Event Transfer (ET) System

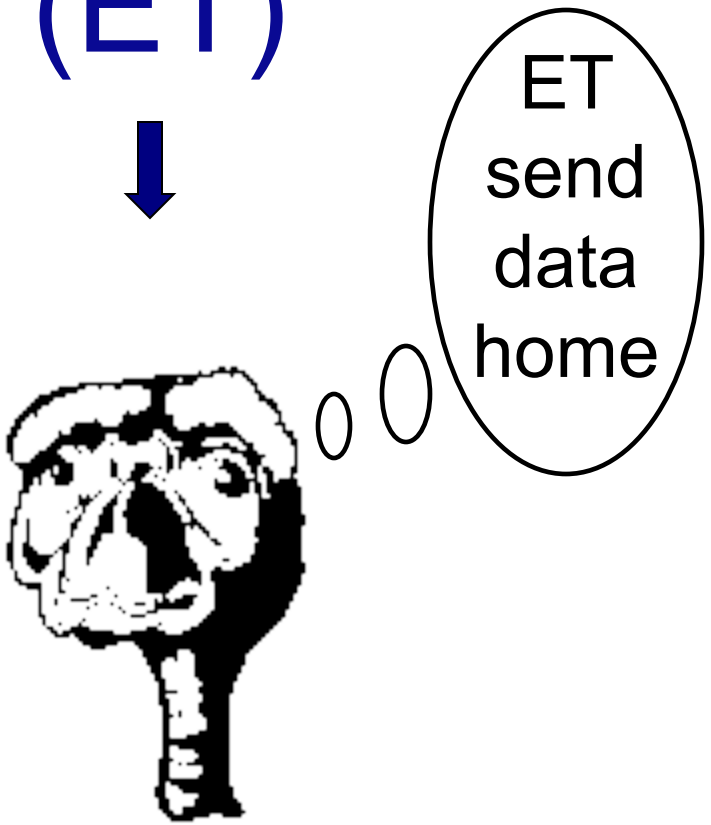

Carl Timmer

Jefferson Lab

What is the ET system?

- Software to move events (data buffers) from process to process
  - One main ET process
    - Stores all events (buffers)
    - Distributes events to users
    - Monitors users
  - Library for users
    - Modify existing ET systems
    - Can even become an ET system
    - Get & put events from ET systems
  - Monitoring programs – text & graphical

ET System Design

- FAST
  - Shared memory
  - Fully multithreaded with POSIX pthreads
  - Written in C on Solaris, Linux, & Mac OS X

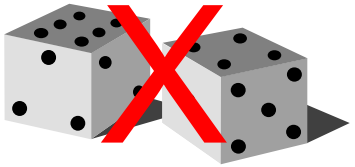

- RELIABLE
  - Complete error recovery for system & users
  - Data from crashed users are recoverable
- FRIENDLY
  - Threads transparent to users
  - Network transparent to users
  - Java version available
  - Reentrant – run as many copies on 1 computer as you like

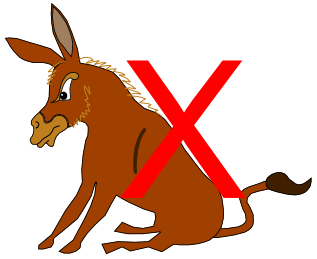

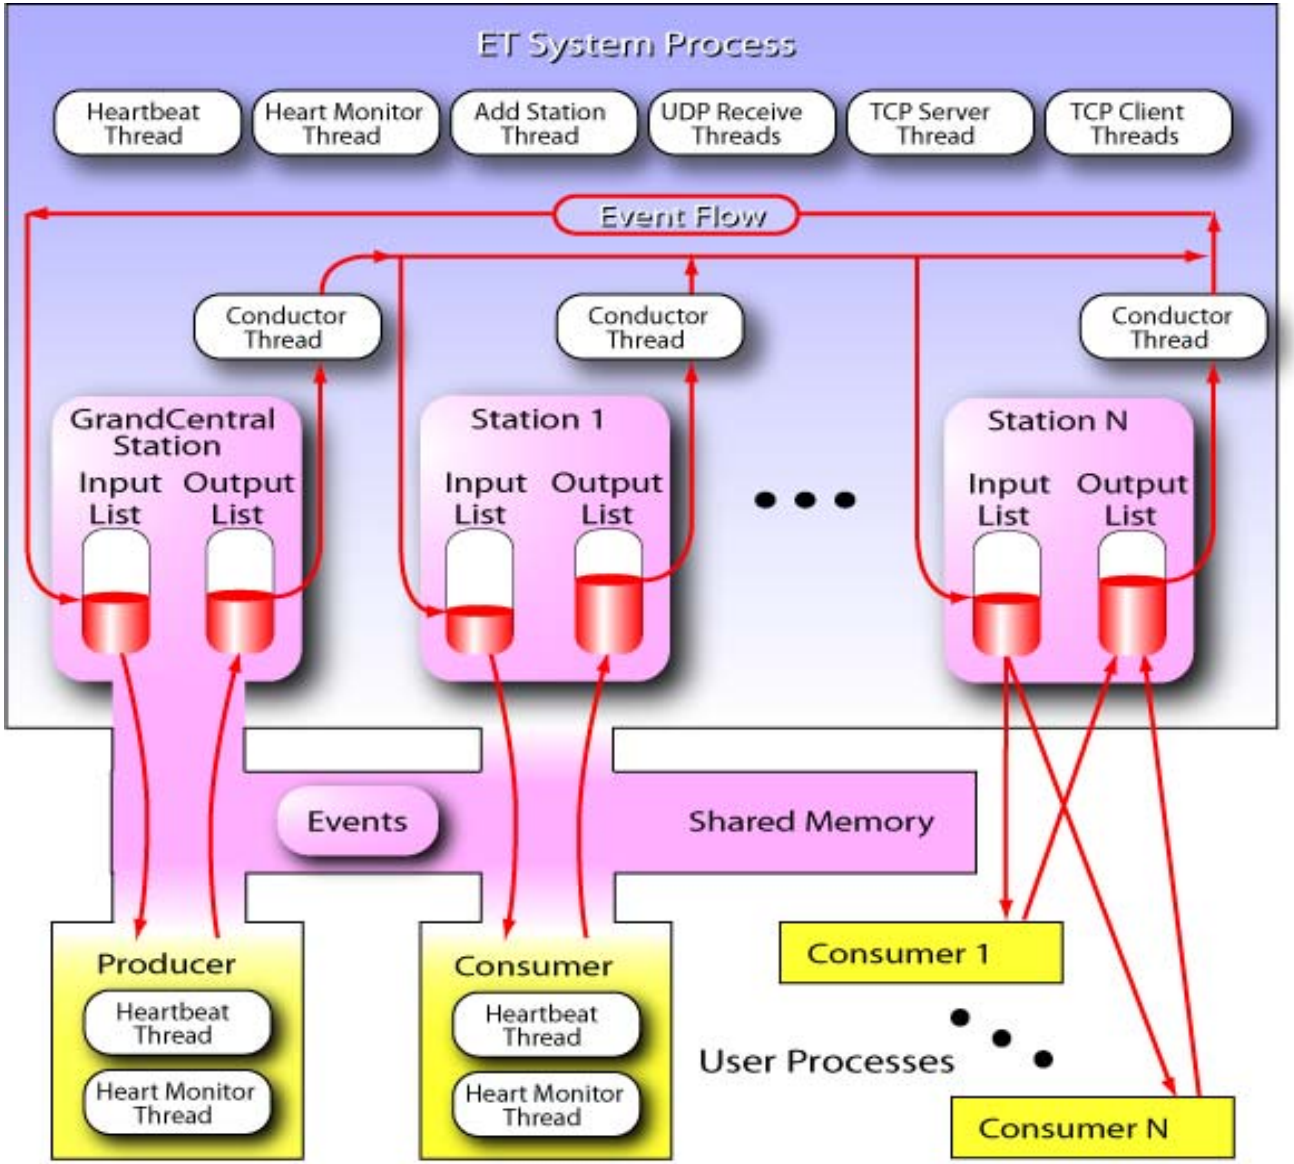

Event Features

- Each event has integer array as metadata
- Events can be any size (within computer's memory limits)
- Keeps track of endian changes
- Users can get and put events in arrays. This increases performance over 10X.
- Data in CODA format can be swapped
- 2 levels of priority

Station Features

- Users can define and add their own stations
- Multiple users can attach to the same station
- Each station can have a user-defined function to select events of interest
- Station order can be changed

  NEW FEATURE
- Stations can be added in parallel with round-robin or load-balancing event distribution algorithms

Parallel Stations

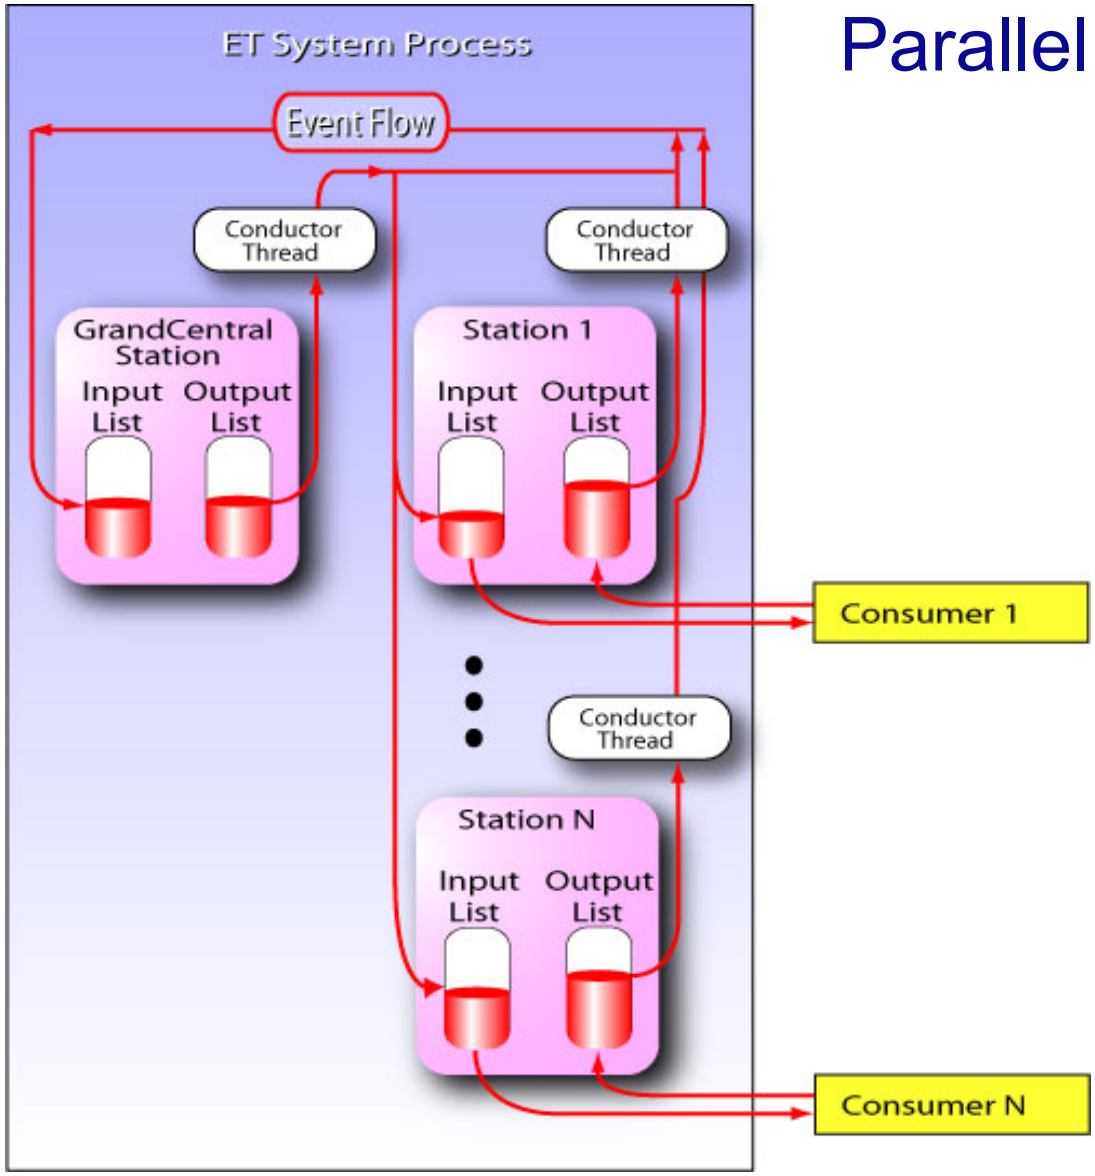

What about performance?

ET System Performance

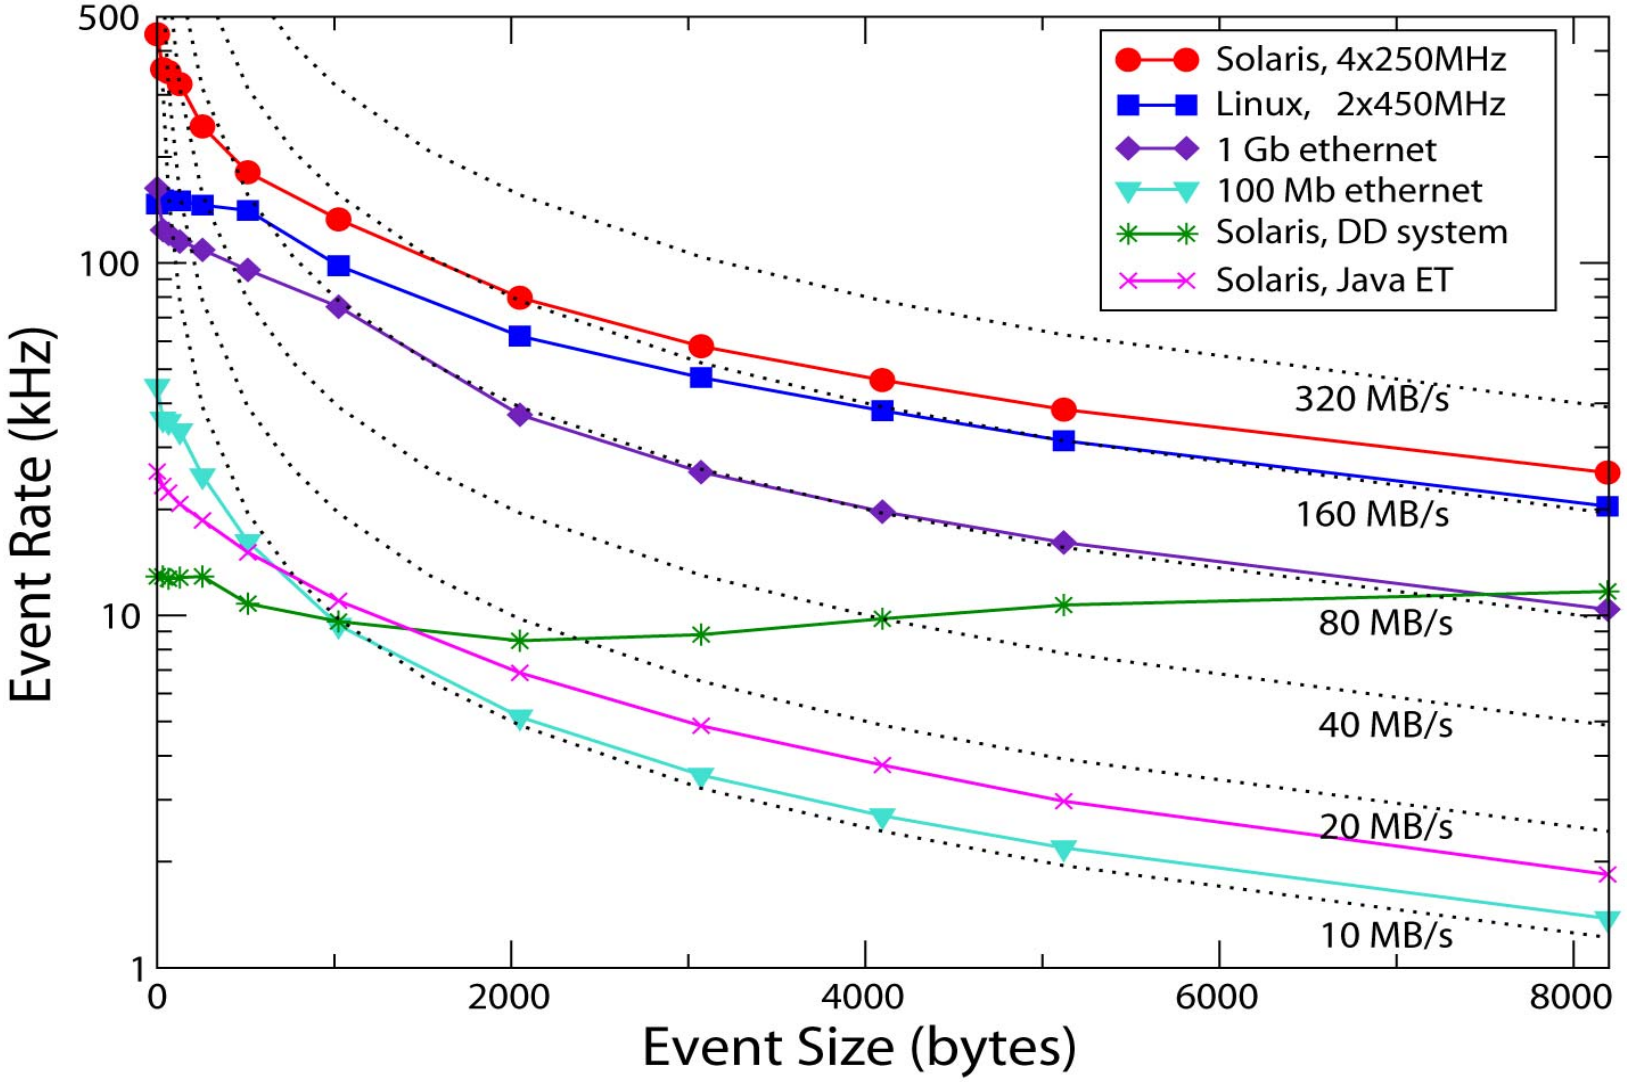

Jefferson Lab

ET System Performance

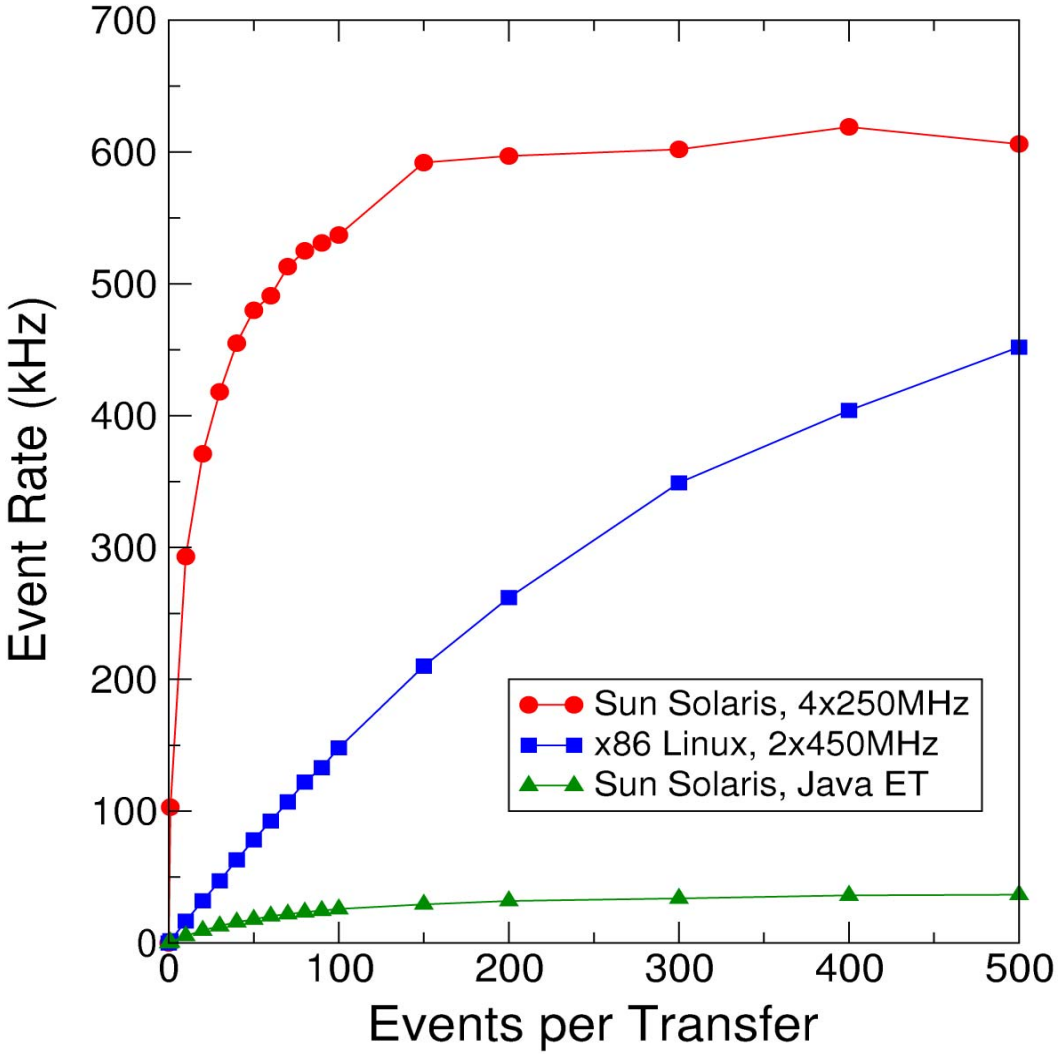

Jefferson Lab

ET graphical Monitor

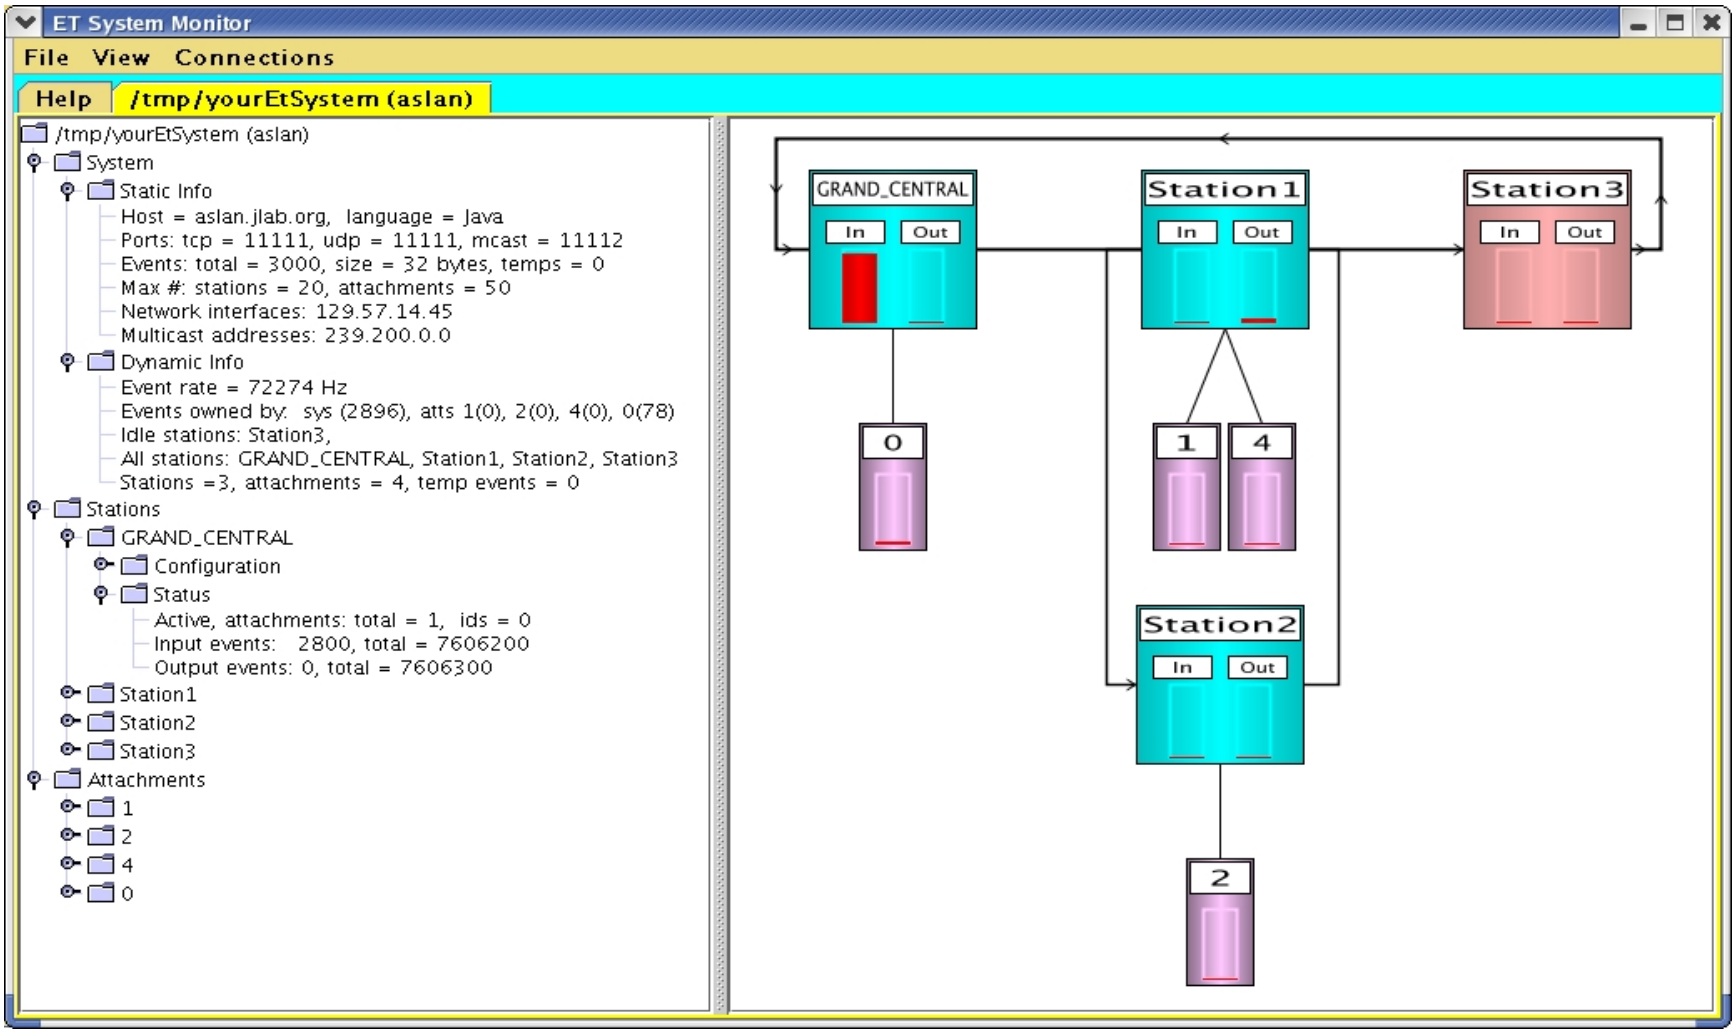

Future Developments

- Rewrite roc, event builder, & event recorder to use ET as communication mechanism
- Run speed tests with and optimize for Gigabit ethernet and TCP Jumbo Frames
- Fix currently broken ET web manual at http://coda.jlab.org
- Bug fixes