

Optimizing Matterhorn System Resources

 This document is only relevant for Avalon versions prior to 7.0

Avalon provides functionality for ingesting and transcoding large batches of digital materials. This ingest process requires significant system resources.

Matterhorn, the software managing ingest workflow processing, can sometimes require more resources than the typical server configuration provides. In particular, with large batches the number of concurrent processes required by Matterhorn can sometimes exceed the typical default of 1024. This results in an unacceptably high batch ingest failure rate.

This limitation can be addressed in two different ways: decreasing the number of concurrently running batch jobs, or increasing the number of allowable concurrent processes for the matterhorn user.

Decreasing the number of concurrently running batch jobs has a negligible effect on the overall throughput of batch processing. That is, the time required to process all of the batches will not be greatly affected. However, because the batches will be run in series instead of in parallel, the later batches will not appear to have been started until much later. This latency can be unsettling for the user. For this reason, it is recommended to take the second approach.

To increase the number of allowable concurrent processes for the matterhorn user, a new configuration file can be added to the system configuration:

```
su - root
touch /etc/security/limits.d/99-matterhorn.conf
echo "matterhorn    hard    nproc    4096" >> /etc/security/limits.d/99-matterhorn.conf
echo "matterhorn    soft    nproc    4096" >> /etc/security/limits.d/99-matterhorn.conf
chmod 644 /etc/security/limits.d/99-matterhorn.conf
```

Cleaning work directory & using external DB:

Mohamed Abdul Rasheed (mohideen@umd.edu) noted on Slack (2017-12-18) that cleaning work dir helped with large batch, using external DB instead of the built-in H2 might help too.

Good morning all, Here is a followup on the matterhorn issues. I found that deleting the matterhorn work directory (with stop/start) makes the encoding run faster and successfully complete again.

When I looked into the past batch ingest email notifications, I discovered that during the loading of the first month data (~750 items) the encoding completed successfully for all the items. After which both the speed of encoding and successful completion of encoding started decreasing rapidly. I suspected that the matterhorn's growing H2 db size might be causing the problem and tried deleting the entire work directory which seems to have helped. I am planning to configure matterhorn with an external DB in hopes that would solve the problem.

Also, I just wanted to note that I have not made any changes to the server cpu/memory configuration.