

Advanced Automated Analytics Using OSS Tools

Fred Elmendorf

Manager, Grid Solutions Services
Grid Protection Alliance, Inc.

felmendorf@gridprotectionalliance.org

Abstract— The exponential increase in data available to analyze power system events is universally recognized, but in many cases the approach to using this data is to do what we already do but do it faster, or get more people to do it. Unfortunately, spinning the hamster wheel faster is not keeping up with the demand to make decisions faster in support of grid modernization. Open source software (OSS) tools offer tremendous opportunity for collaboration that encourages innovation, and the speed and flexibility of development to keep pace with these demands. This paper will describe the OSS approach to developing tools for advanced, automated analytics, and some of the benefits it offers. Next, an overview will be presented to describe the high level building blocks necessary to construct a comprehensive automated system. Finally, a detailed description of the newest of these building blocks will be presented with practical use cases to illustrate the value of this approach. The new building blocks to be covered include remote data acquisition, evaluation of breaker operations during events, and capacitor bank related events. The conclusions will include ‘food for thought’ to encourage the exploration of new ways to solve emerging power system event analytic issues. Attendees will gain practical knowledge of proven and emerging OSS approaches to “Advanced Automated Analytics”.

Keywords—*open source software; OSS; automated analytics*

I. OSS APPROACH

Open source software is software whose source code is available for modification or enhancement by anyone (OSI). The open source software (OSS) approach that we see today is not vastly different from the development community that grew organically around computers in the early days. In an effort to rapidly take advantage of the ‘new’ computational technology of that time, ideas were freely discussed and shared. Over time business models developed to extract financial benefit from the unique features in a company’s software, and by ‘closing’ the source code so that it couldn’t be copied and modified, they were able to maximize their financial return. Some of the down-side results to this approach were inhibited collaboration, isolated development resources, and vendor lock-in. Although this proprietary strategy has served some companies very well, the response in the development community at large was to create organizations that define and promote OSS.

Today we are seeing a dramatic resurgence in the development and acceptance of OSS for many reasons. The rapidly changing electric utility industry can benefit from a number of these potential benefits:

OSS:

- Stimulates innovation
- Encourages and facilitates collaboration
- Reduces time to deployment
- Reduces cycle time for improvements
- Lowers total cost of ownership
- Improves code quality
- Improves security through community review

There are two major licensing strategies within the field of OSS; permissive licensing, and restrictive licensing. A restrictive license subscribes to the notion that all code should be open source forever, including any code used in combination with the OSS code, and all derivatives. A permissive license subscribes to the notion that each developer should be able to decide whether their code is open source or proprietary, even if it is a derivative or prior OSS work. The OSS building blocks described in this paper and used in the case studies presented are licensed under permissive licenses such as the MIT or BSD 3 Clause licenses (MIT License) (BSD License).

II. BUILDING BLOCKS OF AN AUTOMATED SYSTEM

An end-to-end solution to provide fleet-wide information from the aggregation of data recorded by individual devices across the utility’s footprint can be described as three general divisions or building blocks.

ONE: Architecturally, the first of these building blocks is the physical transport of data from the sensing and recording device to a central repository. This task can be accomplished through a spectrum of processes that range from completely manual, to fully automated. A manual process could be as primitive as copying files directly from the device onto storage medium such as a USB drive, carrying it back to the office, and copying it onto a computer, or as sophisticated as opening up a desktop application and initiating a file retrieval operation that ultimately places the data in the same repository as the more primitive process. An automated process could be as simplistic as setting appropriate criteria in a vendor provided software application for each unique brand of device, or as advanced as fully automated retrieval/receiving process that runs autonomously to archive data from all remote devices. Regardless of the approach

used, the first block in the process is to get the data to a central repository.

TWO: The second architectural block is analysis of the data. And once again there is a spectrum of approaches that range from completely manual to completely automated. For this block a manual process often consists of opening a vendor supplied software application, manually locating the data to be analyzed, and visually reviewing the data with a calculator and note pad nearby. At this end of the analysis spectrum, it is often necessary to employ a separate desktop program that is specifically designed to analyze the data from the respective vendor. Moving up the analysis automation spectrum there are third party desktop tools that allow the analysis of a selection of similar data sources in a single software application. The manual or semi-automatic tools for analysis often do not facilitate the extraction and long term trends and system responses but are typically limited to the production of individual manually developed reports. A fully automated analysis system is capable of performing the specified analysis without any manual intervention, and providing automated reports and notifications of conditions that meet predetermined conditions. Ultimately, the automated system also builds a comprehensive database for long term trending analysis in addition to analyzing individual events. A database built from an automated analysis system can also position data for fleet-wide visualization through web based tools and dashboard applications.

THREE: The third architectural block in this view of an end-to-end system is visualization of information derived from analysis of the data. When there are a number of potential consumers of the information that may use the information to make operation or restoration decisions, proactive notifications such as an automated email to a targeted recipient list can be very effective and greatly reduce response time. For an office or control room environment, an ideal method for visualizing large quantities of information is through a dashboard. And just like the dashboard of a vehicle, the annunciations, displays, and images used to present the information can be tailored to the needs of the audience. For example, a crane operator at a construction site would have different displays and information available to him than the driver of a sports car but they both use dashboards. Again using the analogy of a modern dashboard in a car, when some information displayed on the dashboard indicates a cause for concern, such as a low fuel light, the simple click of a "More Information" button should quickly give you

the needed detail of 'Thirty-six miles to Empty'. In a power system event driven dashboard, the information should be presented intuitively with easy access to supporting details.

III. NEW BUILDING BLOCKS

Over the past several years the three basic building blocks described above have been developed in the OSS community. Some of the functions of these building blocks have been described in previous papers, and the evolution of the blocks has taken place in a 2, 3, 1 order rather than 1, 2, 3.

The first of these blocks to be developed in OSS was the automated analytic engine (openXDA). It takes data as quickly as available from the data retrieval process, performs the specified analytics, provides email notifications, and builds a database. In its present form, openXDA can consume any disturbance data file presented in standard COMTRADE or PQDIF file formats. Although some additional proprietary formats have been developed, it is strongly recommended that standard file formats be used in automated processes. In the present form, openXDA has been deployed to recognize and identify disturbances, calculate single and double ended fault distances, and recognize some abnormal equipment operation such as breaker timing and capacitor bank health. These analytics are just the beginning and examples of what can be done in using this approach. Any analytic that is appropriate for disturbance or trended power system data could be added.

Next, the OSS visualization block was developed in the form of a dashboard. The Open PQ Dashboard (PQDashboard) provided the first comprehensive 'fleet view' of all reporting devices together with a summary of performance, and the ability to quickly navigate to full resolution detail of underlying data. In addition to viewing summary and detailed information about event and trended data for the entire fleet of devices, it also provides visualization of the underlying data quality. Visualizations in the dashboard include: GIS or annunciation panel style display for the fleet of devices, a summary chart by day of the characteristic being displayed such as sags or faults, detailed lists of events or trends, and interactive charts of full resolution waveform trended data. The Figure 1. below is a collage that shows a number of functions within the dashboard. As with openXDA, this is just the beginning of the visualization building block. Additional features are actively being added, and any appropriate visualization could be included.

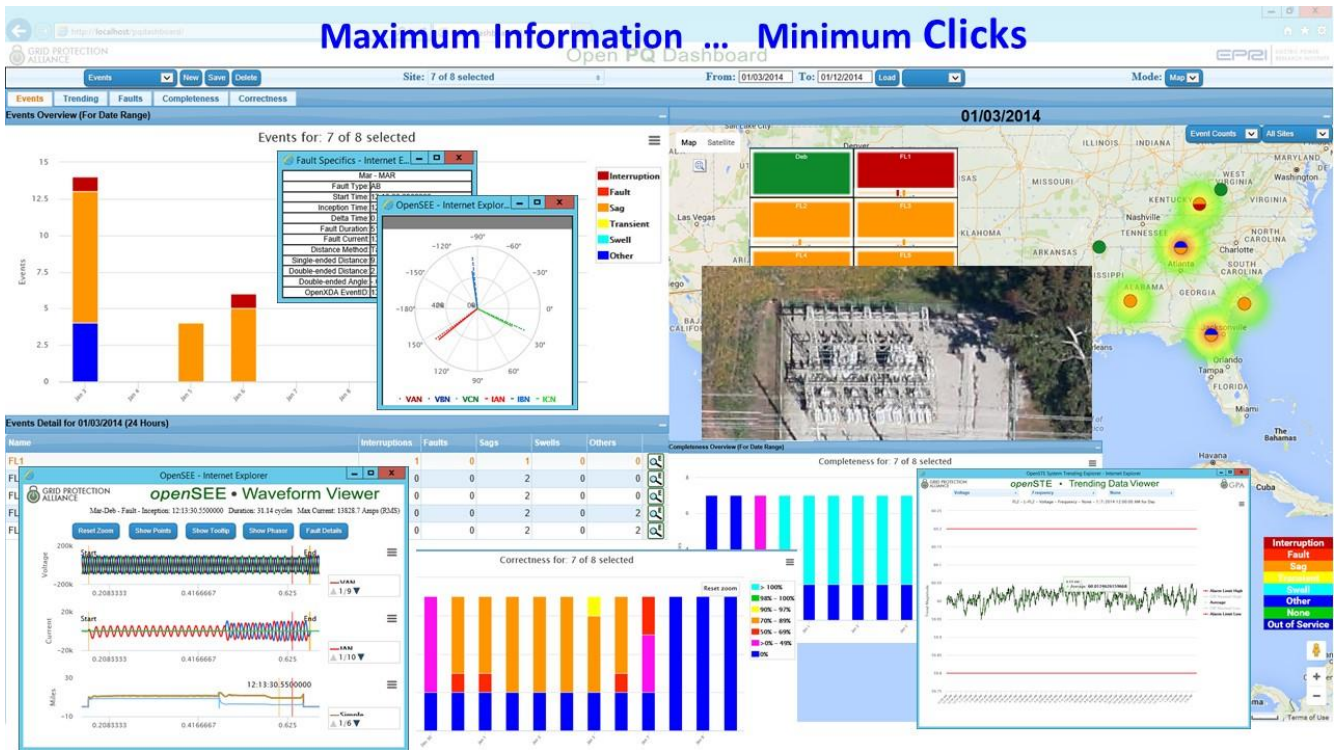


Figure 1. Open PQ Dashboard Collage

The third OSS block to be developed is data acquisition, and has come together over the past year. The development of the data acquisition block was inspired through the discussion of internal work being done at Dominion Virginia Power and the Tennessee Valley Authority (TVA). Both utilities have contributed to its creation. It is in a beta version now, and being tested at TVA. The OSS implementation of this block is called open Meter Information Collection system (openMIC). It uses a web based interface, and can manage any number of network or dial-up connections simultaneously. Presently it uses file transfer protocol (ftp) to retrieve the disturbance files but additional enhancements are being explored.

With these three blocks now available, it is possible to build an end-to-end solution for automated analysis and visualization of power system events entirely with OSS tools.

IV. EXAMPLE USE CASES

Once an automated analytic system is in place many opportunities begin to emerge. Looking at information aggregated from many devices such that every event record is analyzed using the same methods and criteria and a comprehensive database is in place, the value of this new source of information begins to be realized. Here are three examples of new analytic information becoming available by leveraging this approach.

A. Breaker Timing

Georgia Transmission Corporation (GTC) has implemented the analysis and visualization building blocks, and had the need to analyze actual breaker operation times. To accomplish this,

the measured operation time is compared to name plate data. During the past few months the analytic engine has been enhanced to accept the digital trip coil energize signal and compute the delta time between that signal and the time when the current goes to zero. The computed operation time is then compared to the expected operation time, and a flag is set if the measured operating time exceeds the expected time. Requirements are being discussed to extend the capability of this analysis and reporting to compare and report on other quantities such as phase to phase relationships and transients such as breaker restrike.

B. Capacitor Health

TVA has also installed the analysis and visualization blocks, and is testing the data retrieval block. one of the areas of new possibilities being explored to leverage the power of this automated analytic system is capacitor health. By deploying specific monitoring equipment at capacitor bank installations, new waveform data is available to analyze. The analysis of magnitude, step changes, and the timing of changes in the waveform data is providing new information and insights regarding the health of the capacitor bank. As this approach continues to mature alarming and notifications will be automatically generated when thresholds are exceeded.

C. Statistical analysis

Having vast quantities of historical data provides another opportunity for new information. The analytic and visualization blocks have been used to construct a statistical history of normal operating ranges for every trended quantity available from the device. In the initial implementation of this function, historical data was aggregated for a year by hour of the week. Subsequently, every recorded data point is compared to the

normal operating range, then alarms and notifications can be produced when specified limits are exceeded. This provides notification of anomalies such as failing equipment or miscalibration even if established alarm limits have not been exceeded.

V. FOOD FOR THOUGHT

Using these OSS building blocks, fully automated advanced analytic systems can be built to perform any analysis that is appropriate for the type of data available. As they stand today, the blocks are functional and provide dramatically improved access to information compared to previous manual methods. However great the value is that what we see today, it is only the tip of the iceberg. Other functions such as data quality and device availability and performance are also being applied in this environment. Each of these OSS building blocks is freely available for enhancement, extension, and adaption. Through the collaboration and innovation afforded by the OSS approach the functions of these blocks will continue to grow rapidly.

The fact that these tools are freely available in OSS projects means that anyone can use them ‘as is’ or become a part of the community to enhance and extend them. The references below contain links to information regarding the OSS licenses used, and the OSS code repositories where the projects are managed.

VI. REFERENCES

- BSD License
https://en.wikipedia.org/wiki/BSD_licenses
- MIT License
https://en.wikipedia.org/wiki/MIT_License
- openMIC
<https://github.com/GridProtectionAlliance/openMIC>
- openXDA
<http://gridprotectionalliance.org/products.asp#XDA>
- OSI
<https://opensource.com/resources/what-open-source>
- PQDashboard
<https://github.com/GridProtectionAlliance/PQDashboard>

Brief Bio

Fred L. Elmendorf

felmendorf@GridProtectionAlliance.org

Office: (423) 973-4729 Cellular: (423) 280-2464

Fred Elmendorf is Manager of Grid Solutions Services, and Senior Project Manager at Grid Protection Alliance (GPA). He is responsible for expanding GPA's tool suite to include disturbance analytics in addition to their existing synchrophasor tools. Mr. Elmendorf has developed and supported applications in many power system areas, and prior to joining GPA, spent over 30 years with TVA. While at TVA, Mr. Elmendorf was Power Quality Manager for Power Systems Operations, responsible for all long-term power quality monitoring, lightning data systems, and substation IED data integration. He has been actively involved in many research projects, is an IEEE member, and has a B.S. in Computer Science from the University of Tennessee at Chattanooga.