

Common Problems with DaCHS

Author: Markus Demleitner
Email: gavo@ari.uni-heidelberg.de
Date: 2025-11-26
Copyright: Waived under [CC-0](https://creativecommons.org/licenses/by/4.0/)

Contents

Import fails with "Column xy missing" and very few keys	2
Server is Only Visible from the Local Host	3
Transaction Deadlocking	3
'prodtblAccref' not found in a mapping	4
I get "Column ssa_dstitle missing" when importing SSA tables	4
"unpack requires a string argument of length"	4
"resource directory '<whatever>' does not exist"	4
Only RDs from below inputsDir may be imported	5
Not reloading services RD on server since no admin password available	5
I'm getting "No output columns with these settings." instead of result rows	5
gavo imp dies with Permission denied: '/home/gavo/logs/dcErrors'	5
I'm using reGrammar to parse a file, but no splitting takes place	6
IntegrityError: duplicate key value violates unique constraint "products_pkey"	6
relation "dc.datalinkjobs" does not exist	7
(some column) may be null but has no explicit null value	7
Column rave.main.logg_k: Unit dex is not interoperable	7

Column tab.foo is not a regular ADQL identifier	8
Unhandled exception ProgrammingError while importing an obscure table	8
dachs init fails with "type spoint does not exist"	9
'function "spoint_in" already exists with same argument types'	9
psycopg2.ProgrammingError: type "smoc" does not exist	10
relation "ivoa._obscoresources" does not exist	10
duplicate key value violates unique constraint "tables_pkey"	10
Key 'dataprodukt_type' not found in a mapping	11
'dict' object does not support indexing	11
Invalid script type preIndex for resource elements	11
ignoreOn in a rowmaker doesn't seem to work	11
module 'gavo.utils' has no attribute 'pyparsingWhitechars'	12
DistributionNotFound	12
500 FileNotFoundError when trying to pull jquery on bulleye	12
DaCHS limits says "Table X does not have columns DaCHS knows how to get a coverage from."	13
"job hung in EXECUTING" when doing TAP async jobs	13

"ERROR: permission denied to alter role" during dachs upgrade 14

This document tries to discuss some error messages you might encounter while running DaCHS. The rough idea is that when you can grep in this file and get some deeper insight as to what happened and how to fix it.

We freely admit that some error messages DaCHS spits out are not too helpful. Therefore, you're welcome to complain to the authors whenever you don't understand something DaCHS said. Of course, we're grateful if you checked this file first.

Import fails with "Column xy missing" and very few keys

This error is caused by the row validation of the table ingestor – it wants to see values for all table columns, and it's missing one or more. This may be flabbergasting when your grammar yields the fields that are missing here. The reason is simple: You must map them in the rowmaker. If you see this error, you probably wanted to write `idmaps="*" or something like that in the rowmaker.`

Server is Only Visible from the Local Host

When the server is running (`gavo serve start`) and you can access pages from the machine the server runs on just fine, but no other machines can access the server, you run the server with the default web configuration. It tells the server to only bind to the loopback interface (127.0.0.1, a.k.a. localhost).

To fix this, say:

```
[web]
bindAddress:
```

in your `/etc/gavo.rc`.

Transaction Deadlocking

When `gavo imp` (or possibly requests to the server) just hangs without consuming CPU but not doing anything useful, it is quite likely that you managed to provoke a deadlock. This happens when you have a database transaction going on a table while trying to access it from the outside. While DaCHS tries to not leave connections in a state that's called "idle in transaction", bugs or user code may cause this.

To diagnose what's happening, it is useful to see the server's idea of what is going on inside itself. The following script (that you might call `psdb`) will help you:

```
#!/bin/sh
psql gavo << EOF
select pid, username, state, query, date_trunc('seconds', query_start::time)
from pg_stat_activity
order by pid
EOF
```

(this assumes your database is called `gavo` and you have sufficient rights on that database; it's not hard to figure out the `psql` command line for other scenarios). This could output something like:

procpid	username	current_query	date_trunc
9301	gavoadmin	<IDLE>	16:55:39
9302	gavoadmin	<IDLE> in transaction	16:55:39
9303	gavoadmin	<IDLE> in transaction	16:55:39
9306	gavoadmin	<IDLE> in transaction	16:55:43
9309	gavoadmin	SELECT calPar FROM l...	16:55:43

(5 Zeilen)

The `procpid` is the pid of the process handling the connection. As said above, it's the `<IDLE> in transaction` connections you need to watch out for. Simply killing their `procpid` (at the operating system level) will raise an exception in the code that ran the query. Of course, you will need to become the `postgres` or `root` user to do that.

To fix such a deadlock, you will in general have to commit the connection that went idle without a commit. If this badly breaks the atomicity logic, sometimes an alternative is to have the two things that deadlock share a connection.

'prodtblAccref' not found in a mapping

You get this error message when you make a table that mixes in `//products#table` (possibly indirectly, e.g., via SSAP or SIAP mixins) with a grammar that does not use the `//products#define` row filter.

So, within the grammar, say (at least, see the reference documentation for other parameters for `rowgen`):

```
<rowfilter procDef="//products#define">
  <bind name="table">"\schema.dest_table"</bind>
</rowfilter>
```

-- substituting the `dest_table` with the id of the table fed. The reason why you need to manually give the table is that the grammar doesn't know what table the rows generated end up in. On the other hand, the information needs to be added in the grammar, since it is fed both to your table and the system-wide products table.

I get "Column ssa_dstitle missing" when importing SSA tables

The `//ssap#setMeta` rowmaker application does not directly fill the output rowdict but rather defines new input symbols. This is done to give you a chance to map things set by it, but it means that you must `idmap` at least all `ssa` symbols (or map them manually, but that's probably too tedious). So, in the rowmaker definition, you write:

```
<rowmaker idmaps="ssa_*">
```

"unpack requires a string argument of length"

These typically come from a binary grammar parsing from a source with `armor=fortran`. Then, the input parser delivers data in parcels given by the input file, and the grammar tries to parse it into the fields given in your `binaryRecordDef`. The error message means that the two don't match.

This can be because the input file is damaged, you forgot to skip some header, but it can also be because you forgot fields or your `binaryRecordDef` doesn't match the input in some other way.

"resource directory '<whatever>' does not exist"

DaCHS expects each RD to have a "resource directory" that contains input files, auxiliary data, etc. Multiple RDs may share a single resource directory.

By default, the resource directory is `<inputsDir>/<schema>`. If you don't need any auxiliary files, the `resdir` does not even need to exist. In that case, you'll see the said warning. To suppress it, you could just say:

```
<resource schema="<whatever>" resdir="__system">
```

The `__system` resource directory is used by the built-in RDs and thus should in general exist.

However, the recommended layout is, below `inputsDir`, a subdirectory named like the resource schema, and the RD immediately within that subdirectory. In that case, you don't need a `resdir` attribute.

Only RDs from below inputsDir may be imported

RDs in DaCHS must reside below `inputsDir` (to figure out what that is on your installation, say `dachs config inputsDir`). The main reason for that restriction is that RDs have identifiers, and these are essentially the `inputsDir`-relative paths of the file. Out-of-tree RDs just cannot compute this. Therefore, most subcommands that accept file paths just refuse to work when the file in question is not below `inputsDir`.

Not reloading services RD on server since no admin password available

That's a warning you can get when you run `gavo pub`. The reason for it is that the DaCHS server caches quite a bit of information (e.g., the root page) that may depend on the table of published services. Therefore, `gavo pub` tries to make the running server discard such caches. To do this, it inspects the `serverURL` config item and tries access a protected resource. Thus, it needs the value of the config setting `adminpasswd` (if set), and that needs to be identical on the machine `gavo pub` executes on and on whatever `serverURL` points to.

If anything goes wrong, a warning is emitted. The publication has happened still, but you may need to run `gavo serve reload` on the server to make it visible.

I'm getting "No output columns with these settings." instead of result rows

This is particularly likely to happen with the `scs.xml` renderer. There, it can happen the the server doesn't even bother to run database queries but instead keeps coming back with an error message `No output columns with these settings..`

This happens when the "verbosity" (in SCS, this is computed as $10 \times \text{VERB}$) of the query is lower than the `verbLevel` of all the columns. By default, this `verbLevel` is 20. In order to ensure that a column is returned even with `VERB=1`, say:

```
<column name=... verbLevel="1"/>
```

gavo imp dies with Permission denied: '/home/gavo/logs/dcErrors'

(or something similar). The reason for these typically is that the user that runs `gavo imp` is not in the `gavo` group (actually, whatever `[general]gavoGroup` says). To fix it, add that user to the group. If that user was, say, `fred`, you'd say:

```
sudo adduser fred gavo
```

Note that `fred` will either have to log in and out (or similar) or say `newgrp gavo` after that.

To add yourself, type:

```
sudo adduser 'id -nu' gavo
```

I'm using reGrammar to parse a file, but no splitting takes place

This mostly happens for input lines like `a|b|c`; the underlying problem is that you're trying to split along regular expression metacharacters. The solution is to escape the the metacharacter. In the example, you wouldn't write:

```
<reGrammar fieldSep="|"> <!-- doesn't work -->
```

but rather:

```
<reGrammar fieldSep="\|"> <!-- does work -->
```

IntegrityError: duplicate key value violates unique constraint "products_pkey"

This happens when you try to import the same "product" twice. There are many possible reasons why this might happen, but the most common (of the non-obvious ones) probably is the use of updating data items with row triggers.

If you say something like:

```
<!-- doesn't work reliably -->
<table id="data" mixin="//products#table"
...
<data id="import" updating="True">
  <sources>
    ...
    <ignoreSources fromdb="select accref from my.data"/>
  </sources>
  <fitsProdGrammar...
  <make table="data">
    <rowmaker>
      <ignoreOn name="Skip plates not yet in plate cat">
        <keyMissing key="DATE_OBS"/></ignoreOn>
    ...
```

you're doing it wrong. The reason this yields IntegrityErrors is that if the ignoreOn trigger fires, the row will not be inserted into the table data. However, the make feeding the `dc.products` table implicitly inserted by the `//products#table` mixin will not skip an ignored image. So, it will end up in `dc.product`, but on the next import, that source will be tried again – it didn't end up in `my.data`, which is where `ignoreSources` takes its file names from –, and boom.

If you feed multiple tables in one data and you need to skip an input row entirely, the only way to do that reliably is to trigger in the grammar, like this:

```
<table id="data" mixin="//products#table"
...
<data id="import" updating="True">
  <sources>
    ...
    <ignoreSources fromdb="select accref from my.data"/>
  </sources>
  <fitsProdGrammar...
```

```

    <ignoreOn name="Skip plates not yet in plate cat">
      <keyMissing key="DATE_OBS"/></ignoreOn>
    </fitsProdGrammar>
    <make table="data">
    ...

```

relation "dc.datalinkjobs" does not exist

This happens when you try to run asynchronous datalink (the dlasync renderer) when you've not created the datalink jobs table. This is not (yet) done automatically on installation since right now we consider async datalink to be a bit of an exotic case. To fix this, run:

```
gavo imp //datalink
```

(some column) may be null but has no explicit null value

These are warnings emitted by the DaCHS' RD parser – since they are warnings, you could ignore them, but you shouldn't.

This is about columns that have no "natural" NULL serialisation in VOTables, mostly integers. Without such a natural NULL, making VOTables out of anything that comes out of these tables can fail under certain circumstances.

There are (at least) two ways to fix this, depending on what's actually going on:

- (a) you're sure there are no NULLs in this column. In that case, just add `required="True"`, and the warnings will go away. Note, however, that DaCHS will instruct the database to check that you're not cheating, and an import will fail if you try to put NULLs into such columns.
- (b) there are NULLs in this column. In that case, find a value that will work for NULL, i.e., one that is never used as an actual value. "Suspicious" values like 0, -1, -9999 or the like are preferred as this increases the chance that careless programs, formats, or users who ignore a NULL value specification have a chance to catch their error. Then declare that null value like this:

```

<column name="withNull" type="integer"...>
  <values nullLiteral="-9999"/>
</column>

```

Column `rave.main.logg_k`: Unit dex is not interoperable

The [VOUnit standard](#) lets you use essentially arbitrary strings as units – so does DaCHS. VOUUnit, however, contains a canon of units VO clients should understand. If DaCHS understands units, you can, for instance, change them on form input and output using the `displayUnit displayHint` – other programs may allow automatic conversion and similar comforts.

When DaCHS warns that a unit is not interoperable, this means your unit will not be understood in that way. There are cases when that's justified, so it's just a warning, but be sure you understand what you've written and there actually is no interoperable (i.e., using the canonical VOUUnits) way to express what you want to say.

Also note that it is an excellent idea to quote free (i.e., non-canonical) units, i.e., write `unit='Crab'`. The reason is that in the non-quoted case, VOUnit parsers will try to separate off SI prefixes, such that, e.g., `dex` will be interpreted as `dezi-ex`, i.e., a tenth of an `ex` (which happens to actually be a unit, incidentally, although not a canonical VOUnit one).

And yes, `dex` itself would be a free unit. If you look, quantities given with "dex" as a unit string actually are dimensionless. Our recommendation therefore is to have empty units for them.

Column `tab.foo` is not a regular ADQL identifier

This is a message you may see when running `gavo val`. It means that the column in question has a name that will get you in trouble as soon as you open the table in question to TAP queries (and trust me, you will sooner or later). Regular ADQL identifiers match the regular expression `[A-Za-z][A-Za-z0-9_]*` with the additional restriction that ADQL reserved words (including terms like `distance`, `size`, etc) are not allowed either.

If you see the message, just change the name in question. There's so many nice words that there's really no need to use funny characters in identifiers or hog ADQL reserved words.

If you *must* keep the name anyway, you can prefix it by `quoted/` to make it a delimited identifier. There's madness down that road, though, so don't complain to us if you do that and regret it too late. In particular, you may have a hard time referencing such columns from STC declarations, when creating indices, etc. So: Just don't do it.

Unhandled exception `ProgrammingError` while importing an obscure table

This typically looks somewhat like this:

```
ProgrammingError: syntax error at or near "/"
LINE 28:         CAST(/RR/V/ AS text) AS pol_states,
                ^
*** Error: Oops.  Unhandled exception ProgrammingError.

Exception payload: syntax error at or near "/" LINE 28:
CAST(/RR/V/ AS text) AS pol_states,
```

While `ProgrammingErrors` in general happen whenever an invalid query is sent to the database engine, when they pop up in `gavo imp` with `obscure` not far away it almost invariably means that there is a syntax error, most likely forgotten quotes, in the `obscure` mixin definitions of one of the tables published through `obscure`. The trick is to figure out which of them causes the trouble.

The most straightforward technique is to take the fragment shown in the error message and look in `ivoa._obscuresources` like this:

```
$ psql gavo
...
# select tablename
- from ivoa._obscuresources where sqlfragment like '%CAST(/RR/V/%';
```

```
      tablename
-----
test.pgs_siaptable
```

You could `gavo purge` the table in question to fix this the raw way, but it's of course much more elegant to just remove the offending piece from `_obscoresources`:

```
# delete from ivoa._obscoresources where tablename='test.pgs_siaptable';
```

Then fix the underlying problem – in this case that was replacing:

```
<mixin
polStates="/RR/V/"
...
```

with:

```
<mixin
polStates="'/RR/V/'"
...
```

– and re-import the obscure meta; you'll usually use `gavo imp -m && gavo imp //obscure` for that (see also [updating obscure](#))

Note that running `dachs imp //obscure recover` will fix trivial cases without requiring any thought.

dachs init fails with "type spoint does not exist"

This always means that the `pgsphere` extension could not be loaded, and DaCHS can no longer do without it. Actually, we could try to make it, but you really need `pgsphere` in almost all installations, so it's better to fix this than to work around it.

Unfortunately, there is any number of reasons for a missing `pgsphere`.

If, for instance, you see this error message and have installed DaCHS from tarball or git, with manual dependency management, just install the `pgsphere` postgres extension (and, while you're at it, get `q3c`, too); see DaCHS' [installation instructions](#) for details.

If this happens while installing the Debian package, in all likelihood DaCHS is not talking to the postgres version it thinks it is. This very typically happens if you already have an older postgres version on the box. Unless you're sure you know what you're doing, just perform an upgrade to the version DaCHS wants – see [howDol.html#upgrade-the-database-engine](#). If you'd need to downgrade, that's trouble. Complain to the `dachs-support` mailing list – essentially, someone will have to build a `pgsphere` package for your postgres version.

'function "spoint_in" already exists with same argument types'

This would typically show in `dachs init`, and it is again an indication that the state of the `pgsphere` extension is bad. The most likely reason is that postgres has not been acquainted with “new-style” extension metadata of `pgsphere`. In that case, `pg_sphere` will be absent of the output of `select * from pg_extension`.

If that is true, running:

```
psql gavo -c "CREATE EXTENSION pg_sphere FROM UNPACKAGED"
```

should fix the problem (followed by another `dachs init`; if the failure happened during a DaCHS upgrade, run `apt install -f` instead).

psycopg2.ProgrammingError: type "smoc" does not exist

This means that you are using a `pgsphere` (the postgres extension that does the spherical geometry within the database) does not support MOCs. This is true for the `postgresql-pgsphere` type that comes with Debian buster.

To fix this, either don't use `smocs` – or install the `postgresql-11-pgsphere` package coming from GAVO's repository).

relation "ivoa._obscoresources" does not exist

This happens when you try to import an obscure-published table (these mix in something like `//obscore#publish-whatever`) without having created the obscure table itself. The fix is easy: Either remove the mixin if you don't want the obscure publication (which would be odd for production data) or, more typically, create the obscure table:

```
dachs imp //obscore
```

duplicate key value violates unique constraint "tables_pkey"

This typically happens on `dachs imp`. The immediate reason is that `dachs imp` tries to insert a metadata row for one of the tables it just created into the `dc.tablemeta` system table, but a row for that name is already present. For instance, if you're importing into `arihip.main`, DaCHS would like to note that the new table's definition can be found at `arihip/q#main`. Now, if `dc.tablemeta` already says `arihip.main` was defined in `quicktest/q#main`, there's a problem that DaCHS can't resolve by itself.

90% of the time, the underlying reason is that you renamed an RD (or a resource directory). Since the identifier of an RD (the RD id) is just its relative path of the RD to the inputs directory (minus the `.rd`), and the RD id is used in many places in DaCHS, you have to be careful when you do that (essentially: `dachs drop --all old/rd; mv old new; dachs imp new/rd`).

If you're seeing the above message, it's already too late for that careful way. The simple way to repair things nevertheless is to look for the table name (it should be given in the DETAILS of the error message) and simply tell DaCHS to forget all about that table:

```
dachs purge arigh.main
```

This might leave other traces of the renamed RD in the system, which might lead to trouble later. If you want to be a bit more thorough, figure out the RD id of the vanished RD by running `psql gavo` and asking something like:

```
select sourcerd from dc.tablemeta where tablename='arihip.main'
```

This will give you the RD id of the RD that magically vanished, and you can then say:

```
dachs drop -f old/rdid
```

DaCHS will then hunt down all traces of the old RD and delete them.

Don't do this without an acute need; such radical measures will clean up DaCHS' mind, but in a connected society, amnesia can be a strain on the rest of the society. In the VO case, `dachs drop -f` might, for instance, cause stale registry records if you had registered services inside of the RD you force-drop.

Key 'dataprodukt_type' not found in a mapping

You will see this if you use `idmaps="*" in parmakers for tables mixing in timeseries mixins. This is the result of the expansion of the asterisk: It just looks for all params defined in the table and produces maps for each of them.`

The timeseries mixins predefine a few params (`dataprodukt_type`, `dataprodukt_subtype`). You probably will not (and should not) have them in the vars of your parmaker. But that then lets the map undefined.

This means that you cannot use `idmaps="*" in timeseries parmakers. Just enumerate the params you want to map instead (or name them so cleverly that you can catch them with another wildcard).`

'dict' object does not support indexing

You will usually see that when executing some SQL, for instance in combination with Execution of `python script createObscoreView` failed or so. In that situation this simply means that there is a stray percent sign somewhere in the query string; `psycopg2`, our database library, interprets that as a metacharacter and tries to pull fillers from whatever you passed as arguments. Since DaCHS will pass a dictionary for you, this will break (and it would break anyway).

So – find the percent and remove it (or double it if you really want it in).

Invalid script type preIndex for resource elements

This, most likely, is because you have upgraded from DaCHS 1 to DaCHS 2, had dumped the userconfig in version 1 and forgot to read the upgrading hints howDol.html#updates-to-rds-and-similar.

Quick and correct fix: Just remove the script element in your `/var/gavo/etc/userconfig.rd`.

ignoreOn in a rowmaker doesn't seem to work

The most likely reason is that you are testing for the presence of a key that is within the table. This will not work since rowmakers add `key->None` mapping for all keys missing but mentioned in a map (also implicitly via `idmaps`).

If more than one rowmake operate on a source, things get really messy since rowmakers *change* the row dictionaries in place. Maybe this should change at some point, but right now that's the way it is. Thus, you can *never* reliably expect keys used by other tables to be present or absent since you cannot predict the order in which the various table's rowmakers will run.

To fix this, you can check against that key's value being NULL, e.g., like this:

```
<keyIs key="accessURL" value="__NULL__"/>
```

You could also instruct the rowmaker to ignore that key; this would require you to enumerate all rows you want mapped.

module 'gavo.utils' has no attribute 'pyparsingWhitechars'

You are probably running `dachs limits` on the DaCHS that comes with Debian bullseye (11). Its `dachs --version` would give you something like `Software (2.3) Schema (26/26)`. If that is true: Well, don't run DaCHS limits on that version. The functionality has received a fairly fundamental revamp in 2.4 and later (cf. <https://blog.g-vo.org/dachs-2-4-is-out-blind-discovery-pretty-datalink-and-more.html>). Dealing with the limits can wait until after you upgrade.

If you *really* need `dachs limits`, I'm afraid you will have to upgrade, preferably by pulling from [our APT repo](#).

DistributionNotFound

When calling DaCHS, you may see tracebacks like:

```
Traceback (most recent call last):
  File "/usr/local/bin/gavo", line 5, in <module>
    from pkg_resources import load_entry_point
    [...]
  File "/usr/lib/python2.6/dist-packages/pkg_resources.py", line 552, in resolve
    raise DistributionNotFound(req)
pkg_resources.DistributionNotFound: gavodachs==0.6.3
```

This is usually due to updates to the source code when you have installed your source in development mode. Simple do `sudo python setup.py develop` in the root of the source distribution again.

Another source of this error can be unreadable source directories. Please check that the user that's trying to execute the command can actually read the sources you checked out.

500 FileNotFoundError when trying to pull jquery on bulleye

The DaCHS package (version 2.7) in Debian Bullseye (version 12.x) has a broken Javascript combiner. That means that basically everything to do with Javascript in the web interface is broken. To fix this, override the built-in resource like this:

```
cd /var/gavo/web/nv_static/
mkdir js
cd js
curl -O https://dc.g-vo.org/static/js/jquery-gavo.js
sudo service dachs restart
```

Of course, the preferred fix is to upgrade to a newer version; you will have to add our repository then, though (see the installation guide).

DaCHS limits says “Table X does not have columns DaCHS knows how to get a coverage from.”

dachs limits evaluate RD elements like:

```
<updater sourceTable="lines"/>
```

and has heuristics to figure out what area in space, time, and spectrum the referenced table covers. The error message means that these heuristics did not work out. Usually, that is because there is no coverage on these axes; for instance, in LineTAP there is (usually) no point in trying to assign spatial or temporal coverage to a table. But then perhaps DaCHS is too dumb, in which case you might want to report a failure.

Be that as it may: The fix is to restrict the updater to the axes that are actually there. The `updater` element has specific attributes for that, `spaceTable`, `spectralTable`, and `timeTable`. So, for a LineTAP table you would say:

```
<updater spectralTable="lines"/>
```

You can still specify coverage manually, by the way, as perhaps in the case of a LineTAP table resulting from observations of a particular object at a particular time; but there is no way DaCHS could infer such information from a LineTAP table, so there’s no place for `xTable` attributes in that situation.

“job hung in EXECUTING” when doing TAP async jobs

This happens when an async job has spawned the worker process, but that worker process fails to start working on the job. That is usually because that worker process crashes early on. If you want to diagnose that, create a TAP job with, for instance, TOPCAT. Once it has failed, as dachsroot say `psql gavo` and in there:

```
update tap_schema.tapjobs set phase='QUEUED' where jobid='whatever'
```

where “whatever” is the last segment of the job URL (something like “ekzonxsh” or so). Then, as the gavo user, run:

```
dachs --enable-pdb taprun whatever
```

This ought to lead you to where the problem is.

A common cause for this around DaCHS 2.11 in Debian bookworm is fairly easy to remedy. It happens because astropy becomes very unhappy if it cannot write to the current user’s home directory. Since the server generally runs as the user gavo, that user must be able to write to `/home/gavo`. So, the simple fix is:

```
sudo chown gavo /home/gavo
```

“ERROR: permission denied to alter role” during dachs upgrade

This may happen in DaCHS \leq 2.12.1 because DaCHS tries to fix floating point representations for its query roles and its admin role (gavoadmin, cf. [dachs-database-interface](#)). In newer postgreses, this is no longer necessary, so newer DaCHSes don't try it.

To fix the problem, either upgrade to DaCHS $>$ 2.12.1 directly or run:

```
sudo -u postgres gavo -c "ALTER ROLE gavoadmin CREATEROLE"
```

Actually, modern DaCHSes give gavoadmin CREATEROLE privileges during `dachs init`, so you may want to run that command anyway even if DaCHS in normal operation does not need it.