



Extreme Search® Manual

Version 2024.10

<https://lewis-rhodes.com>

support@lewis-rhodes.com

-
- [6 Data Storage Best Practices](#)
 - [6.1 Summary](#)
 - [6.2 File Sizes](#)
 - [6.3 File Quantity](#)

6 Data Storage Best Practices

6.1 Summary

Scans will still return correct results if you don't follow these practices. But you likely won't see speeds around 90 GB/sec unless you follow these.

- Have the majority of your files be between 1 MB and 20 GB.
- Run on at least 3,000 files when using gluster to distribute files. If you don't use gluster to distribute files, make sure you have even amounts of data (bytes) per backend and at least 1,500 files in total. Note that if you have fewer files, its likely the absolute time for scan will be fairly low (a few seconds) even if performance (GB/sec) isn't very good.

6.2 File Sizes

The NPU's performance can vary based off of file sizes. When files are smaller than 1 MB the overhead of opening, closing, and reporting files starts to become substantial. In our tests, we have noticed that scans run on files which average around 200 KB in size have GB/sec speeds half of scans run on files which average around 1200 KB. Having many small files also increases the time functions such as `ls`, `stat`, and `sizes` take. Running `NPUGlusterClient.ls(files)` on 2 million files took ~10 seconds, but running it on ~23,000 files took ~0.14 seconds.

On the other end, having files too large causes too few files as discussed in the following section. I suggest no bigger than 20 GB simply because having files larger than that makes it hard to have enough files to get good performance. In addition, when a large file matches there is still a lot of post analysis work that needs to be done (where in that 20 GB was the match?).

6.3 File Quantity

An ExtremeSearch appliance gets its fast speeds by being "embarrassingly parallel". Depending on your device, there are either 256 (Kuona) or 288 (SmartSSD) "chains" which work in parallel. Having too few files or files not uniformly distributed across the drives reduces the ability for the device to search files in parallel. The 3,000 file number for gluster mentioned earlier is more of a heuristic than a hard number. In relatively small quantities gluster doesn't do the best job uniformly distributing files and it becomes hard to saturate all the 256/288 chains long enough to get good performance. The 1,500 number when forgoing gluster operates on the same principles.

Again, keep in mind that if you have fewer files than this, it likely will not be an issue, given that your scans shouldn't take very long to run on an absolute scale as long as your files aren't too large.