

# The Mutrino 2/15-5/15 Dataset

## About the Mutrino 2/15-5/15 Dataset

This dataset consists of data 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, which is sited at Sandia National Laboratories, is a Cray XC40 system, utilizing the Cray Aries Interconnect, consisting of 100 nodes, which, other than size, is identical to the full Trinity system. This data set is a unique resource for resilience studies, covering the first ~100 days 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. In addition, a system software upgrade occurred during this time, resulting in failures of the power capping functionality and the sensor data reporting functionality. Timelines of the system dataset, annotated with events, such as the software upgrade and event resolution due to system administrator actions are shown below. More information about Mutrino and the events in the dataset can be found in [1]. The dataset is complete except for some items, such as user IDs, IP addresses, and paths, replaced with keys that can still be used to make associations across log lines and some security related events, such as ssh, have been removed.

## Using and Citing the Dataset

The Mutrino 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 Mutrino Dataset:

J. Brandt, A. Gentile, and J. Repik, “Mutrino Dataset 2/15-5/15”. SAND2016-2449 O. 2016. [Online] Available at <http://portal.nersc.gov/project/m888/resilience/datasets/mutrino/logs.051715.cr.tgz>.

## Acknowledgments

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.




Figure 1 (from [1]) Annotated event and log timeline of the Mutrino 2/15-5/15 Dataset, containing network, thermal, electrical, and functional errors. A histogram of the number of log lines per file type per day (raw dataset) is given with events of interest marked.

**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