UT-BATTELLE

Channel Archiver Stats & Problems

Kay Kasemir, Greg Lawson, Jeff Patton (ORNL)

October 2007

SNS SPALLATION NEUTRON SOURCE

OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY

SNS Stats

- Averaging ~2000 samples/sec, from ~80000 channels
  - 1997 design goal: 10000 samples/sec
  - Raw 'write' test: >50000 samples/sec
- Split into ~70 active sub archives
  - CSS DataBrowser makes that usable
- Disk space: ~170 GB/month
  - Maybe 6 month left

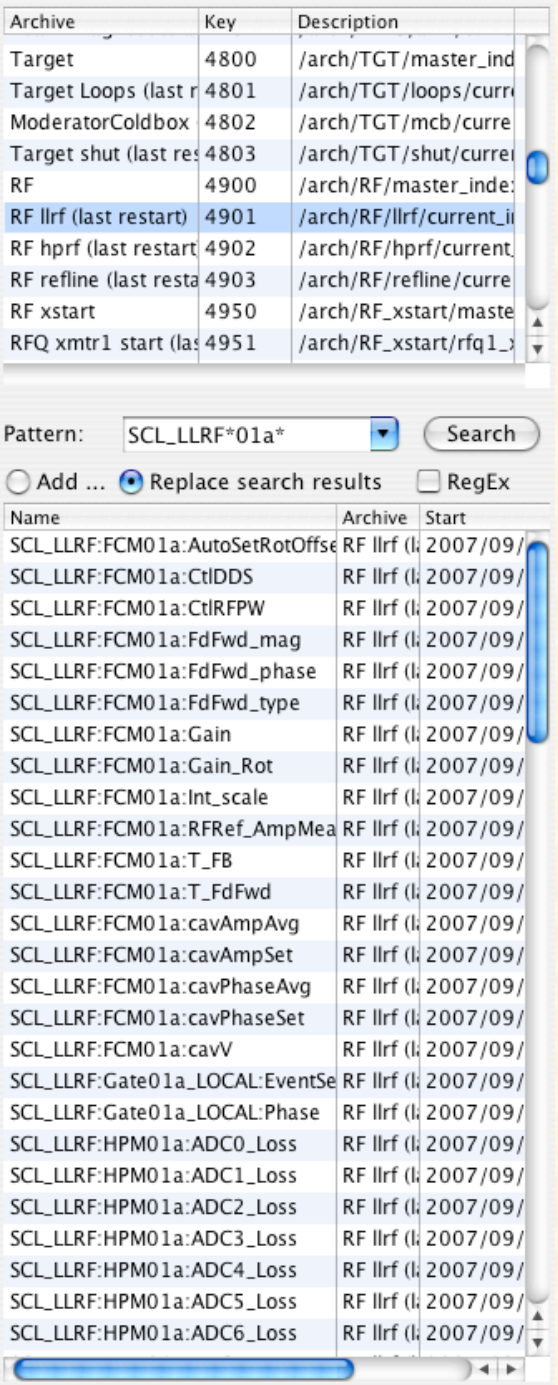

Data Management Limitations

- **Difficult and time consuming**
  - Moving data around requires manual index updates
- **Few Informational Tools**
  - Nothing prevents duplication
  - Which *channels* contribute the most to data growth?
- **Storage only supports "Append new samples"**
  - Removal of selected channels impossible
  - Removal of older data limited to complete 'sub archives'
  - No practical way to use Java or Matlab code to replace original samples with reduced sample count,
  - .. or to insert computed data like daily statistics into "archive"

JLab: MySQL Transition

- Very promising performance tests!
- Limitations by design
  - Stores every update from IOC. No 'sampling'.
  - 'Double' stored as 'float'.
  - Only small arrays.
  - Metadata: Units. No limits, precision. No status/severity.
- MySQL Issues
  - Table size limited
    - Need one table per channel
  - Table count limited
    - Custom code implements 'clustering'
  - SQL "DELETE" doesn't free disk space or is very slow

SNS: Recent Oracle Tests

- Basic JDBC test code:
  up to 8000 inserts/second via network
- Tricks
  - "Batching" ~500 inserts
  - "Partitioning" spreads one big "sample" table over disk partitions
    - Currently one partition each day, automatically added
- Expensive, but looks like the way to go
  - Avoid MySQL workarounds
  - SNS committed to Oracle anyway,
    but will need partitioning license

Great! But what about SLAC?

- Reported promising performance results for Oracle-based data storage
  - Lee Ann Yasukawa, Robert Hall: *"Archiving Into Oracle"*, ICALEPCS2001
- End of 2004: No more.
  - What do we need to learn from that?

SNS Plan

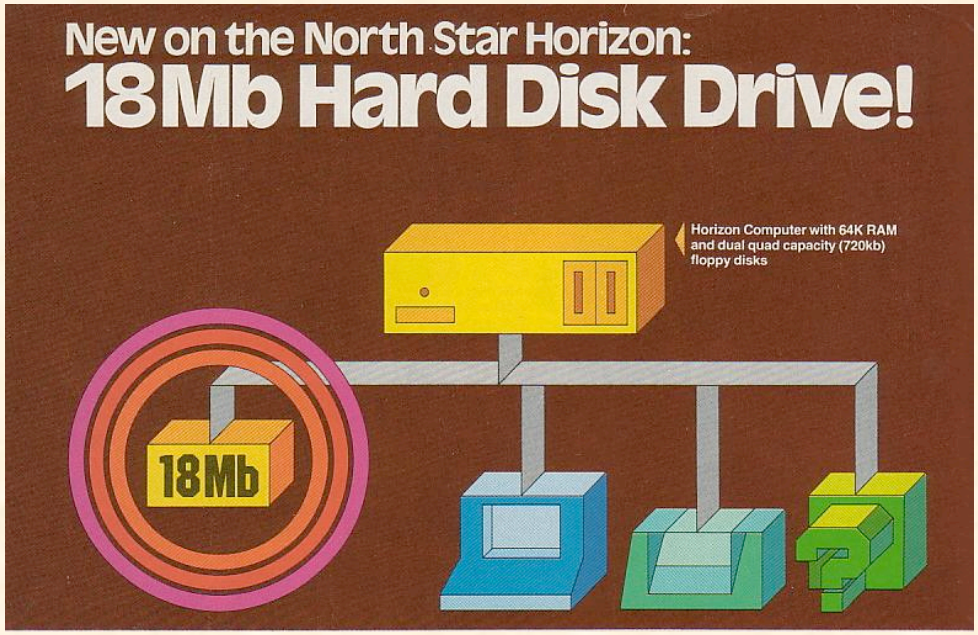

- Buy HP EVA storage array and 30TB disks
  - Shared: Computing Integration group, Controls group
- Can be used for anything
  - Add to current archive server??
  - Add to current Oracle server?
  - Use with designated Oracle server!

Archive Engine Prototype

- Developed in Java
  - Eclipse/CSS command-line app
- Reads existing engine config files
  - OK with LLRF, RCCS, new "permanent" BLM setups
- Writes into Oracle
  - From office and controls network
  - … and MySQL for sites that don't have much data
- Write performance OK for scalar tests
- Current issue: Channel Access connections

Engine's Web Server

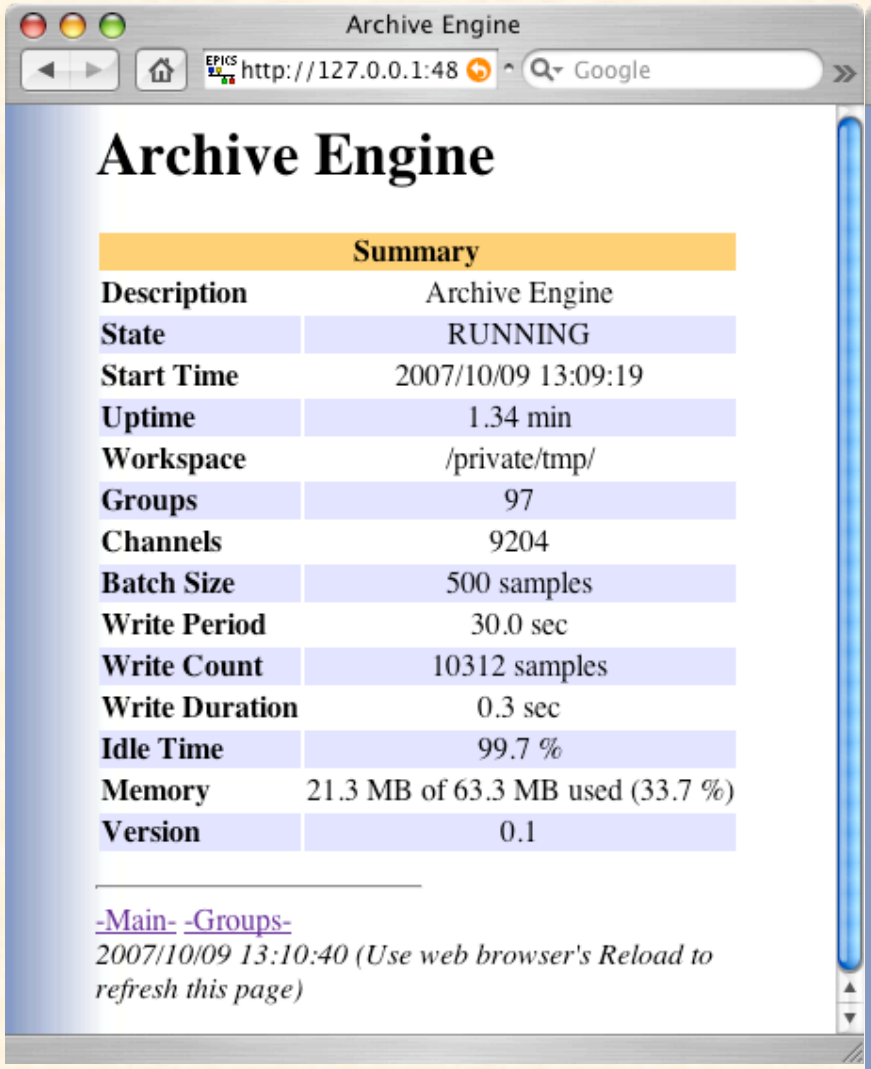

Basic Sample Table Design

- What data types to support?
  - Time stamp detail, enumerated values, arrays, meta data?
- One table per channel (JLab)?
- One table per data type (SLAC)?
- One table for all samples (SNS)?
  - Possibly wasting space, but best to use SQL across various channels of different types

Summary

- Investigating Oracle as archive data storage
- Currently working with Oracle to obtain a quote to add partitioning to license
- Testing Oracle setups, prototyping sampling engine
- Performance expected to be almost comparable to current SNS ChannelArchiver setup
- … but sustainable in the long run.