

Improving the traceability of your code with no effort

Sympathy for the BOFH

Luis Fernando Muñoz Mejías

CERN IT-DI

- 1 Why a log infrastructure
- 2 How we can do it
- 3 Conclusion

Some loose definitions

log A file containing information on the activities of a system, and not related with the output the user expects from it. It is never delivered to users.

system administrators Noun, plural

- 1 Anyone closer to the service instances you use than you (unless you are a systems administrator).
- 2 The guys trying to make that someone else's code runs as a third person expects.

developer Perhaps, YOU. A person writing software that has to run somewhere else than his own laptop and by someone distinct than himself.

- 1 Why a log infrastructure
- 2 How we can do it
- 3 Conclusion

- There is no bugfree software
 - Unit tests find only what you thought about, not what you missed
 - Some bugs make it into production systems
- Users make mistakes
 - They need information to know how to correct their behaviour
- We must have some information on what our process is doing




Figure: Basic human workflow in WLCG

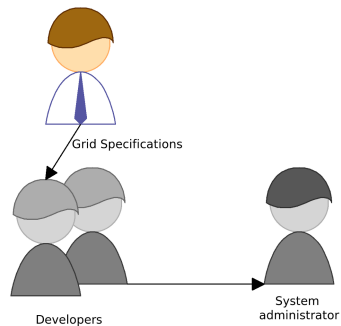


Figure: All logs must be stored centrally and kept for 6 months




Figure: System administrators must transfer logs and keep them secure




Figure: Wait! Aren't we missing someone?

Sympathy for the BOFH



- 1 Why a log infrastructure
- 2 How we can do it
- 3 Conclusion

- Print stuff directly
 - But it interferes with the program's valid output
 - And we can't really choose what we want to see
- Print stuff conditionally
 - And fill our code with if-flag-then-print constructions
- Print to a file
 - And have a filehandle as a global variable (globals are BAD)
 - And what about thread safety?
- Write our own modules/classes for logging
 - But it doesn't solve the central storage problem
- And timestamps are not simple either

Syslog, that unknown

- It's not just a daemon
- It defines a protocol for sending **securely** messages over the network
- It allows for filtering out messages, or sending different events to different places
- And you don't need to be root to use it

```
from syslog import openlog, syslog
openlog("Program_Name") # DON'T USE WHITESPACES!!!
syslog("Hello, world!!!")
```

- But you usually need to be root to read logs, or at least to set up the filters
 - Bad for the development phase
- Only available on Unix-like OS
 - And you might want to program for you-know-who

Apache's log4j and friends

- Modules for different languages that allow for clean logging
- Allow for filters
 - No need to be root for setting them
 - Allow for syslog! :)
- Allow for persistent configuration
 - No need of complex options in the command line
- And they keep the code clean
- Friends (not necessarily maintained by the Apache Foundation)
 - log4cxx
 - log4perl
 - log4py
 - log4r

Examples

In Perl

```
$logger = Log::Log4perl->get_logger(  
    'house.bedrm.desk.topdrwr');  
$logger->debug('this is a debug message');  
$logger->info('this is an info message');  
$logger->warn('etc');  
$logger->error('..');  
$logger->fatal('..');
```

Setting up filters

- In your code

```
logger->setLevel(log4cxx::Level::getInfo());
```

- But users needing some different filters will have to rebuild the code

- In files

- A configuration file chooses where each message should be sent to
- Filters are no longer developer's responsibility :)

```
log4j.rootLogger=DEBUG, A1
```

```
log4j.appender.A1=org.apache.log4j.ConsoleAppender
```

- Lots of flexibility for your users with no effort on your side




- 1 Why a log infrastructure
- 2 How we can do it
- 3 Conclusion**

There is life beyond log4XXX

- You can choose other logging packages
- Don't re-invent the wheel
- But please, **ensure** they can interact with syslog
 - For sysadmins there is no life beyond syslog

- Each time you make a sysadmin cry, God kills a kitten

Thank you!!

-  Francisco de Goya y Lucientes, Fusilamientos del 3 de Mayo de 1808 en La Moncloa
-  RFC5424
-  RFC3339