

OpenSFS Test Cluster Donation



UNIVERSITY OF
LOUISVILLE

J.B. SPEED SCHOOL
OF ENGINEERING

UNIVERSITY of LOUISVILLE

Dr. Nihat Altiparmak, Assistant Professor
Computer Engineering & Computer Science Department
University of Louisville

4/24/2018

Donation Details

- Jan 5, 2018
 - OpenSFS announced call for proposals for donation of their Lustre test cluster
- Feb 9, 2018
 - University of Louisville submitted a proposal
- March 8, 2018
 - University of Louisville has been selected as the recipient of the OpenSFS Test Cluster equipment
- Now we are here to describe our plan with it
- THANK YOU OpenSFS!

Cluster Description

- 32-node cluster
 - 4 MDS nodes, each with
 - 2 x 2.4 GHz 4-core CPU
 - 64 GB memory
 - 4 x 100 GB SSD and 8 x 2 TB HDD storage
 - 4 OSS nodes, each with
 - 2 x 2.4 GHz 4-core CPU
 - 64 GB memory
 - 2 x 100 GB SSD storage and 22 x 2 TB HDD storage
 - 24 compute nodes, each with
 - 1 x 2.4 GHz 4-core CPU
 - 32 GB memory
- Including necessary networking, management, and infrastructure equipment

Our Research Group

- Computer Systems Laboratory (**CSL**) @UofL
 - <http://cecs.louisville.edu/csl/>
- One faculty, two PhD students, two Masters students, and four undergraduate students
- Contemporary topics being investigated:
 - High Performance Parallel I/O
 - [TC '16], [MASCOTS '14], [TOS '13], [CLUSTER '12], [ICPP '12], [TPDS '12]
 - Self-Optimizing Storage Systems
 - [HiPC '16]
 - Energy-efficient Data Retrieval and Task Placement
 - [CLUSTER '17], [BDCAT '17], [MASCOTS '16]

Our Plan with the Cluster

- Main Goals:
 - Promote innovation and adoption of open-source scalable storage technologies
 - Generate practitioners and researchers with expertise in storage systems

Our Plan with the Cluster

- Research Goals:

- Developing novel, high-performance, and energy-aware distributed:

- data placement,
- data reorganization, and
- data retrieval strategies

considering the internal characteristics of

- SATA HDD,
- SATA SSD, and
- NVMe SSD devices.

- Collaborating with other groups and running real world applications on the cluster:

- RNA-Seq pipeline, tumor growth simulation, etc.

Our Plan with the Cluster

- Research Goals:
 - NVMe-aware caching, tiering, buffering, and I/O scheduling techniques, and their effect on system's overall energy consumption
 - MDS and OSS nodes will be upgraded with NVMe SSD devices to enable NVMe-related research
 - Power meters will be used to measure the energy consumption
 - Host managed (aka. Software-defined or “user programmable”) open channel and/or multi-stream SSDs will be investigated for better performance through increased internal parallel I/O and reduced garbage collection

Our Plan with the Cluster

- Teaching/Training Goals:
 - The infrastructure will be integrated into Operating Systems (undergraduate), Distributed Systems (graduate), and Storage Systems (independent study) courses
 - Course projects will be assigned and students will be provided access to the cluster to implement their projects

Research and Data Dissemination

- All developed software and tools will be released under free and open source GNU GPLv3 license
- Teaching and outreach material will be released under Creative Commons (CC) license allowing free adoption, modification, and redistribution
- Produced I/O traces will be released under SNIA IOTTA license allowing free redistribution
- Publications will be released under the copyrights of the publishers
- All materials will be shared publicly through the webpage of the Computer Systems Lab at:
 - <http://cecs.louisville.edu/csl>

Thank You!

Questions?

Related CSL Publications

- **[BDCAT '17]** Big Data Aware Virtual Machine Placement in Cloud Data Centers. Logan Hall, Bryan Harris, Erica Tomes, and Nihat Altiparmak. 4th IEEE/ACM International Conference on Big Data Computing, Applications and Technologies, Austin, Texas, December 2017. Acceptance rate: 29% (27/93).
- **[CLUSTER '17]** A Comparative Study of HDD and SSD RAIDs' Impact on Server Energy Consumption. Erica Tomes and Nihat Altiparmak. 19th IEEE International Conference on Cluster Computing, Honolulu, Hawaii, September 2017.
- **[HiPC '16]** Dynamic Data Layout Optimization for High Performance Parallel I/O. Everett N. Rush, Bryan Harris, Nihat Altiparmak, and Ali Saman Tosun. 23rd IEEE International Conference on High Performance Computing, Hyderabad, India, December 2016. Acceptance rate: 25% (40/160).
- **[MASCOTS '16]** Exploiting Replication for Energy Efficiency of Heterogeneous Storage Systems. Everett N. Rush and Nihat Altiparmak. 24th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems, London, United Kingdom, September 2016. Acceptance rate (Short Paper): 24.7% (43/174).
- **[TC '16]** Multithreaded Maximum Flow Based Optimal Replica Selection Algorithm for Heterogeneous Storage Architectures. Nihat Altiparmak and Ali Saman Tosun. IEEE Transactions on Computers, 65(5):1543-1557, May 2016.
- **[MASCOTS '14]** Continuous Retrieval of Replicated Data from Heterogeneous Storage Arrays. Nihat Altiparmak and Ali Saman Tosun. 22nd IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems, Paris, France, September 2014. Acceptance rate: 20.3% (39/192).
- **[TOS '13]** Generalized Optimal Response Time Retrieval of Replicated Data from Storage Arrays. Nihat Altiparmak and Ali Saman Tosun. ACM Transactions on Storage, 9(2):5:1-5:36, July 2013.
- **[CLUSTER '12]** Replication Based QoS Framework for Flash Arrays. Nihat Altiparmak and Ali Saman Tosun. 14th IEEE International Conference on Cluster Computing, Beijing, China, September 2012. Acceptance rate: 28.8% (58/201).
- **[ICPP '12]** Integrated Maximum Flow Algorithm for Optimal Response Time Retrieval of Replicated Data. Nihat Altiparmak and Ali Saman Tosun. 41st International Conference on Parallel Processing, Pittsburgh, Pennsylvania, September 2012. Acceptance rate: 28.3% (53/187).
- **[TPDS '12]** Equivalent Disk Allocations. Nihat Altiparmak and Ali Saman Tosun. IEEE Transactions on Parallel and Distributed Systems, 23(3):538-546, March 2012.