FLOW OF MONITORING DATA THROUGH THE MeerKAT TELESCOPE CONTROL SOFTWARE

M. Slabber, M.T. Ockards
SKA SA, Cape Town, South Africa

ICALEPCS 2015
WEPGF065

The MeerKAT telescope, under construction in South Africa, is comprised of a large set of elements. The elements expose various sensors to the Control and Monitoring (CAM) system and the sampling strategy set by CAM per sensor varies from several samples a second to infrequent updates. This creates a substantial volume of sensor data that needs to be stored and made available for analysis. The flow of sensor data through the CAM system, shows the various memory buffers, temporary disk storage and mechanisms to permanently store the data in HDF5 format on the network attached storage (NAS).

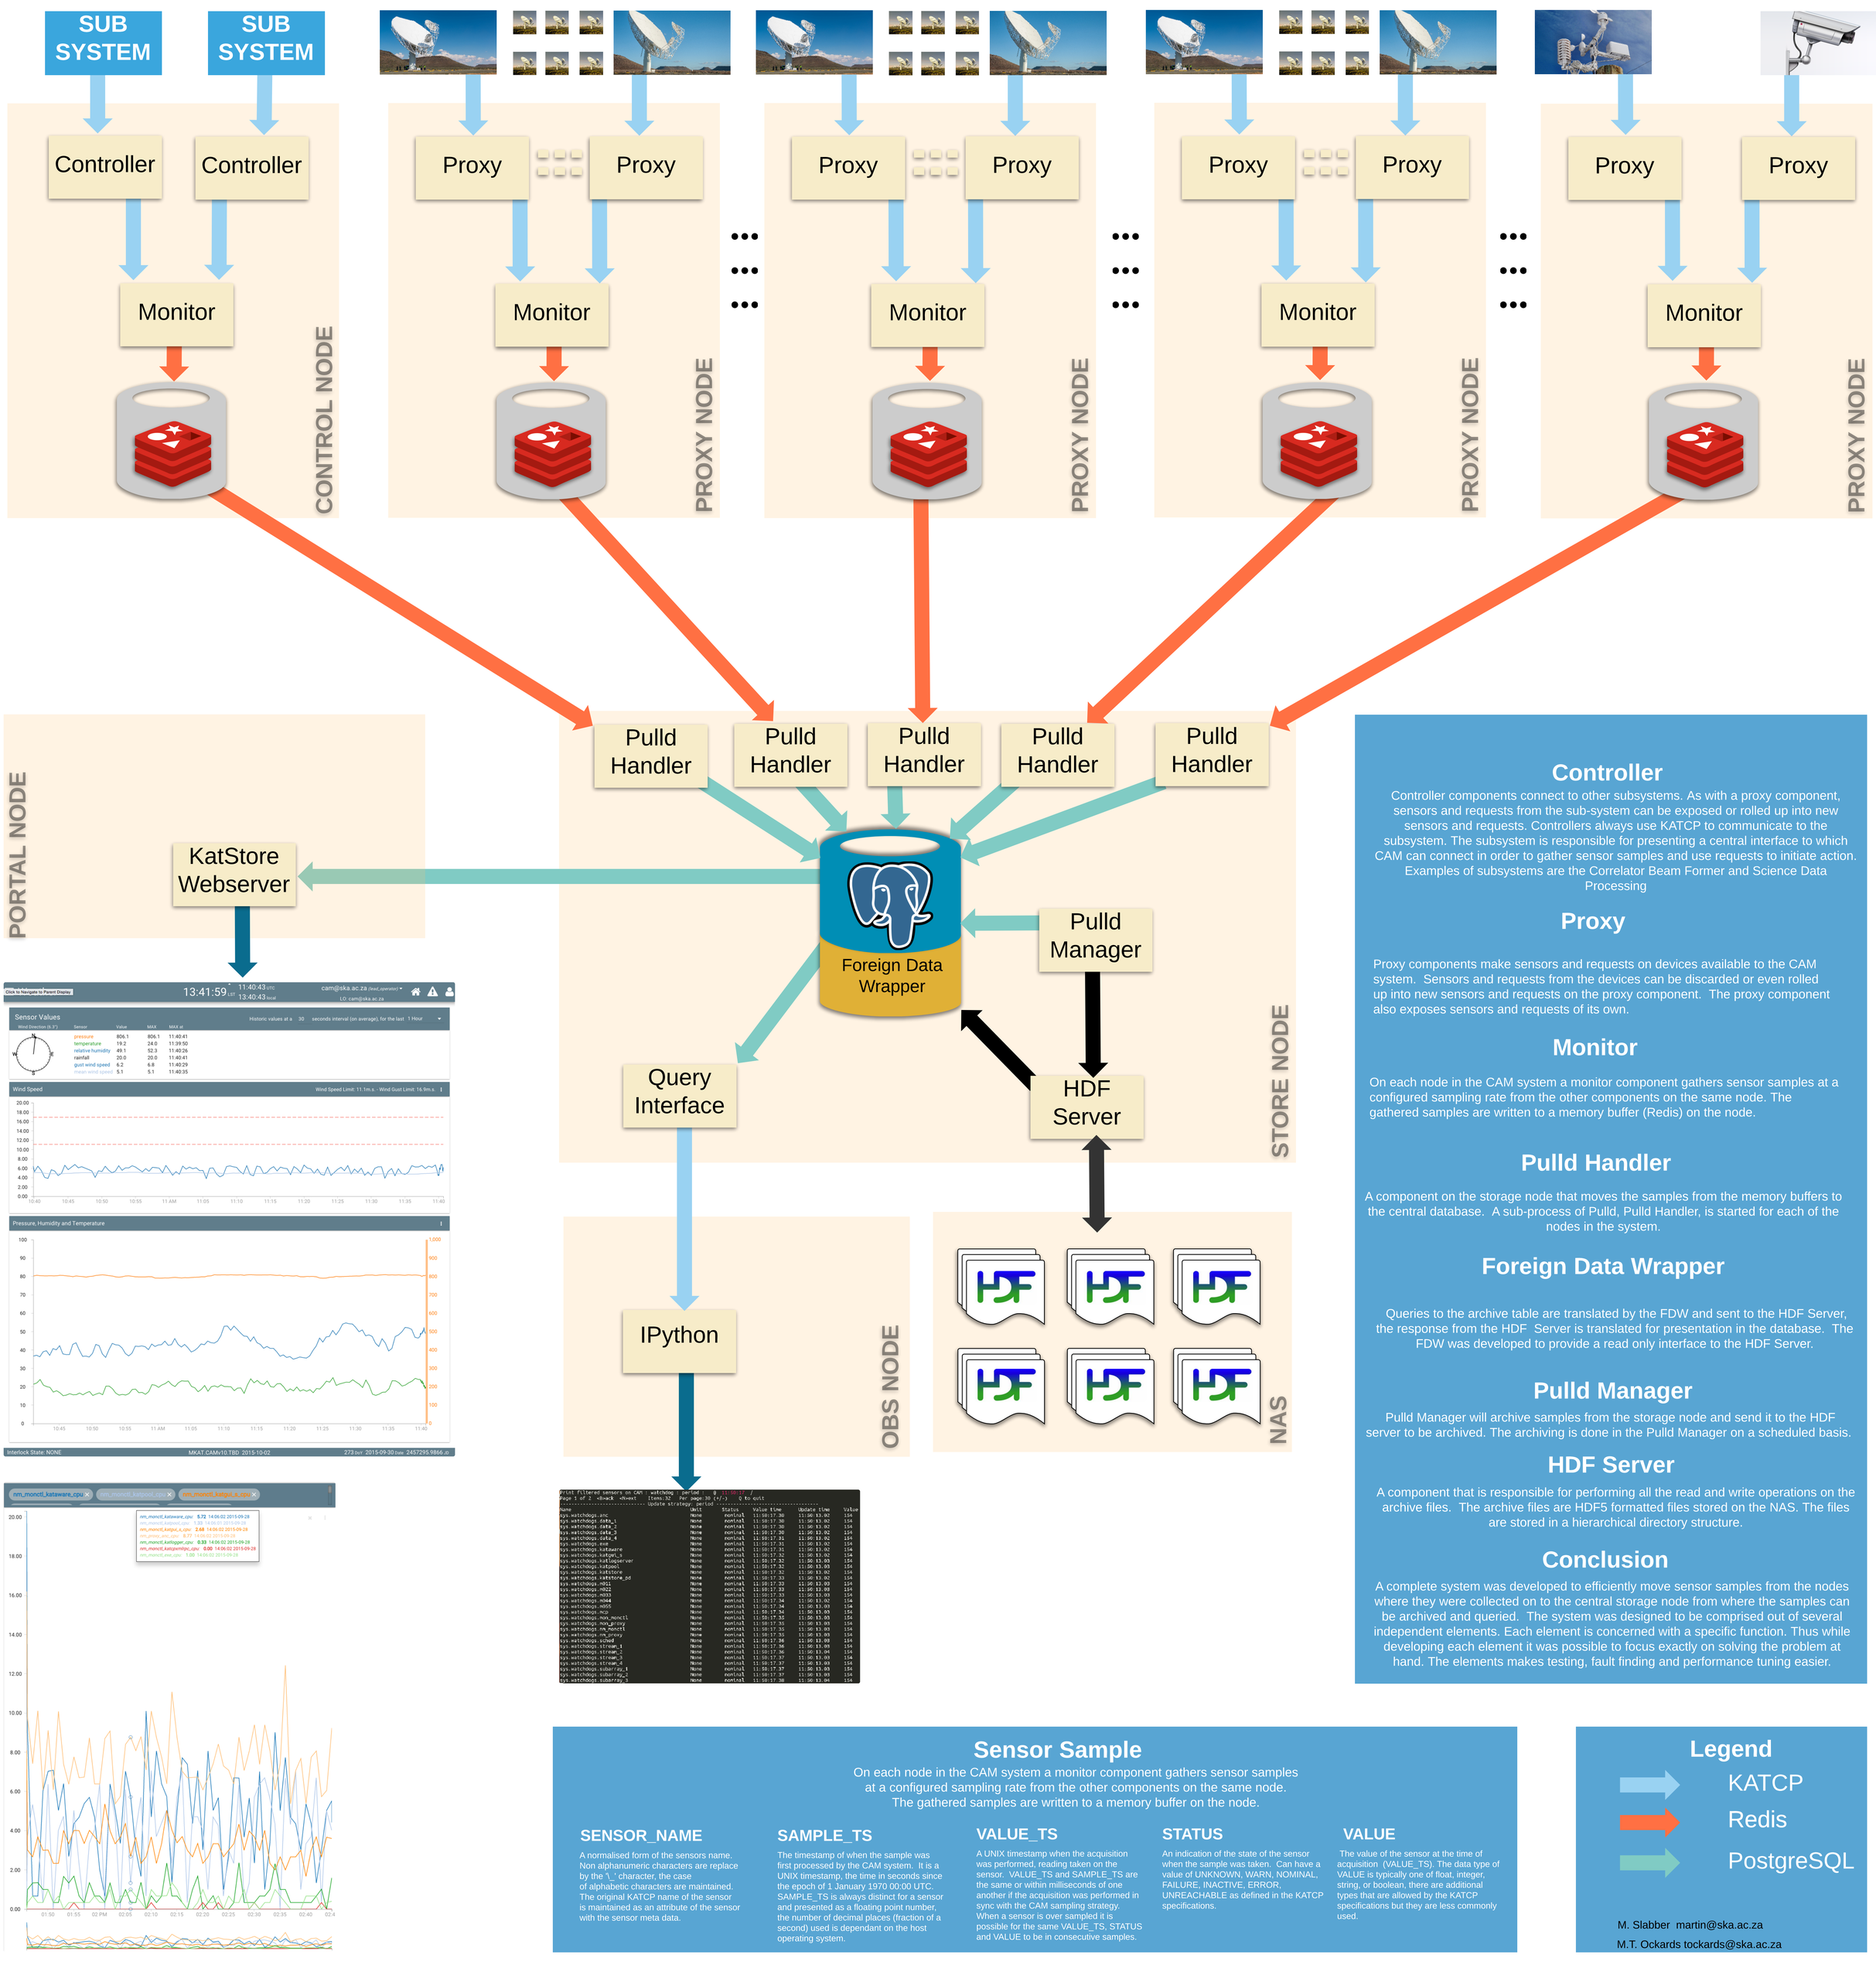

Controller

Controller components connect to other subsystems. As with a proxy component, sensors and requests from the sub-system can be exposed or rolled up into new sensors and requests. Controllers always use KATCP to communicate to the subsystem. The subsystem is responsible for presenting a central interface to which CAM can connect in order to gather sensor samples and use requests to initiate action. Examples of subsystems are the Correlator Beam Former and Science Data Processing

Proxy

Proxy components make sensors and requests on devices available to the CAM system. Sensors and requests from the devices can be discarded or even rolled up into new sensors and requests on the proxy component. The proxy component also exposes sensors and requests of its own.

Monitor

On each node in the CAM system a monitor component gathers sensor samples at a configured sampling rate from the other components on the same node. The gathered samples are written to a memory buffer (Redis) on the node.

Pulld Handler

A component on the storage node that moves the samples from the memory buffers to the central database. A sub-process of Pulld, Pulld Handler, is started for each of the nodes in the system.

Foreign Data Wrapper

Queries to the archive table are translated by the FDW and sent to the HDF Server, the response from the HDF Server is translated for presentation in the database. The FDW was developed to provide a read only interface to the HDF Server.

Pulld Manager

Pulld Manager will archive samples from the storage node and send it to the HDF server to be archived. The archiving is done in the Pulld Manager on a scheduled basis.

HDF Server

A component that is responsible for performing all the read and write operations on the archive files. The archive files are HDF5 formatted files stored on the NAS. The files are stored in a hierarchical directory structure.

Conclusion

A complete system was developed to efficiently move sensor samples from the nodes where they were collected on to the central storage node from where the samples can be archived and queried. The system was designed to be comprised out of several independent elements. Each element is concerned with a specific function. Thus while developing each element it was possible to focus exactly on solving the problem at hand. The elements makes testing, fault finding and performance tuning easier.

Sensor Sample

On each node in the CAM system a monitor component gathers sensor samples at a configured sampling rate from the other components on the same node. The gathered samples are written to a memory buffer on the node.

SENSOR_NAME

A normalised form of the sensors name. Non alphanumeric characters are replace by the '_' character, the case of alphabetic characters are maintained. The original KATCP name of the sensor is maintained as an attribute of the sensor with the sensor meta data.

SAMPLE_TS

The timestamp of when the sample was first processed by the CAM system. It is a UNIX timestamp, the time in seconds since the epoch of 1 January 1970 00:00 UTC. SAMPLE_TS is always distinct for a sensor and presented as a floating point number, the number of decimal places (fraction of a second) used is dependant on the host operating system.

VALUE_TS

A UNIX timestamp when the acquisition was performed, reading taken on the sensor. VALUE_TS and SAMPLE_TS are the same or within milliseconds of one another if the acquisition was performed in sync with the CAM sampling strategy. When a sensor is over sampled it is possible for the same VALUE_TS, STATUS and VALUE to be in consecutive samples.

STATUS

An indication of the state of the sensor when the sample was taken. Can have a value of UNKNOWN, WARN, NOMINAL, FAILURE, INACTIVE, ERROR, UNREACHABLE as defined in the KATCP specifications.

VALUE

The value of the sensor at the time of acquisition (VALUE_TS). The data type of VALUE is typically one of float, integer, string, or boolean, there are additional types that are allowed by the KATCP specifications but they are less commonly used.

Legend

KATCP
Redis
PostgreSQL

M. Slabber martin@ska.ac.za
M.T. Ockards tockards@ska.ac.za