



Yale

ALICE OVERWATCH: Status and Plans

Offline Week - March 2017

Run 264261

TIME SLICES

[Run List](#)

[Basic QA](#)

[Gamma Trigger Low](#)

[AmpEdgePosEMCGALOnline](#)

[MaxEdgePosEMCGALOnline](#)

[Gamma Trigger High](#)

[AmpEdgePosEMCGAHOOffline](#)

[AmpEdgePosEMCGAHOOffline](#)

[AmpEdgePosEMCGAHRRecalc](#)

[MaxEdgePosEMCGAHOOffline](#)

[MaxEdgePosEMCGAHOOffline](#)

[MaxEdgePosEMCGAHRRecalc](#)

[Jet Trigger Low](#)

[AmpEdgePosEMCJELOnline](#)

[MaxEdgePosEMCJELOnline](#)


[Jet Trigger High](#)

Run 264261

Run 264261 started at approximately Tuesday, 25 Oct 2016 11:18:59 (CERN time zone).

[Logbook entry](#)

AmpEdgePosEMCGALOnline



Raymond Ehlers on behalf of the OVERWATCH Collaboration

OVERWATCH Collaboration:

Raymond Ehlers (Yale University)

Markus Fasel (ORNL)

Sarah LaPointe (Frankfurt Institute for Advanced Studies)

Mar 28, 2016

Outline

OVERWATCH

- Introduction and Status
- Upgrade

Recent Activities

- Deployment of the Overwatch Upgrade
- Elasticsearch Investigations

Future Plans

- Apache Kafka
- Triggering Alarms

OVERWATCH/1

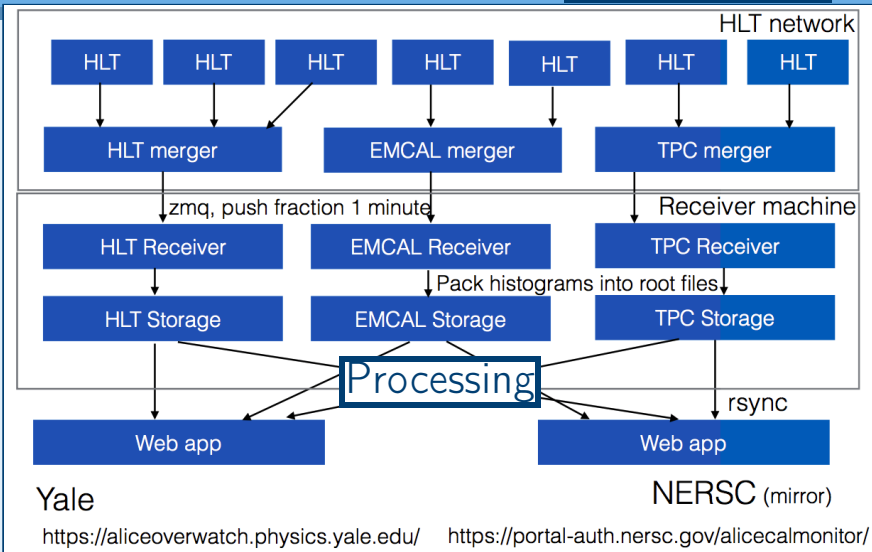
- ▶ **OVERWATCH**: Online Visualization of Emerging tRends and Web Accessible deTector Conditions using the HLT.
- ▶ Created to **monitor and visualize** QA data from the HLT.
 - ▶ Data is stored **persistently**.
- ▶ OVERWATCH handles **spectra, 2D histograms** (for example, EMCal cell amplitudes), etc.
- ▶ Allows slicing of data in **time windows** ("time slicing").
- ▶ Based around two main components:
 - ▶ **Processing** based on PyROOT.
 - ▶ **WebApp** based on flask.
- ▶ Code available at:
<https://github.com/raymondEhlers/OVERWATCH>.

OVERWATCH/2

- ▶ It began as a project during the 2015 Pb–Pb run to provide online **real-time** feedback on the EMCal (beginning with James Mulligan and RJE).
 - ▶ Has since expanded to support **additional detectors and additional features**.
 - ▶ Includes a **plug-in framework** to extend nearly any part of the processing.
- ▶ Currently supports the **EMCal, TPC, and HLT** (includes assorted histograms without a dedicated HLT merger).
- ▶ First version deployed during the 2015 Pb–Pb run, and then ran for all of 2016 with few issues.
- ▶ Collected approximately **100 GB** during this period.

Data Flow from HLT

Image from: M. Fasel



Return to Run List Top of Page User: emcal Logout

OVERWATCH Run List Basic QA

Run 267166	EMC Histograms	
	EMC ROOT Files	←
	TPC Histograms	←
	TPC ROOT Files	←
	HLT Histograms	←
	HLT ROOT Files	
Run 267165	EMC Histograms	
	EMC ROOT Files	
	TPC Histograms	
	TPC ROOT Files	
	HLT Histograms	
	HLT ROOT Files	

User access to data files

Index

Return to Run List Top of Page User: emcal Logout

Run 267166

Run 267166 started at Monday, 05 Dec 2010 04:01:10 (CEBN time zone).
L30000 entry

← Time slice

GeVana Trigger Low

← Histogram selection

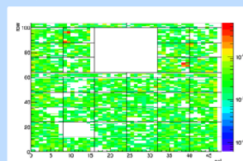
GeVana Trigger High

Jet Trigger Low

Jet Trigger High

Run view

MaxEdgePosEMCALOnline

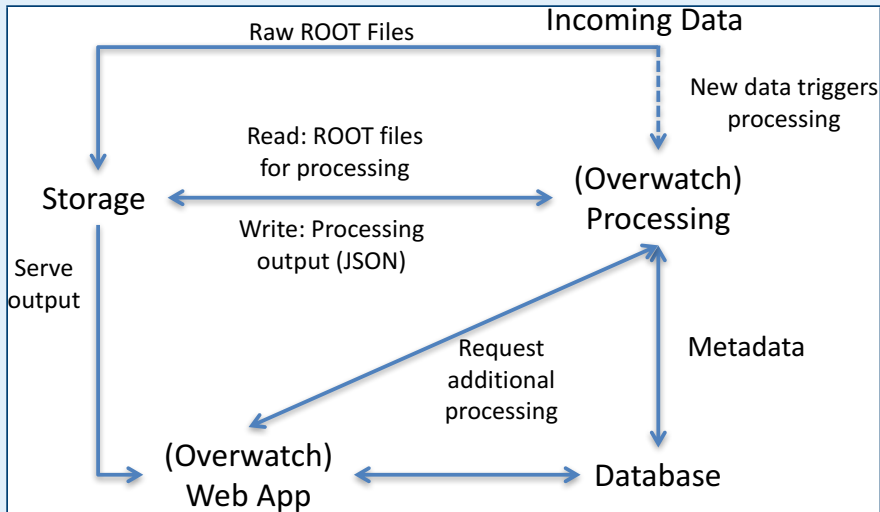


Data

Upgrade to Overwatch

- ▶ Major upgrade developed to improve on current version.
- ▶ Substantial improvements to user experience.
 - ▶ Improved interactivity and visualization with **JSRoot and AJAX**.
 - ▶ **Cleaner and more consistent interface** built with Google's Polymer.
- ▶ Much more **capable and extensible** back end
 - ▶ Much **better performance**.
 - ▶ Provides **plug-ins to all stages** of processing.
 - ▶ Moves towards micro-services architecture (some improvements still needed here).
- ▶ Major infrastructure improvements
 - ▶ Migrating **all infrastructure to CERN** by utilizing OpenStack and coordinating with Offline.
 - ▶ **Commonly accessible** data storage.
- ▶ Hoping to deploy for the beginning of the 2017 pp data taking.

Data Flow in New Version



Next Version Look-and-Feel

The image displays four screenshots of the OVERWATCH software interface, illustrating the look-and-feel for the next version.

Top Left Screenshot: OVERWATCH Run List
This screenshot shows the main interface with a sidebar on the left containing a "Run List" and "Basic QA" sections. The main area displays a list of runs with details for Run 300005, Run 264261, Run 259888, Run 259868, Run 257389, and Run 246991. Each run entry includes the run number, date, and time, along with links to "EMC Histograms", "EMC ROOT Files", "HLT Histograms", "HLT ROOT Files", "TPC Histograms", and "TPC ROOT Files".

Top Right Screenshot: Run 264261
This screenshot shows the details for Run 264261. It includes a "Run List" sidebar, a "Basic QA" section, and a "Gamma Trigger Low" section with a "Logbook entry" link. The main area displays the run number, start time (approximately Tuesday, 25 Oct 2016 11:18:59 CERN time zone), and a link to "AmpEdgePosEMCGALOnline". Below this is a heatmap titled "Integrated amplitude EMCGAL patch Online" showing event rates per event.

Bottom Left Screenshot: Run 264261 (Processing Settings)
This screenshot shows the details for Run 264261 with a "Processing Settings" dialog box open. The dialog box has a "Time Dependent Merge" section with a slider for "Hot channel threshold (µ1000)" set to 32. There are also "Scale by nEvents" and "CANCEL" / "ACCEPT" buttons. The background shows the same heatmap as the top right screenshot.

Bottom Right Screenshot: Run 264261 (FastORL1)
This screenshot shows the details for Run 264261 with a "FastORL1" section. It includes a "THRESHOLD" of 0.248 and a "FAST OR HOT CHANNELS ID" of [1006, 1875, 4825]. Below this is a scatter plot titled "L1" showing "Confidence Region ID" vs "EMC/TPA_HotFastORL1". A table in the bottom right corner provides summary statistics for the FastORL1:

EMC/TPA_HotFastORL1	
Entries	4390/1693
Mean	2132
Std Dev	1390

Outline

OVERWATCH

- Introduction and Status
- Upgrade

Recent Activities

- Deployment of the Overwatch Upgrade
- Elasticsearch Investigations

Future Plans

- Apache Kafka
- Triggering Alarms

Deployment and Commission of the Overwatch Upgrade

- ▶ Migrated receivers and data to **CERN OpenStack**.
 - ▶ Created a **dedicated project** for ALICE HLT QA.
 - ▶ Yale site will stay until the transition is completed.
- ▶ Data storage on EOS is required for all data to be accessible from the **HLT receivers, processing, and web app**.
 - ▶ We now have **dedicated storage** on EOS.
 - ▶ The previously collected data was **recently migrated** to EOS.
- ▶ Currently adapting some code to properly access data on EOS.
- ▶ Data taking is fast approaching and there is still some work to do, but we hope to deploy soon!

Investigation into Elasticsearch for Histogram Storage

- ▶ Overwatch stores full histograms, so we investigated the possibility of storing both metadata and histograms in Elasticsearch.
- ▶ Single database source could simplify operation (such as interactively accessing data via Swan).
- ▶ Challenging due to the **large amount of redundant histogram information**, amongst other issues.
- ▶ Developed a possible data layout (code linked here).
- ▶ Early tests observed the Elasticsearch database was **~10 times larger** than storing raw data.
- ▶ Storage of histograms in Elasticsearch **does not** seem to be feasible.
 - ▶ **Still useful for trending** information.
 - ▶ Perhaps for metadata, **depending on additional external factors**.
- ▶ For further information, see Markus' presentation at the February DPG General Meeting:
<https://indico.cern.ch/event/611050/#sc-2-2-online-qa>.

Outline

OVERWATCH

- Introduction and Status
- Upgrade

Recent Activities

- Deployment of the Overwatch Upgrade
- Elasticsearch Investigations

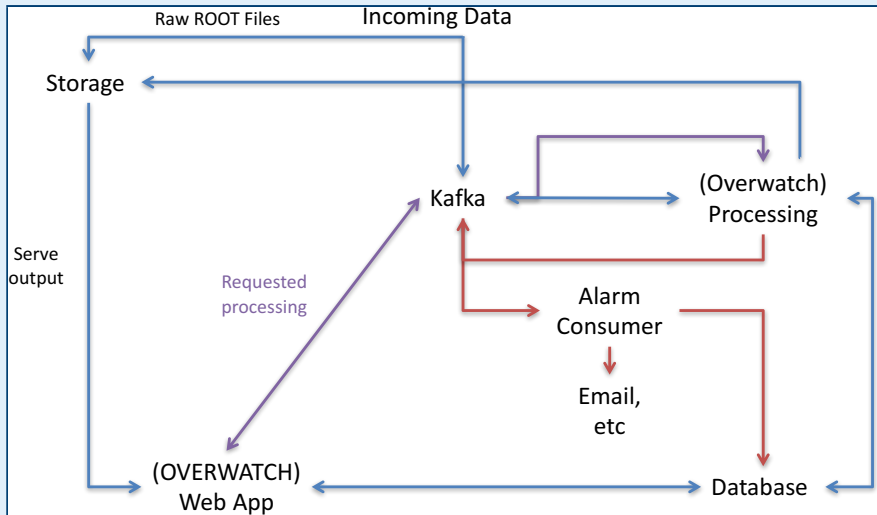
Future Plans

- Apache Kafka
- Triggering Alarms

Apache Kafka

- ▶ Data processing is currently triggered by receiving data.
 - ▶ This aspect of the system is **not particularly flexible**.
- ▶ We are interested in using **Apache Kafka** to managing the processing of incoming data.
- ▶ Currently **setting up a small Kafka test system** in the Overwatch OpenStack cluster.
 - ▶ Initial tests will attempt to receive histograms into the system and then process them.
- ▶ Investigating the potential of **utilizing elements** of the **existing Overwatch processing** functionality.

Architecture with Kafka (Proposed)



Towards Triggering Alarms

- ▶ There is already some **very basic** alarm code in the Overwatch processing framework.
- ▶ Basic examples include:
 - ▶ Check if there are **outliers** from a distribution.
 - ▶ Check for values above a **threshold**.
- ▶ Work is needed to:
 - ▶ **Improve and generalize** the processing framework, particularly with respect to **extracting and saving trending information**.
 - ▶ Implement **additional processing functions**, such as comparing to reference spectra.
 - ▶ Implement **alarms based on the processing functions and trending**.
- ▶ **Service task proposed** to address this need, with a focus on implementing processing functions and alarms.
 - ▶ Amongst other activities, need to work with some detector experts to determine how best to implement some example alarms.
- ▶ Should be completed with an eye towards integration with Kafka.

Conclusions

- ▶ Overwatch provides **online monitoring and visualization** of data from the HLT.
 - ▶ The data is **stored persistently**.
- ▶ The current version has run for **all of 2016** with few issues.
- ▶ We are working **towards deploying the upgraded version** soon.
- ▶ Continuing to work on:
 - ▶ Investigation of **Apache Kafka**.
 - ▶ Improving processing and implementing **trending and alarms** framework.


Backup

Service Task Proposal

Development of a trending and alarm framework for online HLT QA

Several QA components on the HLT provide QA data for different detectors at discrete times as simple 1 or multi-dimensional histogram. These data are visualized using the Overwatch web application. The scope of the task is to implement a component processing the histograms and extracting trending information. These trending information will be sent to an Elasticsearch database for visualization. In addition, by comparing to limits defined by users, automatic alarms should be raised (E-Mail to detector responsible with alarm message + Log in database). The alarm handling should be in testing phase during the pp data taking period end of 2017.

OVERWATCH Architecture



OVERWATCH Architecture

- ▶ OVERWATCH is python and ROOT based.
- ▶ Split into two main parts:
 - ▶ Processing
 - ▶ The Web App
 - ▶ Depends on the processing module for user requested processing.
- ▶ Receiver from HLT written in C++ and utilizes ZMQ.
 - ▶ Data is received approximately every minute and time stamped.
 - ▶ Large, but not unreasonably large, amount of volume. (≈ 100 GB/year for EMCal + HLT + ≈ 3 months of TPC data).
- ▶ Designed to run as micro-service, so can start instances as needed.
- ▶ Since OVERWATCH processes data from the HLT, our architecture is similar to data processing for Run 3 when the HLT->Event Processing Node.

OVERWATCH Processing

- ▶ Processing utilizes PYROOT and runs every minute on newly received data.
- ▶ Manages run and subsystem data via ZODB (Zeo Object Database)
 - ▶ Makes management of python objects straightforward.
 - ▶ Also used by Indico.
 - ▶ We aren't strongly attached to this DB.
 - ▶ Any appropriate SQL or NoSQL database would be fine.
 - ▶ Code is not really reliant on ZODB - easy to switch elsewhere.
 - ▶ Structure is hierarchical.
 - ▶ Run->Subsystems->HistogramGroups->Histograms.
- ▶ Actual data is just stored on disk in root files.
- ▶ Output of processing is stored in json files.

OVERWATCH Processing

- ▶ Additional processing available per detector/histogram.
 - ▶ Can check values in particular histograms, stack hists, create additional hists to summarize, etc.
 - ▶ Can also handles alarms.
- ▶ Time slices
 - ▶ Can make arbitrary selections in time (0-10 minutes, 5-17, whatever, etc) within a run*
 - ▶ Can also select processing options. Hot channel thresholds, scale by number of events, etc.
 - ▶ Caches result - only reprocess if absolutely necessary.
- ▶ Basic trending support for extracted values. More to come.

* - subject to intrinsic time resolution of HLT of 2 minutes encompassing \approx 3 mins.

OVERWATCH Web App

- ▶ The Web App is built on Flask.
 - ▶ Interface built using Google's Polymer.
 - ▶ Pages are built using the Jinja2 template engine. Each detector can build their own.
 - ▶ JSROOT used for presenting histograms, with fall back to static images.
 - ▶ Navigation handled by AJAX, with fall back to full page reloads.
- ▶ Display histograms according to detector specification.

Try it yourself using docker

- ▶ Docker image available at <https://hub.docker.com/r/rehlers/overwatch/>.
- ▶ Can be tested using the following procedure (to be streamlined - we don't deploy processing like this at the moment).
 - ▶ Download test data from: <https://aliceoverwatch.physics.yale.edu/testingDataArchive>.
 - ▶ `docker run -it -v data:/overwatch/data -e deploymentOption=devel overwatch /bin/bash`
 - ▶ `cd /overwatch && python runProcessRuns.py`
 - ▶ `cd deploy && python updateDBUsers.py`
 - ▶ `cd /overwatch && python runWebApp.py`
- ▶ Still testing some cases - please let us know if you run into any trouble!

OVERWATCH

Online Visualization of Emerging Trends and Web Accessible Detector Conditions using the HLT

Additional improvements

- ▶ Improve time series summary support.