the VOSI endpoints for such purposes so far, you probably don't need to bother.

You can directly access the administration panel for an RD by accessing `/seffe<rdId>`, e.g , `seffe/__system__/services`.

There are several more or less introspective resources within DaCHS that do not need authentication.

Among those, there's `/browse`. That's a list of all RDs that have (ivo or local) published services or data in them. Links on the RDs lead to info pages on the RDs, in particular giving tables and services within the RD.

robots.txt

DaCHS answers to requests for `robots.txt` with a built-in resource that forbids to index URLs with `/seffe` and `/login`. You may want to keep other pages out

of indices. In particular, `/browse` will let robots find unpublished services. To exclude those, add a file `robots.txt` in your `webDir` (run `gavo config webDir` to find out where that is) and add lines like:

```
Disallow: /browse
```

The built-in rules will be prepended to whatever you specify in your user `robots.txt`. For more information on what you can put into `robots.txt`, see [Robot exclusion standard](#)

Admin CLI Interfaces

You can also perform various housekeeping operations using `gavo admin`. Try `gavo admin --help`. This includes user management (there's a bit on it in the tutorial), precomputing previews for images, and create registry records for deleted services that got lost.

An admin tool that comes in handy is `gavo admin tapabort`. Call it with a TAP job id and a helpful (!) text to abort a TAP job and set it to an error state giving a short explanation what happened and the helpful text. The idea is that when users run queries against large tables without using indices (or do other stupid things), you can send them messages in this way (and clear away their resource hogs at the same time).

Housekeeping

There are a few things you should do to keep your data center function well. Most importantly, see how your services do in the VO ecosystem. Some places operate validators that tell you if there is something wrong with your site. If you are logged in as `gavoadmin`, you can run some of these interactively from the service info pages. We recommend, however, to peruse your results on <http://heasarc.gsfc.nasa.gov/vo/validation/vresults.pl> (the interface lets you select just your services).

Also, if you run a TAP service, you should run something like:

```
stilts taplint tapurl=<root-url>/tap | grep "^E-"
```

– `stilts` you can get from [Mark Taylor's TOPCAT page](#); this will output errors with your service. A fairly typical reason for errors is that you changed metadata without running `gavo imp -m`. If the metadata is wrong for just a few tables, just `gavo imp -m` their respective RDs. An easy way to fix all such problems is to just run:

```
gavo admin updateTAPSchema
```

which goes over all RDs that have contributed metadata to TAP_SCHEMA and re-ingests that metadata.

Two-Server Operation

When you want to run the database server on a different machine than the DaCHS server, the setup becomes a bit more involved. First, you'll only install python-gavodachs (rather than gavodachs-server) on the application server.

Then, you need to configure the two servers to talk to each other. Details depend on your setup, and you may want to apply extra hardening over what's below, but the following should get you started; it assumes the machine DaCHS runs on is called appserver.local, the database server is dbserver.local

To enable remote queries on the database server, on that machine, do something like:

```
$ cd /etc/postgres/<version>/main
$ sudo vi postgresql.conf
# edit so that listen_addresses = '*'
# -- without that, the server only listens on the loopback address
$ sudo vi pg_hba config
# add a line like
# host      all             all             appserver.local/32      md5
# below the access line enabling access for 127.0.0.1
# you can be more restrictive (e.g., "gavo" instead of the first "all",
# but that's left as an exercise to the reader
$ sudo service postgresql restart
# the following assumes you have the same username on appserver and on
# dbserver, otherwise use the username on *appserver*.
$ sudo -u postgres createuser -sP 'id -nu'
$ createdb -Ttemplate0 --encoding=UTF-8 --locale=C gavo
```

After that, you have a database DaCHS can configure from your account on appserver, using the password you gave in the createuser step.

To do that, on appserver say:

```
$ dachs init -d "host=dbserver.local user=<yourname> password=<pwd> dbname=gavo"
```

After that, everything should work as in the one-machine case.

Hint: do something like:

```
export PGHOST=dbserver.local export PGUSER=<yourname>
(or gavoadmin, if you prefer)
```

in your appserver's `.bashrc` (or equivalent), and `psql gavo` will give you a database shell on dbserver. If you trust your appserver, you can even give your password in `.pgpass`. But don't complain if this bites you later...

Upgrading

Upgrading DaCHS

In general, we try to make upgrades painless, but with a system allowing people to play tricks with intestines like DaCHS guarantees are hard. Be sure to subscribe to [DaCHS-users](#). We'll announce new releases there, together with brief release notes pointing to possible spots of trouble. Ideally, you'll have a development system and regression tests in place that let you diagnose problems before going to production. Poke us for hints on good and easily-maintained setups.

Upgrading Installations from Debian Package

- (1) Make sure you have enabled the intended distribution (*release* or *beta*) in your `/etc/apt.sources` (or equivalent)
- (2) Make sure all RDs DaCHS sees are in order:

```
gavo val -c ALL
```

Warnings you can usually ignore, but try to understand messages you get. You can hide known-broken RDs from DaCHS by dropping a file named `DACHS_PRUNE` into their directories.

Broken RDs are very likely to break upgrades. Fix them or remove them.

- (3) Do the actual upgrade:

```
apt-get update
apt-get upgrade
```

- (4) Run:

```
gavo val -tc ALL
```

This will complain if any of our changes break your services. If that is true and the Changelog did not alert you to the issue, please complain immediately. We may still be able to fix things for other people.

This last command might complain about mismatches between RD and on-disk metadata; there are several reasons why that may happen, including dumb or clever things we've done in the software. In any case, you should fix the problem, usually by re-importing the respective table.

When we change dependencies, individual DaCHS packages may be held back. `apt-get` will tell you this much, and you'll get `DistributionNotFound` errors from `pkg_resources` when starting DaCHS.

If it's only our packages that are held back, it should be safe to fix the problem by just running `apt-get dist-upgrade` (which will pull in the new dependencies). We can't give a generally recommendable recipe for how to deal with holdbacks for other packages.

Upgrading Installations from SVN

The basic thing to remember: After every update, do, as a user with ingestion privileges, the restart sequence:

```
$ gavo val -c ALL
$ sudo /etc/init.d/dachs stop # (or whatever you use to stop the server)
$ gavo upgrade
$ sudo /etc/init.d/dachs start # (or whatever you use to start the server)
$ gavo test ALL
```

In principle, `gavo upgrade` can run while the server is active, and with most updates, users won't even see errors, but since you need to restart anyway, why bother. On possible failures of the `gavo val` command, see the text in the packaged upgrade case.

The steps to update depend on what you did to install out of the subversion checkout. If you initially said `setup.py develop` (which we recommend), all it takes to upgrade is:

```
$ cd <checkout dir>
$ svn update
<run the restart sequence given above>
```

If you instead initially said `setup.py install`, do:

```
$ cd <checkout dir>
$ svn update
$ sudo python setup.py install
<run the restart sequence given above>
```

Upgrading Postgres

There's a howto over at [howDoI](#)