

## File System Tracing, Replaying, Profiling, and Analysis on HEC Systems



Erez Zadok  
Klaus Mueller  
Stony Brook University  
[www.cs.stonybrook.edu](http://www.cs.stonybrook.edu)



Ethan L. Miller  
UC Santa Cruz  
[www.cs.ucsc.edu](http://www.cs.ucsc.edu)

## Researchers

- Erez Zadok (Stony Brook University)
  - Avishay Traeger (PhD)
  - Ivan Deras Tabora (MS)
  - Nikolai Joukov (PhD, now at IBM Watson)



- Klaus Mueller (Stony Brook University)
  - Aili Li (PhD)



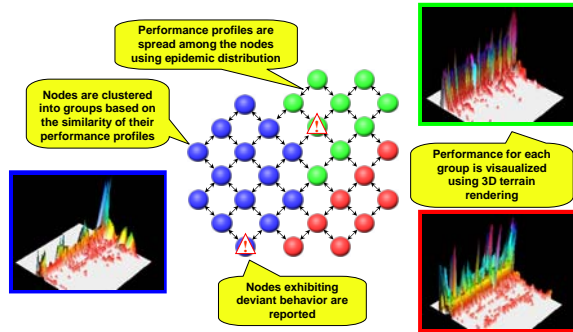
- Ethan L. Miller (UC Santa Cruz)
  - Andrew Leung (PhD)
  - Eric Lalonde (MS)



Erez Zadok and Klaus Mueller (Stony Brook University)  
Ethan L. Miller (UC Santa Cruz)



## The Big Picture



Erez Zadok and Klaus Mueller (Stony Brook University)  
Ethan L. Miller (UC Santa Cruz)

3



## Outline

- Anomaly Detection
- Visual Analytics
- Online Storage Visualization System
- Contributions
- Future
- Needs

Erez Zadok and Klaus Mueller (Stony Brook University)  
Ethan L. Miller (UC Santa Cruz)

4



## Anomaly Detection: Current State

- Profiling with OSprof [OSDI '06]
- Began profiling three file systems on a small cluster: UCSC's Ceph [OSDI '06], IBM's GPFS, NFSv4
- Manually profiling these systems helps us gain experience to use in our automated methods
- Explored several options for profile comparison algorithms
- Implementing prototype of DARC (Dynamic Analysis of Root Causes)
  - ◆ Dynamically instruments code to search for root causes of any behavior seen in a profile (user-space and kernel-space).
- Epidemic profile propagation
  - ◆ Efficient profile distribution & aggregation

Erez Zadok and Klaus Mueller (Stony Brook University)  
Ethan L. Miller (UC Santa Cruz)

5



## Visual Analytics

- Put the expert into the loop
- Create the *Performance Analysis Cockpit*
- Crucial:
  - ◆ Interactive exploration, interrogation
  - ◆ Visual expressiveness
  - ◆ Intuitive user interface
  - ◆ Intuitive mapping of data to visuals
  - ◆ Feedback loop: overview first, zoom + filter, then details on demand (The Shneiderman Mantra of Information Visualization)



Erez Zadok and Klaus Mueller (Stony Brook University)  
Ethan L. Miller (UC Santa Cruz)

6



## Visual Analytics: Current State

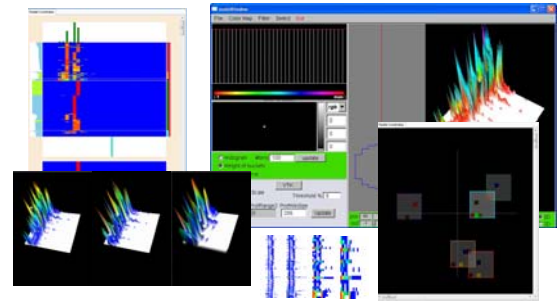
- Interface supporting:
  - ◆ Data visualization as 2D maps and 3D terrain
  - ◆ Filtering with various feature sensitive metrics
  - ◆ Feature-preserving data overviews
  - ◆ Feature-preserving zoom lenses
  - ◆ Various maps for data highlighting
  - ◆ Overview plots created via Multi-Dimensional Scaling
- Exploring algorithms for profile visualization
  - ◆ Zoom into terrain plots to view profiles over time
  - ◆ View cluster plots to obtain an overview of profile population

Erez Zadok and Klaus Mueller (Stony Brook University)  
Ethan L. Miller (UC Santa Cruz)

7



## Stop By the Poster Tonight to See More on the Visualizations!



Erez Zadok and Klaus Mueller (Stony Brook University)  
Ethan L. Miller (UC Santa Cruz)

8



## Online Storage Visualization System

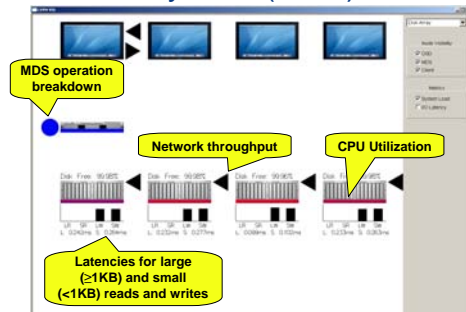
- View metrics of a distributed file system in real time
- Can view the effects of inter-nodal relationships and management policies
- File system nodes collect performance information and forward it to visualization nodes
- Visualization nodes sort data chronologically and send updates to the visualization client application

Erez Zadok and Klaus Mueller (Stony Brook University)  
Ethan L. Miller (UC Santa Cruz)

9



## Online Storage Visualization System (cont.)



Erez Zadok and Klaus Mueller (Stony Brook University)  
Ethan L. Miller (UC Santa Cruz)

10



## Publications

- Operating System Profiling via Latency Analysis
  - ◆ OSDI 2006
  - ◆ N. Joukov, A. Traeger, R. Iyer, C. P. Wright, E. Zadok
- Ceph: A Scalable, High-Performance Distributed File System
  - ◆ OSDI 2006
  - ◆ S. A. Weil, S. A. Brandt, E. L. Miller, D. D. E. Long, C. Maltzahn
- Using Comprehensive Analysis for Performance Debugging in Distributed Storage Systems
  - ◆ MSST 2007
  - ◆ A. Leung, E. Lalonde, J. Telleen, J. Davis, C. Maltzahn
- Scalable Security for Petascale Parallel File Systems
  - ◆ SC 2007
  - ◆ A. Leung, E. L. Miller, S. Jones
- Round-Trip Privacy with NFSv4
  - ◆ StorageSS 2007
  - ◆ A. Traeger, K. Thangavelu, E. Zadok

Erez Zadok and Klaus Mueller (Stony Brook University)  
Ethan L. Miller (UC Santa Cruz)

11



## Broader Impact

- Releases to SNIA IOTTA repository (iotta.snia.org)
  - ◆ Tracfs
  - ◆ Replays
  - ◆ File system traces
- UCSC now hosting IOTTA repository mirror
- Releases on the Web
  - ◆ Ceph (<http://ceph.sourceforge.net>)
  - ◆ OSprof + profile selection tools ([www.fsl.cs.sunysb.edu](http://www.fsl.cs.sunysb.edu))
- Education
  - ◆ Storage visualization system was a UCSC class project
  - ◆ Currently training 3 PhD and 2 masters students
  - ◆ UCSC storage and distributed systems classes are broadcasted to LANL

Erez Zadok and Klaus Mueller (Stony Brook University)  
Ethan L. Miller (UC Santa Cruz)

12



## Future

- Dynamic data analysis
- Feedback from visualization controller to nodes in real-time
  - ◆ CAVE at Stony Brook (CEWIT center)
- Epidemic profile propagation
  - ◆ Algorithms for automatic grouping
- Integrate DARC with visualization
- Test on published HEC-FSIO data
  - ◆ Testbed cluster: 250 nodes, 100TB
- Release to national labs, test

Erez Zadok and Klaus Mueller (Stony Brook University)  
Ethan L. Miller (UC Santa Cruz)

13



## Help Us to Help You...

Feedback on real problems you face

- Feedback on useful visualizations
- Making software and data sets available to us
- Possibility of you running our software
- Main HECURA portal site
  - ◆ Links to supported projects
  - ◆ Links to released software
  - ◆ Papers published
- Coordinate visits to national labs

Erez Zadok and Klaus Mueller (Stony Brook University)  
Ethan L. Miller (UC Santa Cruz)

14



## Q&A

### File System Tracing, Replaying, Profiling, and Analysis on HEC Systems

<http://www.fsl.cs.sunysb.edu/hpcvperf/>

Erez Zadok  
Klaus Mueller  
Stony Brook University  
[www.cs.stonybrook.edu](http://www.cs.stonybrook.edu)

Ethan L. Miller  
UC Santa Cruz  
[www.cs.ucsc.edu](http://www.cs.ucsc.edu)