

# Monitoring metadata refresh

The SP Proxy (<https://login.terena.org/wayf>) uses metadata that is generated by another host: <https://pioneer.terena.org/mr>.

This host polls various metadata sources from cron, using SimpleSAMLphp's metarefresh module.

The reason for this dual VM is that the metarefresh module sometimes would get stuck validating XML signatures, and then the SP proxy would also hang, which sucks.

This behaviour has been fixed, but the dual VM set-up is still a good idea, so it will stay this way.

When [pioneer.terena.org/mr](https://pioneer.terena.org/mr) has finished, it uses rsync to synchronise the directory that holds the metadata to [login.terena.org/wayf](https://login.terena.org/wayf).

[pioneer.terena.org/mr](https://pioneer.terena.org/mr) is currently configured to poll a couple of dozen URLs for metadata.

This process works, and it has been improved so that when there are errors with the refreshing/polling, the previous/cached metadata is re-used, instead of it being nuked.

This is already an improvement, and when it happens, an error is logged, however this requires an administrator to look at the log files, which does not happen.

So, the error goes unnoticed, and eventually the cached metadata will expire, because it has a set lifetime embedded.

At this stage IdPs will start to disappear from the metadata until it's completely empty, until there are no valid entries any more and the complete set will have disappeared.

At this point the service will be unavailable.

During the SimpleSAMLphp hackathon on 27 May 2015 an attempt has been made to add some form of monitoring capabilities to metarefresh.

This can be done in many ways, however not all were considered:

- Write any custom Nagios checks and let the Nagios server poll the URLs. This option was rejected because the network checks should be done from the same host that does the metarefresh process. Otherwise subtle changes in network access lists, firewall, OS and library versions might yield different results.
- Run separate Nagios checks through/from the same host that does the metarefresh process. This approach was rejected because with dozens of URLs, any non-responsive URLs would significantly increase the time it takes to run the Nagios plugin.

In the end the follow approach was chosen. The metarefresh process already logs all possible errors, and also stores any Conditional GET values for each URL in a state file.

By slightly adapting the metarefresh module, it was possible to create an additional state file that holds error information about the URL.

This state file is then parsed by Nagios through `check_by_ssh`. This is not the best approach security wise, and could later be improved so that the data is exposed over HTTPS and protected by supplying a secret variable.

Several issues were encountered:

- When the metarefresh config is changed, for example URLs are added/deleted/changed, the Nagios configuration needs to be changed as well, the Nagios process needs to be reloaded, and a forced check of the services needs to be scheduled.
- The name of the service would be the src URL, but this contains lots of 'illegal' characters as far as Nagios is concerned. A first attempt to fix work around this was by stripping the path from the URL. This means that the name would be just the base URL, which makes sense to admin, but does not (yet?) contain illegal characters. However, this turned out to be a problem because metadata is polled from several sources that reside on the same base URL (<http://my.org/saml/ldap1>, <http://my.org/saml/ldap2>, etc). The base URL is the same in this case, which is also not allowed by Nagios. The final work-around was to append a hash of the URL to the base URL. This way it is easy to distinguish what domain/IdP has an issue, and the string will be unique for each URL.
- Because the way metarefresh is implemented, the URL state file is being overwritten for each URL, and only at the end will it yield the full status array. This causes a race condition with the Nagios check, it can happen that the Nagios check for a specific URL happens during the metarefresh, when there is no status information yet on that URL. A work-around was implemented by scheduling a simple `cp` command at the end of metarefresh cronjob. This will make sure the state file will always have all the URLs.

Because the logic of metarefresh is used, the same errors can be detected. This comes down to any connection errors, and all HTTP response codes other than 200 (OK) and 304 (Not Modified).

A few basic tests were done to confirm the functionality, such as having a URL return an HTTP 500 error:

- [illegible]