

Performance tests of Hypertable



**Mauro Giacchini, Robert Petkus
Group Meeting
10/27, 2009**

Target

- Make the EPICS Archiver embedded DB more reliable and safe especially in large application with a large number of PVs to store
- First Approach: re-write the Rtree and double-linked-list embedded DB structure
- Second approach: look for a complete replacement of the embedded DB with new one

Why not Beauty ? (Best Ever Archiver Utility)

Beauty = Classical Archiver over Oracle DB

Pros:

- SQL statements accept
- Professional Solution

Cons:

- Not so fast
- Really expensive (licenses)



Why not MySQL ?

- Potential scalability concerns
- Designed for a single machine (not distributed)
- What it takes to make it scale
 - Major engineering effort
 - Solutions are usually ad hoc
 - Solutions usually involve horizontal partitioning + replication
 - Solutions involve expensive hardware




BigTable ?

- Proprietary DB used by Google over their proprietary Google File System
- High Performance
- Large amount of data
- Stable
- Reliable
- Distributed

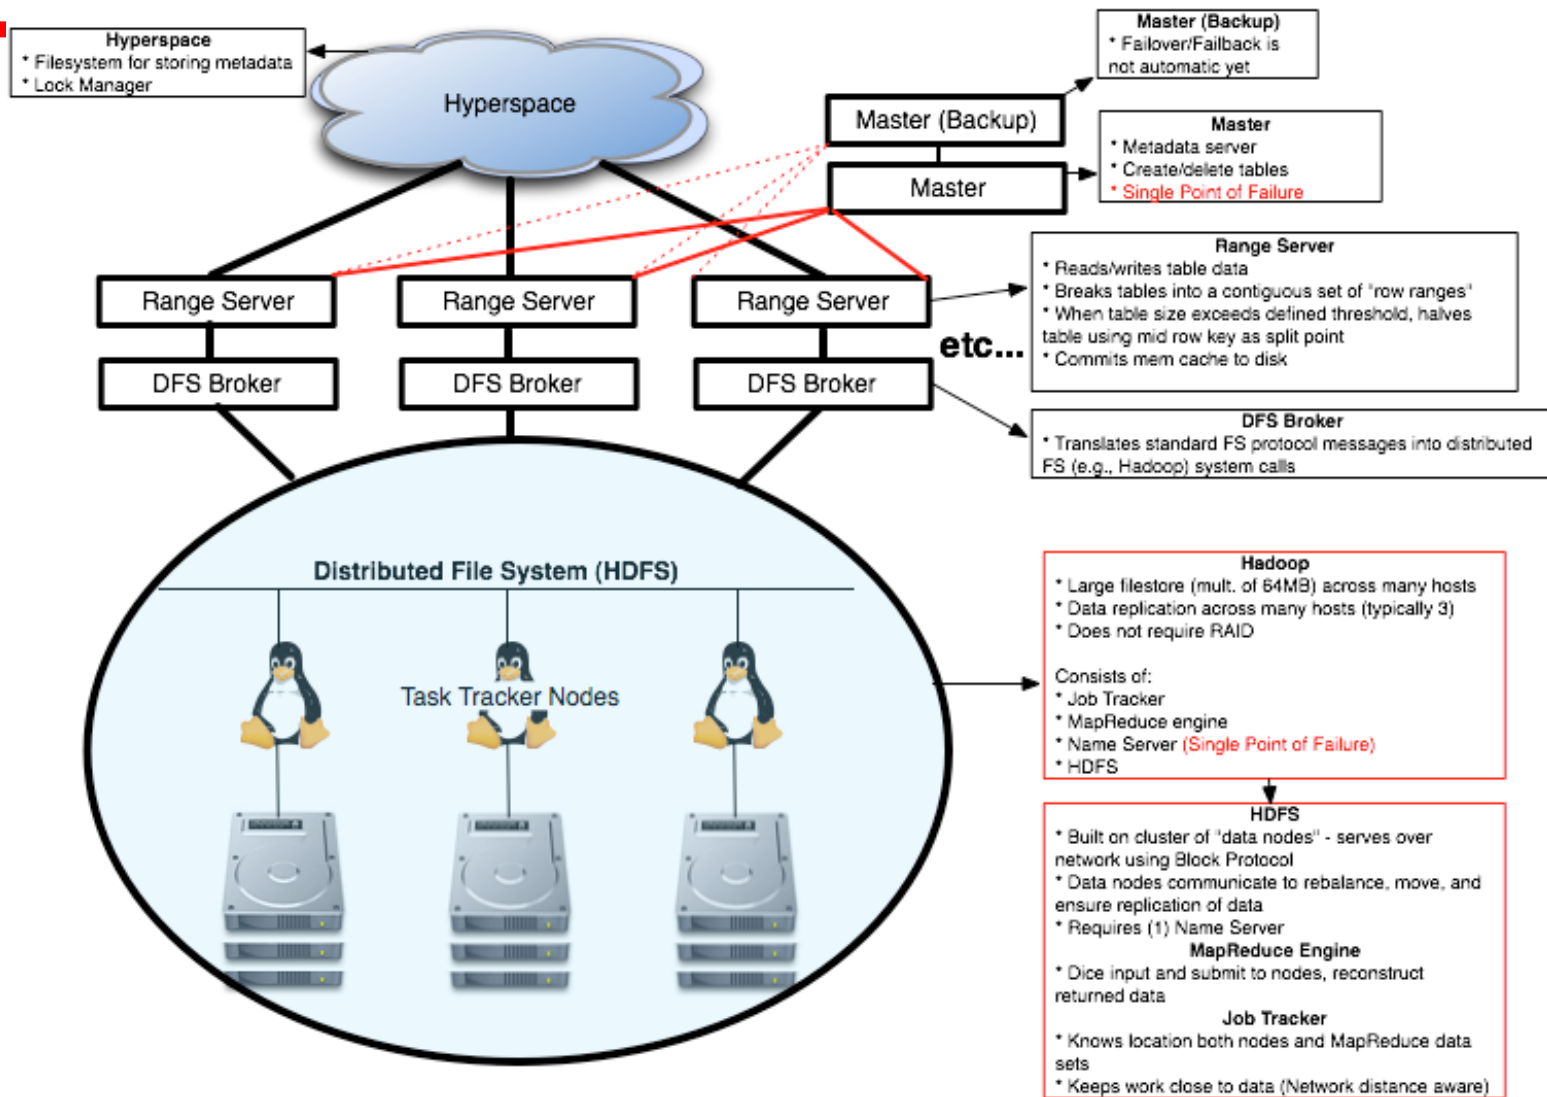



HYPERTABLE...what is it ?


- Hypertable has been developed as an in-house software at Zvents Inc. In January 2009, [Baidu](#), the leading Chinese language search engine, became a project sponsor. Rediff.com is one of the premier worldwide online providers of news, information, communication, entertainment and shopping service.
- Hypertable is licensed under the GNU General Public License Version 2.
- “Our goal is nothing less than that Hypertable become one of the world’s most massively parallel high performance database platforms.”
- Not RDB



Hadoop/Hypertable Architecture

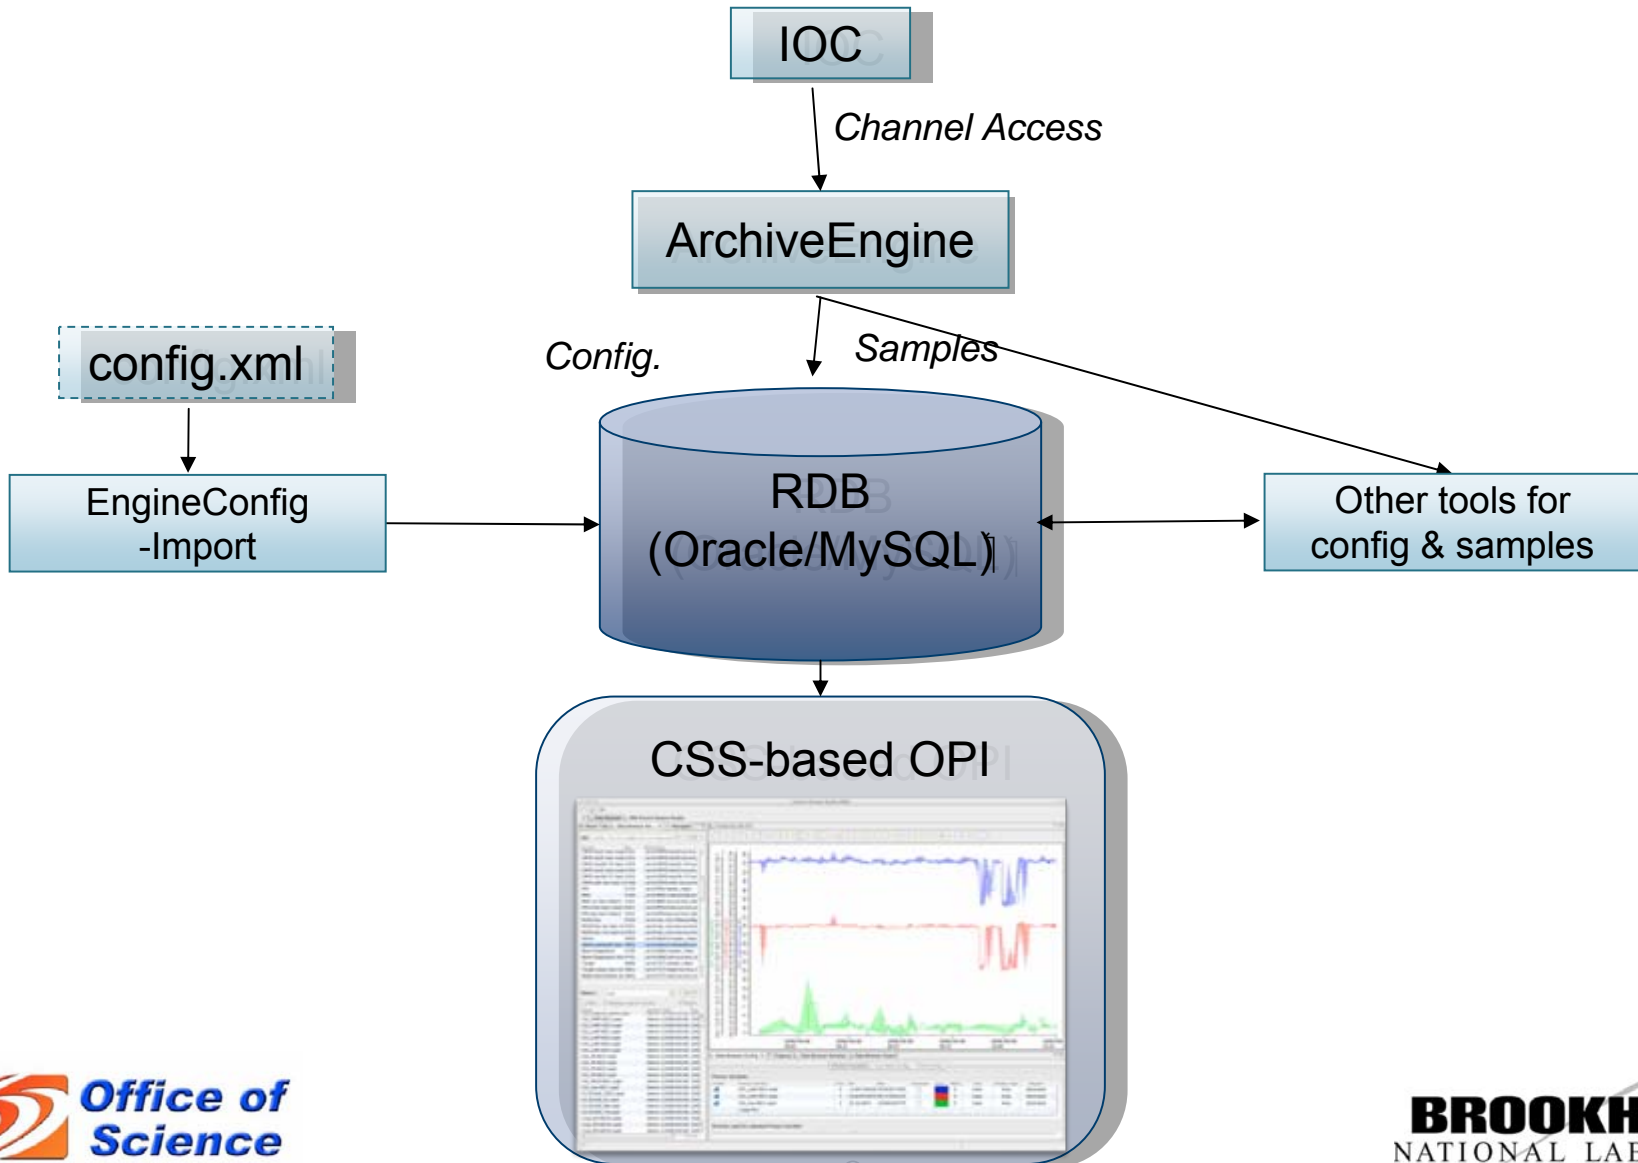


Hadoop/Hypertable Solution Objective

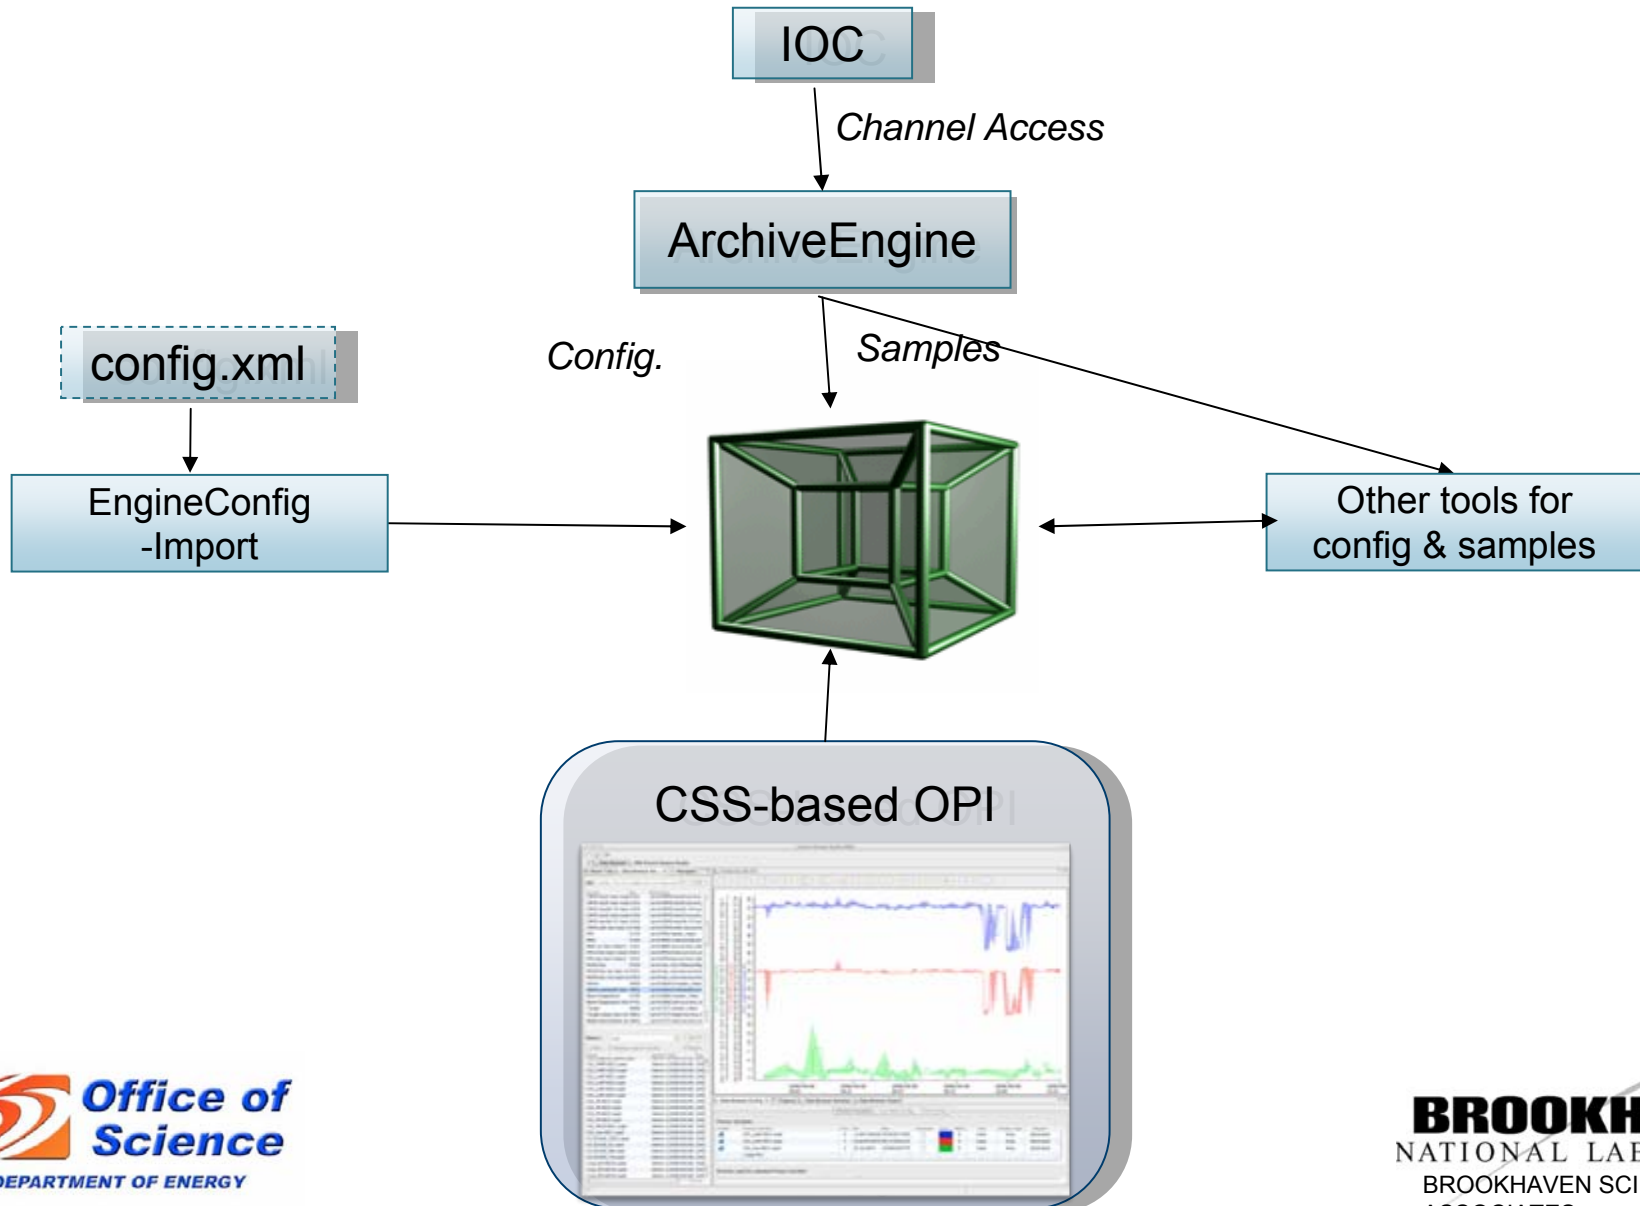


- Handles applications with large datasets
- Fast fault detection and fail-over
- High throughput processing
- Centralized scheduling of tasks and execution of batch processes
- Can be deployed on low cost, commodity hardware
- Eliminates or reduces the need for expensive table joins

Beauty

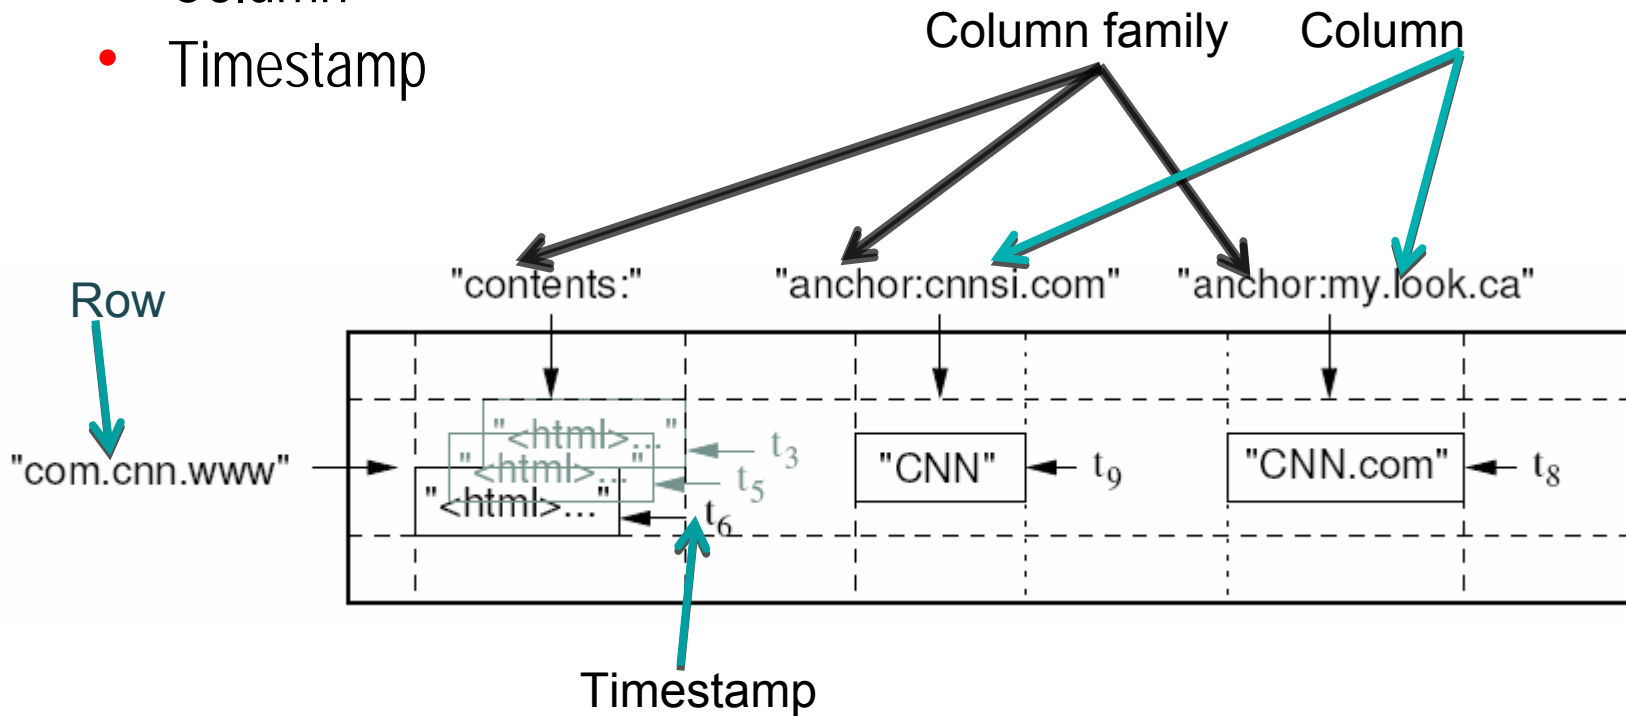


Archiver w Hypertable ..?



Data model in Hypertable

- Four dimension table:
 - Row
 - Column Family
 - Column
 - Timestamp



Archiver table on Hypertable

| PV | TS | VAL | STAT | SEVR |
|----|----|-----|------|------|
| | | | | |
| | | | | |

```
<Schema>
  <AccessGroup name="default">
    <ColumnFamily>
      <Name>pv</Name>
      <deleted>>false</deleted>
    </ColumnFamily>
    <ColumnFamily>
      <Name>ts</Name>
      <deleted>>false</deleted>
    </ColumnFamily>
    <ColumnFamily>
      <Name>val</Name>
      <deleted>>false</deleted>
    </ColumnFamily>
    <ColumnFamily>
      <Name>stat</Name>
      <deleted>>false</deleted>
    </ColumnFamily>
    <ColumnFamily>
      <Name>sevr</Name>
      <deleted>>false</deleted>
    </ColumnFamily>
  </AccessGroup>
</Schema>
```

Machines used to the tests

- Virtualbox virtual machine
- Debian Linux Stable x86
- 4 cores (shared among 3 VMs)
- 4GB RAM
- (1) 5.4k RPM SATA Disk
- \$4k (to 4 Vbox)



- Debian Linux Stable AMD64
- 8 cores
- 64GB RAM
- (7) 10k RPM SAS Disks (RAID 5)
- \$11k



Hyp Performance ½ hour



- Insertion Time Test (1/2 h data acq SNS)
- 1min 31 sec to insert
- 1 777 600 samples
- 19.5K samples/sec

Insertion data amount block:
 $4.8\text{K} + 49\text{K} + 496\text{K} + 5.2\text{M} = 5.75\text{M}$

- Extraction Time Test (all pv parameters)
- 0.042 sec to extract
- 133K samples of 4 records
- 3.17M samples/sec

Hyp Performance ½ hour

8

- Insertion Time Test (1/2 h data acq SNS)
 - 1min 13.6 sec to insert
 - 1 777 600 samples
 - 24K samples/sec
-
- Extraction Time Test (all pv parameters)
 - 0.024 sec to extract
 - 133K samples of 4 records
 - 5.54 Msamples/sec

Hyp Performance on a week

8

- Insertion Time Test (7dy data acq SNS)
- 11.2 hours to insert
- 597M samples
- 14.8K samples/sec
- Extraction Time Test (all parameters)
- 0.031 sec to extract
- 133360 samples of 4 records
- 4.3M samples/sec
- Extraction Time Test (all pv)
- 19min 51 sec
- 597.273.600 pv
- 500K samples/sec

Hyp Performance on a month

8

- Insertion Time Test (30dy data acq SNS)
- 37.86 hours to insert
- 2557M samples
- 18.7K samples/sec

Insertion data amount block:

$$4.8\text{K} + 49\text{K} + 496\text{K} + 5.2\text{M} = 5.75\text{M} * 1440 = 8.3\text{G}$$

- Extraction Time Test (all parameters)
- 0.030 sec to extract
- 133360 samples of 4 records
- 4.3M samples/sec

Numbers

- Oracle insertion time : 3.746.847 samples per hour (1Ksamples/sec)
- Hypertable insertion time: 1.777.600 per 1min 13 sec (24Ksamples/sec)

- Oracle retrieve time: 36.657 samples in 79.294 sec (~462 vals/sec)
- Hypertable retrieve time: 57624 samples in 0.36 sec (160K vals/sec) in half hour