# BoF: Analyzing Parallel I/O

**Julian Kunkel**[1] **Philip Carns**[2] Shane Snyder Huong Luu
Matthieu Dorier Wolfgang Frings Glenn Lockwood

1 German Climate Computing Center
2 Argonne National Laboratory

Supercomputing 2015

# I/O Monitoring: Introduction

## Goals for I/O monitoring

- Explore and diagnose I/O behavior e.g. performance and availability
- Notify if something unusual/critical happens
- Prescribe (enact) behavioral changes to mitigate issues & optimize

## Facets of monitoring and analysis

- Analysis: Descriptive, predictive, prescriptive
- Perspective: System-level vs. application-level
- Scope: Client side, middleware, OS, file system, block devices
- Temporal resolution: Profile (statistics) vs. app phases vs. tracing

# Introduction

## A typical HPC I/O stack

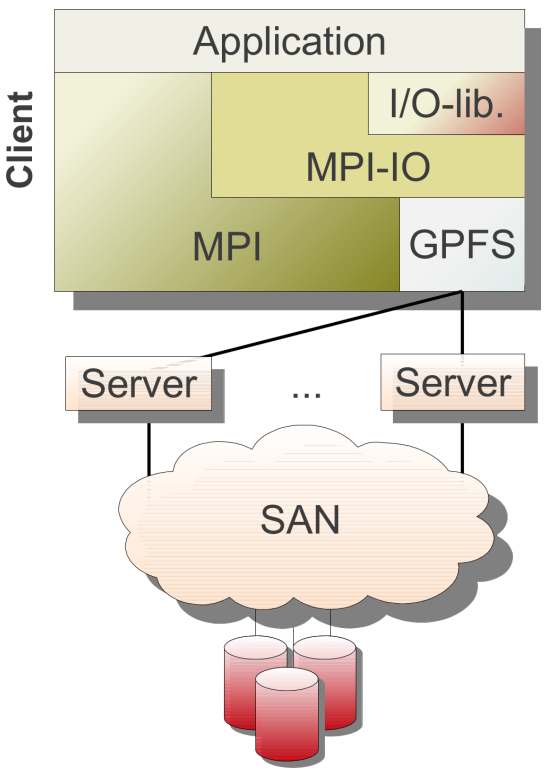


## I/O and its cause and effect chain

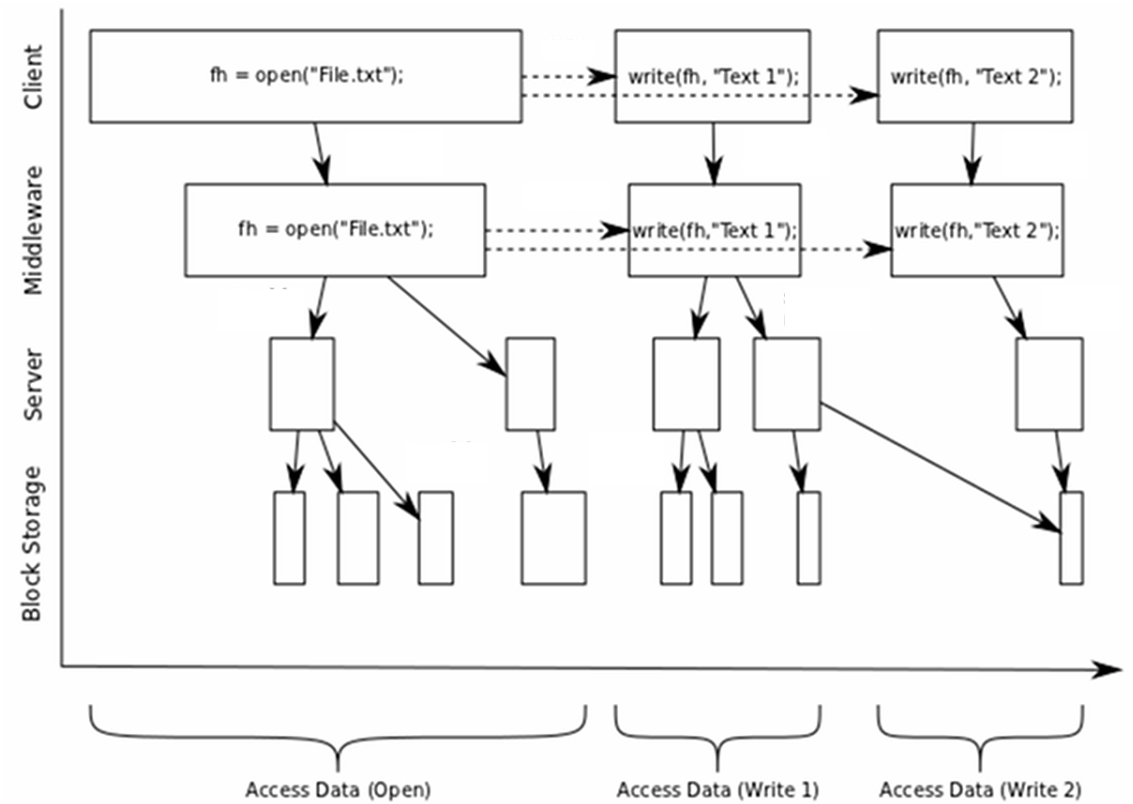

# Existing Tool Landscape

- Local tools: blktrace, */proc*, ...
- System-wide I/O hardware/software statistics
  - May be linked to application information
  - Examples: Ganglia, Nagios, Lustre Monitoring Tool (LMT)
  - Observation: every data center builds its own monitoring tools
- Tracing tools for HPC I/O
  - Typically POSIX layer is captured
  - I/O servers are usually out of scope
  - Examples: TAU, IPM, Score-P, LANL-Trace, IOSIG, PAS2P-I/O, RIOT, ScalaIOTrace, PIOvis

# (Some) Requirements for I/O Analysis Tools

- Data center wide monitoring
  - At least capture relevant/statistical information
- On demand application instrumentation
  - Comprehensive information e.g. via tracing
  - Detailed temporal resolution at best focusing on relevant behavior
- Portability
  - incl. support for non-HPC environments (for wide adoption)
- Ease of use
- Low overhead
- Allow guided or automatic tuning
- Extensibility (support new I/O libraries & I/O research)

# Agenda

1. Introduction
   - I/O Monitoring: Introduction
   - I/O Stacks and Dependencies
2. Agenda for the BoF
   - SIOX: A flexible approach
   - Darshan: state of the project and new features
3. Invited talks
   - The Dashboard: HPC I/O Analysis Made Easy
   - How predictable are HPC application, and why should we care?
   - Parallel I/O Monitoring at JSC
   - Tools and techniques towards a holistic understanding of I/O demands at NERSC
4. Monitoring plans @ DKRZ and discussion about monitoring standards
5. Discussion

# Discussion

- Requirements for future tools?
- How much overhead is acceptable? (1% runtime?)
- What kind of information about I/O accesses is of interest for users?

---

Talk to us!

Darshan Philip Carns <carns@mcs.anl.gov>

SIOX Julian Kunkel <kunkel@dkrz.de>