Stanford Linear Accelerator Center

Oracle Storage for the Channel Archiver

Managing Channel Archiver data with Oracle partitions Overview

Topics

- Goals
- Oracle table design for Channel Archiver data
- Partition management algorithms
- Partition compaction algorithms
- Oracle support tables
- Oracle views to retrieve data
- Status of implementation

Goals

- Increase data storage and retrieval performance
- Use a commercial RDB with its associated data management tools
- Support the current functionality of the Channel Archiver
- Allow flexibility for each site to manage their data their own way.

Table Structure Overview

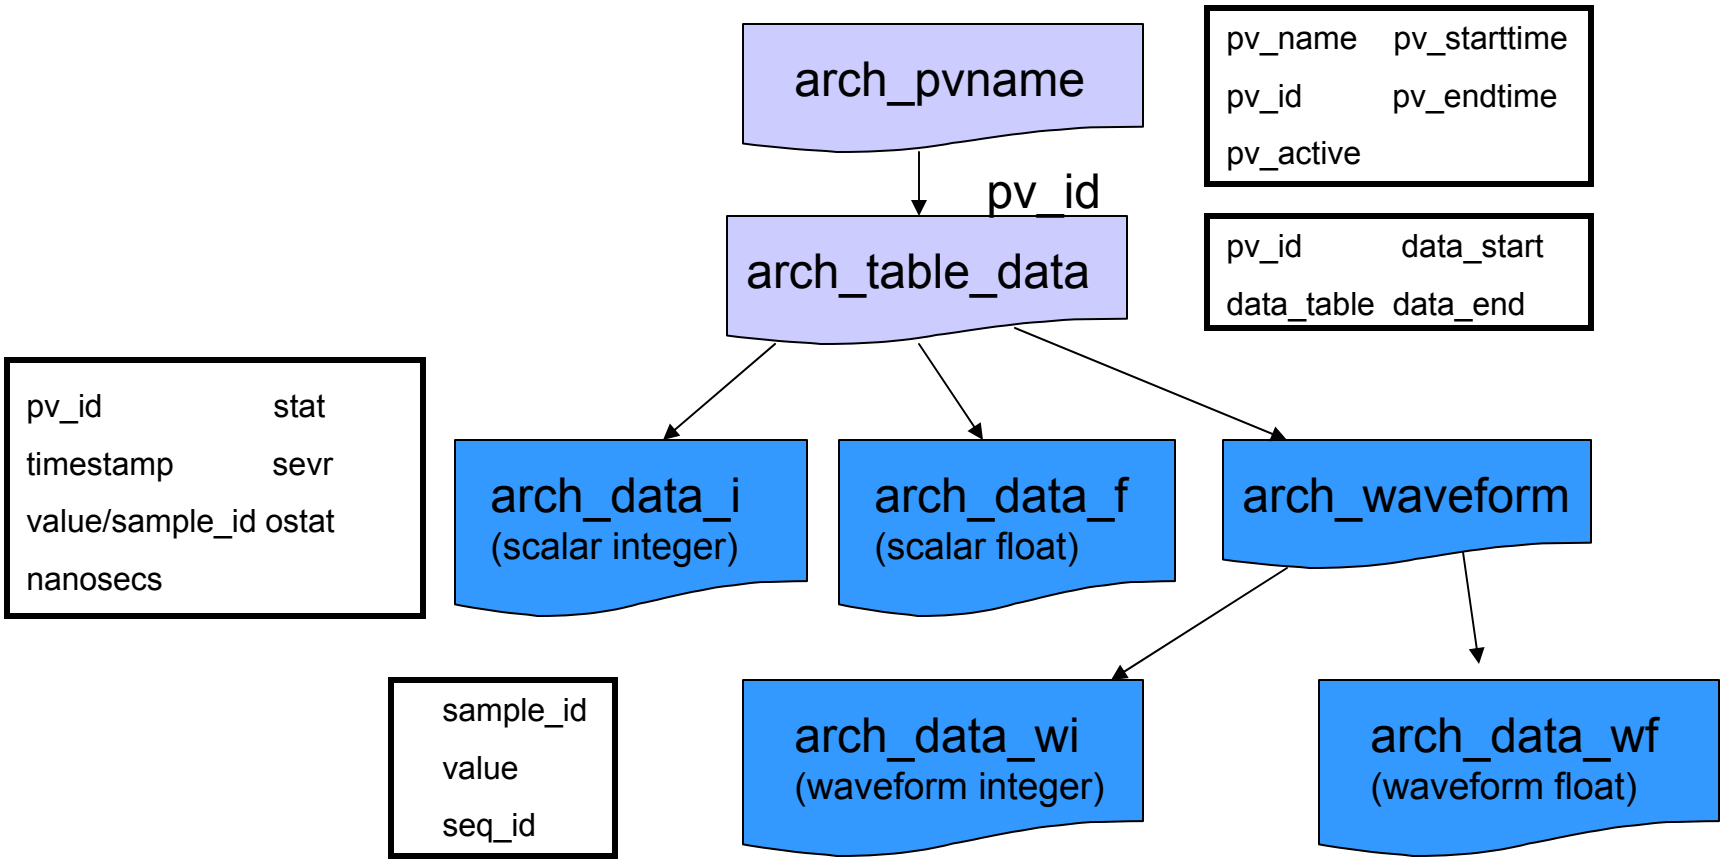

Partitioning Syntax

```
CREATE TABLE ARCH_DATA_I

(     timestamp              date,
      pv_id                  number(38),
      value                  number(38),
      nanosecs               number(9),
      stat                   number(8),
      sevr                   number(8),
      ostat                  number(16) )

  PARTITION BY RANGE (timestamp)

      ( partition MAY0702_0001 values less than
          (TO_DATE('05/07/2002 00:10:00','mm/dd/yyyy hh24:mi:ss')),

        partition MAY0702_0002 values less than
          (TO_DATE('05/07/2002 00:20:00','mm/dd/yyyy hh24:mi:ss')),

        partition bin values less than (MAXVALUE) );
```

Partitioning of Oracle Tables

The arch_data_f, arch_data_i and arch_waveform tables will be partitioned into small (~10 minutes but can be specified) time intervals for the day. These tables are NOT indexed.

EX:

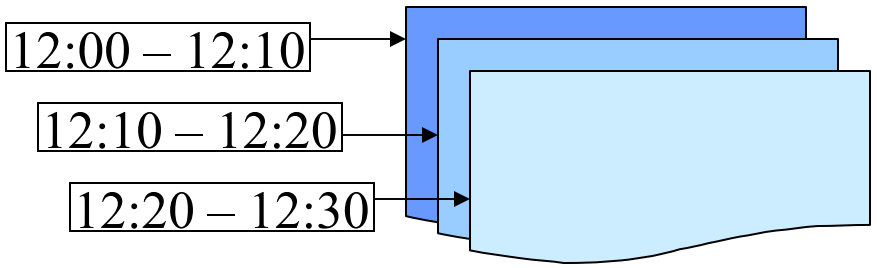

Oracle will track which partition to store the data in so no additional overhead is performed by the Archive Engine.

Daily Processing of Partitions

(Scalar Data – similar processing for waveforms)

arch_data_f arch_data_i	temp_data_f temp_data_i	archive_data_f archive_data_i

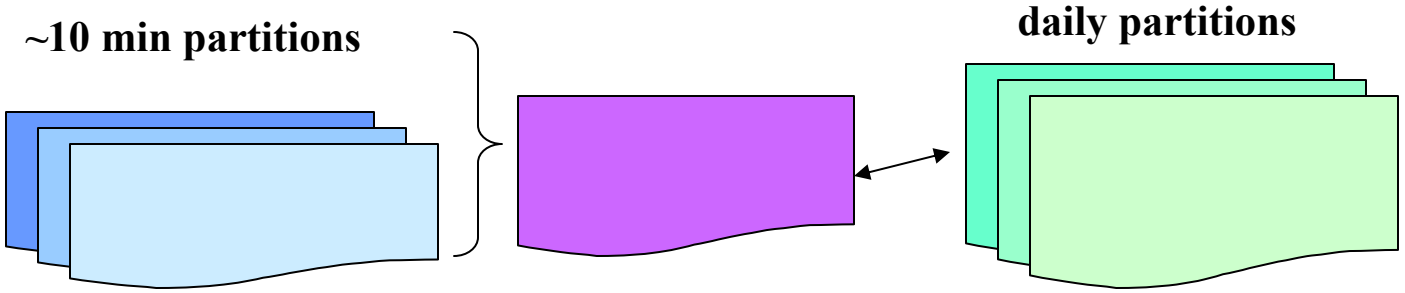

All partitions for yesterday are copied to a temporary table

The data is validated and indexes are created

The temporary table, along with indexes, is exchanged with a partition in the archived data tables

Partition Compaction Algorithm

- Each night a partition compaction script will run which processes archive_XXX tables and associated indexes.
- The compaction algorithm uses the arch_part_durations table to determine the way in which partitions will be compacted.
- Eventually, it will also be used to handle the "rolling out" of partitions from the current location into a near-storage device.

Oracle Partition Compaction

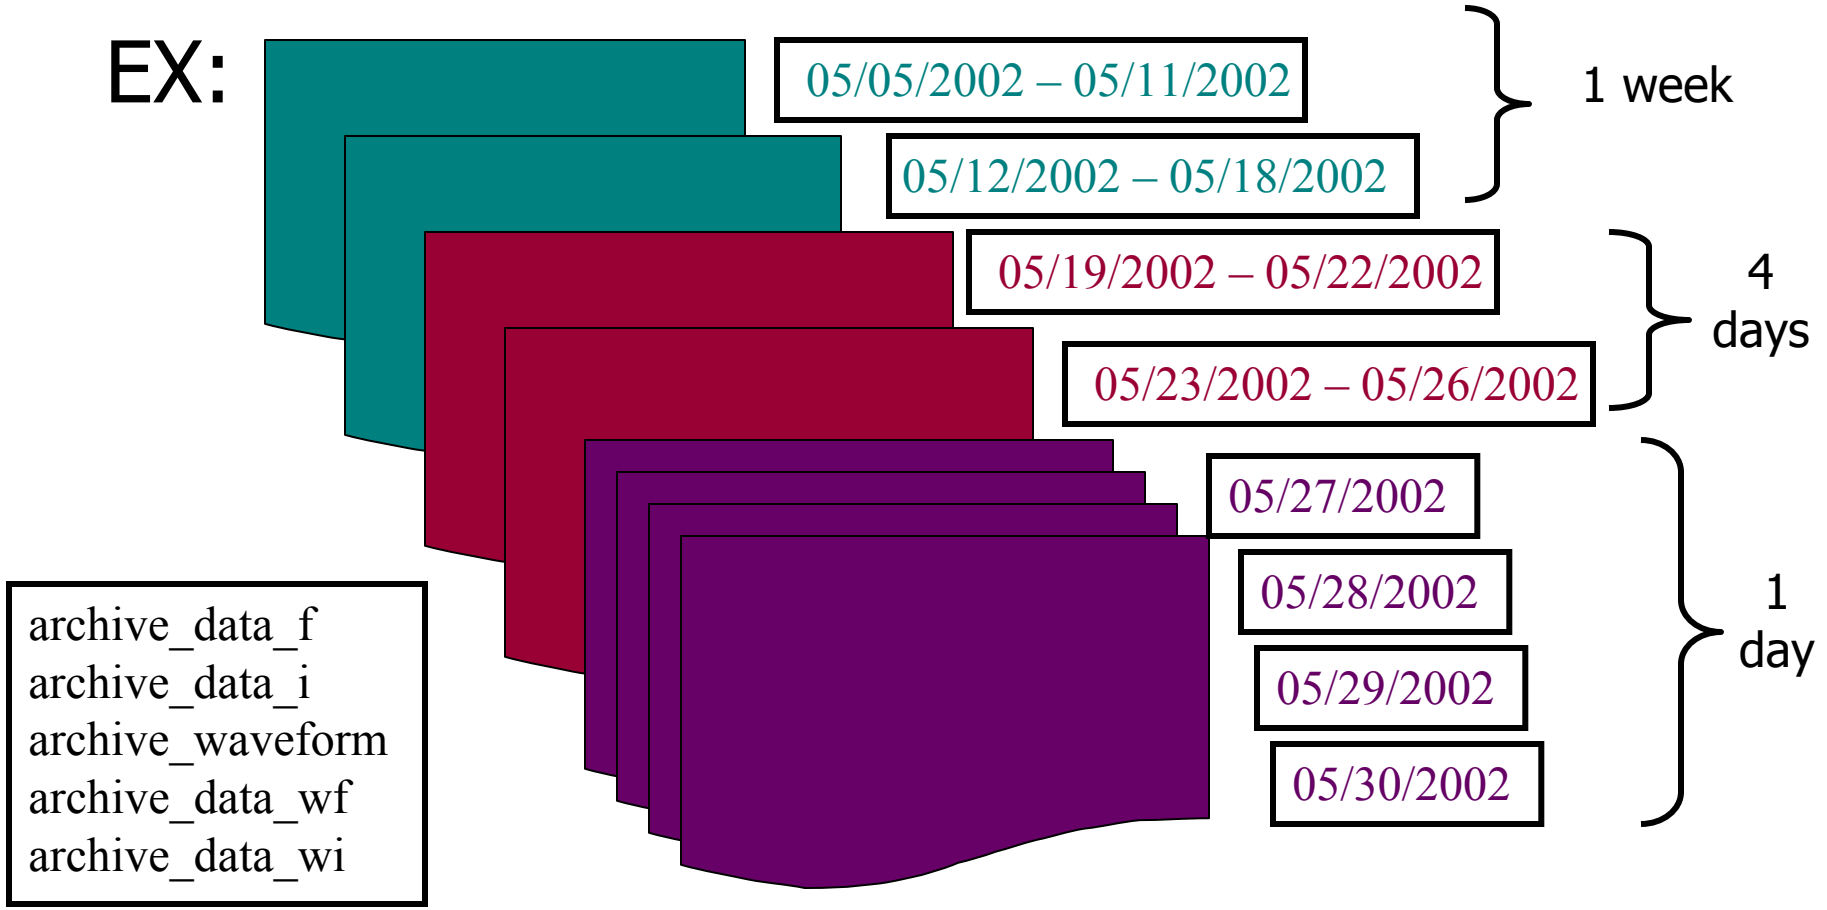

Oracle Views for Retrieval

- Since there will be two tables for each data type, an Oracle view will be created to retrieve data from both tables for each data type
- The retrieval SQL will query the view instead of querying the tables directly.
- The views will be created read-only
- Views allow flexibility as to what data the user has access to
- Views allow access to scalar and waveform data to be the same

Status of Implementation

- We have an Oracle machine available but we currently don't have enough disk storage available for long term archiving.
- The OCI interface to Oracle is defined
- The support scripts are written but some of the partition management functions are still being worked out
- We have a limited number of licences for the Oracle partitioning option
- We are waiting on the changes to the Channel Archiver LibIO code for the integration and test phase to begin.

Additional Notes

- We have tried to keep most of the processing flexible so other labs can use it "out of the box".
- Other labs may use bits and pieces of the Oracle table processing algorithms and are not required to handle their data the same way we plan on handling our data at SLAC. The only hard and fast requirement is for the *initial* table structure to be the same so the Channel Archiver knows where to store the data.
- We are open to any suggestions and ideas for improvement.