

The Mutrino 2/15-6/16 Dataset (12/16 Release)

About the Mutrino 2/15-6/16 Dataset (12/16 Release)

This dataset consists of cleansed logs from the Alliance for Computing at Extreme Scale (ACES) Trinity supercomputer Application Readiness Testbed (ART) system, Mutrino. Trinity is a joint effort of the Alliance for Computing at Extreme Scale (ACES) partnership between Los Alamos National Laboratory (LANL) and Sandia National Laboratories (SNL) as part of the NNSA Advanced Simulation and Computing (ASC) Program.

Mutrino [1], is sited at Sandia National Laboratories. It is a Cray XC40 system, utilizing the Cray Aries Interconnect and, other than size, is identical to the full Trinity system. During this dataset time frame, Mutrino consisted of 100 Haswell nodes. Subsequent to this time frame, Mutrino has been increased to include an additional 100 KNL nodes.

This data set is a unique resource for resilience studies, covering the over the first year of the lifetime of the system, including standup, and a variety of naturally occurring and induced network, electrical, thermal, and functional failures. These include Aries errors; DIMM failures; induced facilities failures, including power loss and fan failure; and thermal-related clock throttling. The dataset is complete except some items, such as user IDs, IP addresses, and paths, have been anonymized and some security related events, such as ssh, have been removed. The cleansing process may result in some side effects or artifacts.

This dataset is a superset of the previously released Mutrino 2/15-5/15 Dataset [2], however a different cleansing methodology has been used. A timeline of events and related analysis for that 3 month time frame can be found in [1].

Using and Citing the Dataset

The Mutrino 2/15-6/16 Dataset is being released by the “Holistic Measurement Driven Resilience (HMDR): Combining Operational Fault and Failure Measurements and Fault Injection for Quantifying Fault Detection, Propagation and Impact” project to the research community in support of Extreme-Scale HPC Resilience Research. HMDR is funded by the U.S. Department of Energy Office of Advanced Scientific Computing Research.

Use the following citation for the dataset:

J. Brandt, A. Gentile, and J. Repik, “Mutrino Dataset 2/15-6/16 (12/16 Release)”. SAND2016-12310 O, Sandia National Laboratories. 2016. [Online] Available: <http://portal.nersc.gov/project/m888/resilience/datasets/mutrino/mutrino1yr-v122016.tgz>.

Acknowledgements

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research.

References:

[1] J. Brandt, A. Gentile, C. Martin, J. Repik and N. Taerat, "New Systems, New Behaviors, New Patterns: Monitoring Insights from System Standup" in *Proc. 2015 IEEE Int'l Conf. on Cluster Computing (CLUSTER)*, Chicago, IL. 2015 pp. 658-665. DOI:10.1109/CLUSTER.2015.116

[2] J. Brandt, A. Gentile, and J. Repik, "Mutrino Dataset 2/15-5/15". SAND2016-2449 O, Sandia National Laboratories, 2016. [Online] Available:

<http://portal.nersc.gov/project/m888/resilience/datasets/mutrino/logs.051715.cr.tgz>