

Tools for understanding file I/O patterns produced by LSST

Fabio Hernandez

fabio@in2p3.fr

IN2P3/CNRS computing center, Lyon , France



- Ongoing work aiming at **understanding the file I/O patterns** induced by LSST stack
- Ultimate goal: to identify what data storage platform best suits the needs of LSST

or alternatively, modify LSST software to take into account the intrinsic limitations of the available storage platforms

Current status

- Developed FUSE-based synthetic file system to **intercept, trace** and **forward** I/O operations to the target underlying file system

`open, read, write, close, flush, stat, fsstat, ...`

- Inspired by existing tools: systemTap, strace, dtrace, ...

none of them fully satisfying for this particular use case

- Features

Works on both Linux and MacOS X

No superuser privilege required

Stand-alone executable

Focus exclusively on I/O operations

Generates easy-to-parse text files: CSV and JSON initially

Usage example

- To trace operations on files under `$HOME/Temp`, mount the file system at `$HOME/trace`

```
$ tracefs --target=$HOME/Temp --mount=$HOME/trace
```

- All the operations referring to files under `$HOME/trace` will be traced

```
$ cat $HOME/trace/hello.txt
```

will produce

```
{ "Uid":501, "Gid":20, "Pid":8058, "User":"fabio", "Group":"staff",  
  "Process":"cat", "OpType":"access", "Start":"2015-02-11T14:38:52.807023059+01:00", "End":"2015-02-11T14:38:52.807069488+01:00", "Path":"/Users/fabio/Temp/hello.txt", "Mask":4}  
  
{ "Uid":501, "Gid":20, "Pid":8058, "User":"fabio", "Group":"staff",  
  "Process":"cat", "OpType":"openfile", "Start":"2015-02-11T14:38:52.807250079+01:00", "End":"2015-02-11T14:38:52.807299616+01:00", "Path":"/Users/fabio/Temp/hello.txt", "IsDir":false, "Flags":0, "Perm":0}  
  
{ "Uid":501, "Gid":20, "Pid":8058, "User":"fabio", "Group":"staff",  
  "Process":"cat", "OpType":"read", "Start":"2015-02-11T14:38:52.807557872+01:00", "End":"2015-02-11T14:38:52.808424381+01:00", "Path":"/Users/fabio/Temp/hello.txt", "Offset":0, "Size":19, "BytesRead":19}  
  
{ "Uid":501, "Gid":20, "Pid":8058, "User":"fabio", "Group":"staff",  
  "Process":"cat", "OpType":"flush", "Start":"2015-02-11T14:38:52.8086749+01:00", "End":"2015-02-11T14:38:52.808700451+01:00", "Path":"/Users/fabio/Temp/hello.txt"}  
  
{ "Uid":501, "Gid":20, "Pid":8058, "User":"fabio", "Group":"staff",  
  "Process":"cat", "OpType":"releasefile", "Start":"2015-02-11T14:38:52.80884041+01:00", "End":"2015-02-11T14:38:52.808854485+01:00", "Path":"/Users/fabio/Temp/hello.txt", "IsDir":false}
```

Current work

- Developing tools to exploit trace files

Python notebook

Source code

- Source code and documentation will be soon on Github