Failure Trace Archive

FOR IMPROVING THE RELIABILITY OF DISTRIBUTED SYSTEMS

http://fta.inria.fr

Purpose

The **Failure Trace Archive** (FTA) is centralized public repository of availability traces of distributed systems, and tools for their analysis.
The purpose of this archive is to facilitate the design, validation, and comparison of fault-tolerant models and algorithms.
In particular, the FTA contains the following:

- availability traces of distributed systems, differing in scale, volatility, and usage
- scripts and tools for analyzing these traces
- links to complementary traces (for example, workload or application traces)

http://fta.inria.fr

Format

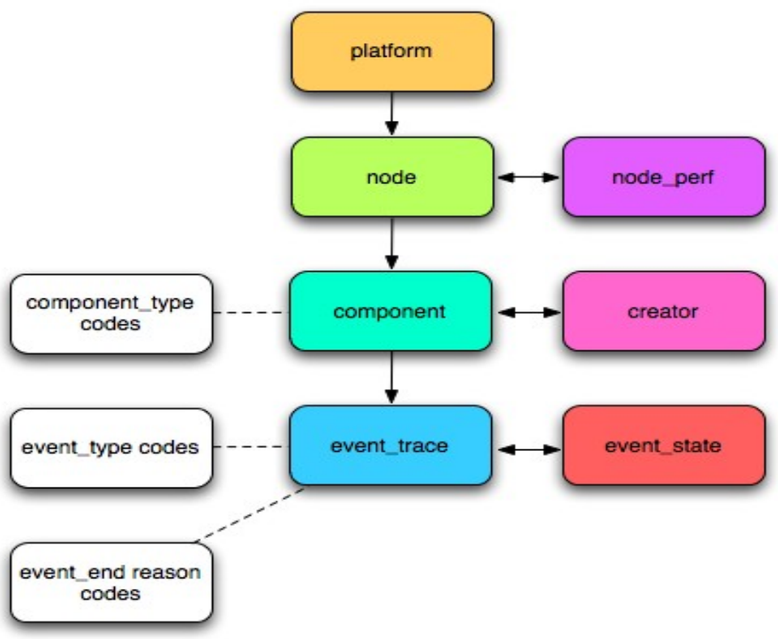

- A **platform** contains a set of nodes. Examples of a platform include SETI@home, desktops at Microsoft.
- A **node** contains a set of components, which is a software module or hardware resource of the node. Each node can have several components (e.g. CPU speed, availability memory, client availability), each of which has a corresponding trace.
- A **component** describes attributes of a software module or hardware resource of a node.
- **component_perf** is the component performance, as measured through benchmarks for example.
- A **creator** is the person responsible for the trace data set. This table stores details about citations and copyright.
- An **event_trace** is the trace of an event, with all of corresponding timing information.
- **event_state** is the state corresponding to an event_trace. For example, for CPU availability, the event_state could be the idleness of the CPU.

http://fta.inria.fr

Data Sets

System	Type	# of Nodes	Target Component	Period	Year
SETI@home	Desktop Grid	226208	CPU	1.5 years	2007-2009
Overnet	P2P	3000	host	2 weeks	2003
Microsoft	Enterprise	51663	host	35 days	1995
LANL	HPC Clusters	3577	host	9 years	1996-2005
HPC2	HPC Clusters	256	IO	2.5 years	1996-2005
HPC4	Supercomputers	152516	every thing	~1 year	2004-2006
PNNL	HPC Cluster	980	CPU, IO, memory	4 years	2003-2007
NERSC	HPC Clusters	NA	IO	5 years	2001-2006
Skype	P2P	4000	host	1 month	2005
Web sites	Web servers	129	host	8 months	2001-2002
DNS	DNS servers	62201	host	2 weeks	2004
PlanetLab	P2P	200-400	host	1.5 year	2004-2005
Grenouille03	DSL	4800	host	1 year	2003
Grenouille05	DSL	4800	host	1 year	2005
EGEE	Grid	2500 queues	CE queue	1 month	2007

- Each trace data set is made available in:
  - raw format
  - tabbed-delimited format
  - MySQL format
- Perl scripts/module are provided for:
  - the FTA format
  - parse the raw data and format it with tabs
  - convert tabbed format into MySQL format

http://fta.inria.fr

Trace Analysis

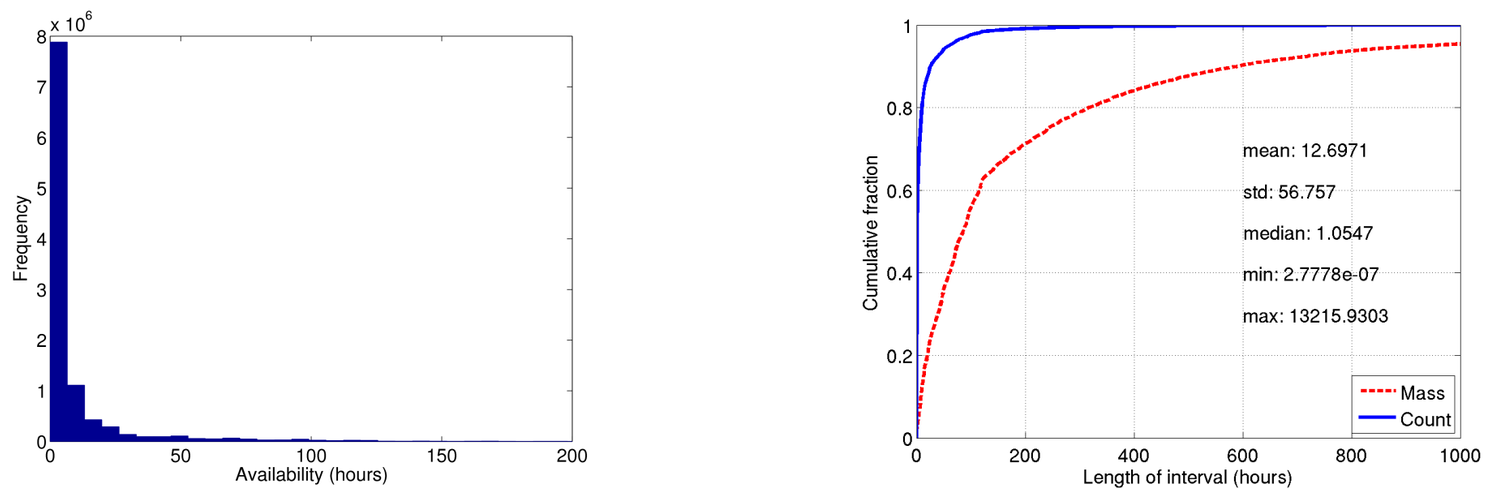

SETI@home data set, (left) histogram (right) mass-count of availability intervals

TU Delft Technische Universiteit Delft · INRIA · Osaka University · ZIB