

# MAM

## How can we clean out old records from the MAM database?

**Issue:** I would like to be able to remove old records from the database.

**Solution:** Since version 7.2.0, you can use the System Prune action in mam-shell (or goldsh, for older versions still using Gold) to clean out old records from the MAM database. Currently (as of Q4 2020), "System Prune" is considered to be in beta, and Adaptive has not documented it due to the lack of site feedback. It has been tested against a large set of production data, and appears to function well. However, we strongly recommend making a backup of your database before making changes to your system with this command.

For more information see the CHANGES file found in the root of the MAM download tarball.

### MAM Example

Switch to the mam user and create a backup of the database

```
#####  
[root]# su - mam  
[mam]$ pg_dump mam > mam-db-backup-nov-18-2020.sql  
#####
```

To prune all data before Nov 1, 2019 run the following command. Be aware that this could take a long time if you have a lot of old data.

```
#####  
[mam]$ /opt/mam/sbin/mam-shell System Prune RetentionTime:="2019-11-01"  
Confirm:=True  
#####
```

After pruning it is advisable to perform a vacuum to allow Postgres to reclaim storage space from the deleted tuples. This could also take a long time.

```
#####  
[mam]$ exit  
[root]$ su - postgres  
[postgres]$ psql -d mam -c "VACUUM (VERBOSE, ANALYZE)"  
#####
```

### Gold Example (for older versions still using Gold)

# MAM

General goldsh usage applies, as explained in the MAM Admin Guide.

To remove 3 months worth of data, relative to now, run this command from the goldsh prompt:

```
gold> System Prune RetentionPeriod:="3 months !" Confirm:=True
```

As stated below, if you leave off the Confirm option, it won't actually perform the operation; instead, it will just show you what the command would have done had it included Confirm.

Cluster admins may also run Postgresql's vacuumdb tool to do garbage collection on the database to help with the size.

<https://www.postgresql.org/docs/9.2/sql-vacuum.html>

Enabled database pruning via a new System Prune action. This action requires a RetentionPeriod option and a Confirm option. Without the Confirm option asserted this command tells you the date prior to which data will be deleted. The RetentionPeriod option is of the form [^|~|!] where the period may be something like 1 year, 3 months or 5 minutes. The period of data that is retained (not deleted) will not be less than the time period specified. The modifiers indicate whether the time period should be relative to now (!), or relative to the start of this (~) designator (month or year, etc.), or relative to the start of the first (^) designator (month or minute, etc.). So assuming the command was issued at "2012-08-20 12:34:56", if you specified "3 months !" as the retention period, data after "2011-08-20 12:34:56" would be retained, if you specified "3 months ~" as the retention period, data after "2012-05-01 00:00:00" would be retained, and if you specified "3 months ^" as the retention period, data after "2012-03-01 00:00:00" would be retained.

Database pruning may be useful to some sites because database queries and reports may become slower over the years after millions of transactions, usage records and allocation log records accumulate. This command attempts to safely trim back historical and log data prior to a specified time period while maintaining database self-consistency. After pruning the database to a certain retention time, reports and queries must only be requested for data after that time or the results may be inconsistent or undefined. Also added a RetentionTime attribute so you can specify the time that you want to be kept directly. RetentionTime and RetentionPeriod are mutually exclusive options. One of the two must be specified (along with Confirm).

Unique solution ID: #1077

Author: Rick McKay

Last update: 2020-11-18 21:05