Data, Information, and Code Availability

James Nunez
Los Alamos National Lab

LA-UR 06-5734

Data/Code/Information Availability

An important area to enable us to help ourselves and enable others to help us. One of the issues we heard loud and clear from the HECIWG FSIO 2005 Workshop

- **Machine/Environment information**
- **Failure**
- **Usage**
- **Event**
- **Synthetic code and traces**
- **Real app traces**
- **Synthetic benchmarks**
- **Real app I/O routines**
- **Real app simulated I/O fingerprints**

Machines/Environment Info

- Most HEC sites make machine lists, sizes, types available
- Many other sites will make machine lists, sizes, types available as well.

LANL: Machines/Environment Info for Data Provided

system CMU paper number	system data machine number	system type	number nodes	number cpus	cpus/node	install date	production date	decommision date	fru	mem per node	cpu type	number of interconnects	use type
1	7	smp	1	8	8	before tracking	before tracking	Dec-99	part	16	3	0	compute
2	24	smp	1	32	32	before tracking	before tracking	Dec-03	part	8	7	1	compute
3	22	smp	1	4	4	before tracking	before tracking	Apr-03	part	1	6	0	compute
4	8	cluster	164	328	2	Mar-01	Apr-01	current	part	1	4	1	compute
5	20	cluster	512	2048	4	Oct-01	Dec-01	current	part	16	2	2	compute
6	21	cluster	128	512	4	Aug-01	Sep-01	Jan-02	part	16	2	2	compute
7	18	cluster	1024	4096	4	Mar-02	May-02	current	part	16	2	2	compute
8	19	cluster	1024	4096	4	Aug-02	Oct-02	current	part	16	2	2	compute
9	3	cluster	128	512	4	Aug-03	Sep-03	current	part	4	2	1	compute
10	4	cluster	128	512	4	Aug-03	Sep-03	current	part	4	2	1	compute
11	5	cluster	128	512	4	Aug-03	Sep-03	current	part	4	2	1	compute
12	6	cluster	32	128	4	Aug-03	Sep-03	current	part	16	2	1	compute
13	14	cluster	128	256	2	Aug-03	Sep-03	current	node	4	2	1	compute
14	9	cluster	256	512	2	Aug-03	Sep-03	current	node	4	2	1	compute
15	10	cluster	256	512	2	Aug-03	Sep-03	current	node	4	2	1	compute
16	11	cluster	256	512	2	Aug-03	Sep-03	current	node	4	2	1	compute
17	13	cluster	256	512	2	Aug-03	Sep-03	current	node	4	2	1	compute
18	12	cluster	512	1024	2	Aug-03	Sep-03	current	node	4	2	1	compute
19	16	cluster	16	2048	128	Oct-96	Dec-96	Sep-02	part	32	1	4	compute
20	2	cluster	49	6152	128	Nov-96	Jan-97	current	part	128	1	12	graphics.compute
21	23	cluster	5	544	128	Oct-98	Oct-98	Dec-04	part	128	1	4	compute
22	15	numa	1	265	256	Nov-04	Nov-04	current	part	1024	5	0	compute

Failure/Usage/Event Data

- **Over the last two decades, very little data has been provided on failure, usage, and events. Many papers cite a DEC VAX environment for 128 weeks.**
- **A quote from a recent CMU paper:**
  - Unfortunately, obtaining access to failure data from modern, large-scale systems is difficult, since such data is often sensitive or classified. Existing studies of failures are often based on only a few months of data, covering typically only a few hundred failures [14, 17, 12, 13, 11, 5]. Many of the commonly cited studies on failure analysis stem from the late 80's and early 90's, when computer systems where significantly different from today [1, 2, 4, 10, 14, 7, 8]. Finally, none of the raw data used in the above studies has been made publicly available for use by other researchers.
- **Recent LLNL paper on BG/L Failure**

(Ramendra K. Sahoo, Yinglung Liang, Yanyong Zhang, Morris Jette and Anand Sivasubramaniam. Blue Gene/L Failure Analysis and Prediction Models. *International Conf. of Dependable Systems and Networks (DSN2006)*. IEEE, March 2006.)

- **HP Labs Mary Baker is helping get HP motivated**
- **Library of Congress is interested in providing their data**
- **We hope other HEC sites will begin to provide data soon**
- **We hope to coalesce into a clearinghouse somehow**

LANL: Failure/Usage/Event Data Sets Available (almost 9 years coverage in some cases)

description	size	records	name
all systems failure/interrupt data 1996-2005	2963538	23741	LA-UR-05-7318-failure-data-1996-2005.csv
system 20 usage with domain info	51675641	489376	LA-UR-06-0803-MX20_NODES_0_TO_255_DOM.TXT
system 20 usage with node info nodes number from zero	43926669	489376	LA-UR-06-0803-MX20_NODES_0_TO_255_NODE-Z.TXT
system 20 event info nodes number from zero	33120015	433490	LA-UR-06-0803-MX20_NODES_0_TO_255_EVENTS.csv
system 15 usage with node info nodes number from zero	2416139	17823	LA-UR-06-0999-MX15-NODE-Z.TXT
system 16 usage with node info nodes number from one	321293488	1630479	LA-UR-06-1446-MX16-NODE-NOZ.TXT
system 23 usage with node info nodes number from one	60674531	654927	LA-UR-06-1447-MX23-NODE-NOZ.TXT
system 8 usage with node info nodes number from one	67291020	763293	LA-UR-06-3194-MX8-NODE-NOZ.TXT

Real Application Traces, Benchmarks, and Real Application I/O Fingerprints

- **Traces, I/O kernels/footprints etc.**
  - LLNL: Application I/O kernel and benchmarks
  - U Chicago: Flash I/O and other open codes
  - We hope other HEC sites will begin to provide data/code soon
  - We hope to coalesce into a clearinghouse somehow
- **Benchmarks**
  - Dozens of I/O benchmarks (xdd, bonnie, …)
  - DARPA HPCS has I/O scenarios and someday benchmarks
  - Other orgs like SNIA have benchmarks
  - Should we coalesce benchmarks into a clearinghouse somehow?

LANL: Real Application Traces, Benchmarks, and Real Application I/O Fingerprints

- **MPI-IO based synthetic**
  - Open Source Benchmark at http://public.lanl.gov/jnunez
- **POSIX based synthetic**
- **Synthetic Traces**
- **Real Application Traces**
  - working towards real app simulated I/O fingerprint replay
  - UC Santa Cruz Rosie Wacha: Synthetic Parallel Application Generator
  - Intel/CMU Mike Mesnier

Available LANL Data

- http://institutes.lanl.gov/data/
- http://www.lanl.gov/computing_research_data

Computer Failure Data Repository (CFDR)

- **Central place to put failure data, usage data, event data etc., or pointers to sites that contain these data**
- **Similar to Plant Lab or CRAWDAD (the Community Resource for Archiving Wireless Data At Dartmouth, and is operated by Dartmouth College under a grant from the National Science Foundation.**
- **Probably will be a USENIX site for neutrality**
- **Will be looking for lots of donors of data**
- **May need some modest $ for support at some point**
- **Working licensing issues/liabilities**
- **Who is working to make this happen**
  - bianca@cs.cmu.edu
  - breed@yahoo-inc.com
  - fox@cs.berkeley.edu
  - garth@cs.cmu.edu
  - ggrider@lanl.gov
  - kimberly.keeton@hp.com
  - mary.baker@hp.com
  - remzi@cs.wisc.edu