21, rue d'Artois, F-75008 PARIS
http : //www.cigre.org

CIGRE US National Committee
2017 Grid of the Future Symposium

Multi-Vendor Data Concentration of 7000 Distribution Devices to Facilitate Asset Optimization Analytics

A. LYTZ
SUBNET Solutions, Inc.
Canada

SUMMARY

A major eastern US Utility company invested significantly in FLISR over the past years across its electric distribution system by installing thousands of overhead Reclosers to automate power restoration and reduce the impact and length of power interruptions. A key goal to Distribution Automation is historizing the electronic Recloser relay data to facilitate optimized inspection and maintenance practices for effective Asset Management using the PI System.

Deploying more than 7000 devices from various vendors, all of them configured to report DNP3 unsolicited data over a cellular network, and aggregating all 100,000 data steams to a singular head end Virtual RTU was achieved through the SUBNET system. This project utilized the PI SDK to interface the SUBNET Virtual RTU with the PI Data Archive for continuous data historization.

The ultimate goal of this project was to connect data concentrators to the PI System and leverage non-operational data to support key business case analytics. The subsequent analysis produced a high return on investment by providing engineers and asset management with rich data dashboards and the ability to build event frames to track FLISR performance and asset performance, which saves periodic inspection and maintenance costs.

KEYWORDS

Distribution Automation, Asset Management, FLISR, Data Concentration, Data Analytics, DNP3

alan.lytz@subnet.com

Introduction

Two primary functions of a Distribution Automation (DA) system are; to improve reliability by expediting power restoration, and providing grid operators with robust operational data analytics. Every edge device is a sensor that relays critical grid performance data that can reveal a diverse number of analytical insights. However, before any analysis can be performed, DA data needs to be aggregated and historized, which requires integration of many IT and OT systems.

Large scale Distribution Automation systems with multi-vendor devices present several IT and OT challenges. IT challenges often include network load balancing, security patching, hardware constraints, and application integration. On the OT side, managing device communication protocols, device configurations, passwords, relay maps, and organizing the immense volume of data produced are only a few of the challenges. Once data is acquired and historized, the organization needs to be able to transform raw data into actionable insights. This paper details a Utility's design decisions to overcome these challenges by implementing a scalable solution to provide the business with rich data and displays while under heavy IT cost constraints.

Defining the Problem

The Utility needed to provide engineers and asset managers with historical DA data in a centralized location. Asset Management's desire was for inspection and maintenance data from reclosers such as battery level checks, fault current magnitudes, number of device operations, and contact wear calculations. Grid engineers were interested in the effectiveness of FLISR, time to first switch move, Volt-Var characteristics, and low impedance fault data. Leadership was interested in understanding how their automation investments were returning value in the various performance metrics. All departmental goals focused on historical data analysis of non-operational data[1] and can be summarized with four overarching goals:

1. Improve System Reliability
2. Support Condition Based Maintenance programs
3. Provide Operational Visibility
4. Build a foundation for future predictive analytics

In order to achieve these goals, the first step was effectively acquiring the relevant data from the edge devices. Distribution pole-top mounted automation devices, included electronic relcosers, capacitor banks, and motorized switches from various vendors had been continually installed since 2010 with a combined total population reaching nearly 7000. The SCADA Master (DMS) maintained operational control of these DA devices and operational data[2] was displayed on operator's screens and historized in the PI Data Archive, however, non-operational data was not yet captured. In order to acquire both data sets and historize them in the PI Archive, a new strategy was needed to aggregate and organize this data using limited staffing resources. The problem the Utility faced was; how to we cost effectively aggregate data from 7000 devices and get it into the PI historian.

---

[1] Non-operational data is defined as data that operators do not need such as maintenance related data (e.g. number of operations, wear,

[2] Operational data is defined as data that System Operators use to control and maintain the grid (e.g. Volts, amps, status)

Solution Design

Three key business focuses guided this project's design process. First, any solution must be easy to maintain with limited staff. A solution cannot require several engineers or IT staff to spend weeks per year maintaining hardware or patching interfaces. Second, the solution must be NERC CIP compliant and provide robust security options. Lastly, the solution must be cost conscious and aim for the lowest maintenance expense costs possible.

Using the business focus as a guide, IT and OT teams met to better understand the limitations of the current architecture and where significant improvements could be made. Before this project, hundreds of simple rack mounted physical RTUs in a centralized data center worked to aggregate operational data and port it to DMS solely for system operators to use. Only after data was in the DMS system could it be relayed to other systems such as the PI Archive or other systems. This architecture created a need to maintain hundreds of physical RTUs and inflated the number of pass through data streams above the DMS license threshold (which was only 100,000 data streams at the time). With the addition of non-operational data, both RTU infrastructure and DMS data licensing would need to be expanded, thus growing staffing needs and increasing yearly O&M.

After consideration of many options, the Utility decided to remove the hundreds of physical RTUs and deploy virtual server RTUs that can port data to DMS, PI, or any other OT system using slave ports. This option allows easy scalability of new RTUs when new edge devices are added and system expansion is required. Also, O&M costs are very low and off-the-shelf interfaces are prebuilt for many OT applications.

Each Virtual Machine RTU polls 200 devices and roughly 5,000 data points via DNP3 over unsolicited cellular network, thus for 7000 total devices 35 VM's were required.

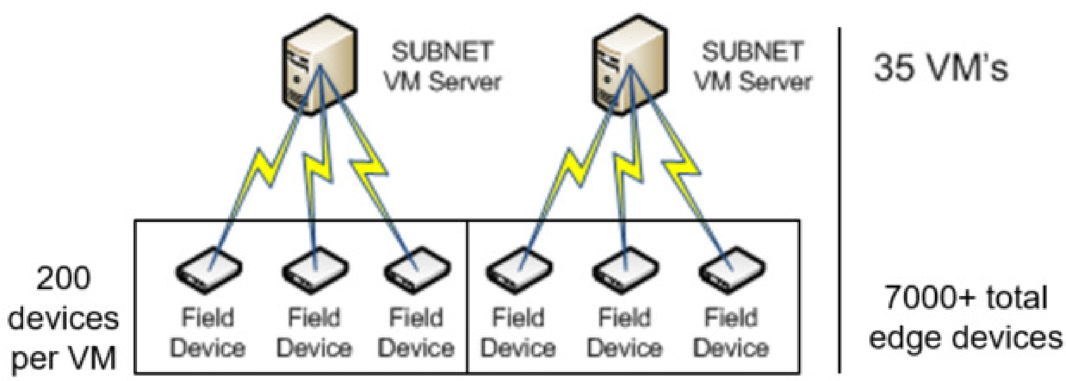

Figure 1: Basic Virtual RTU Architecture

One main advantage is the ability to modify, stop, and start DNP polling individually. Other solutions require restarting connections of all devices in order to change configurations or add devices, which can be disruptive since the device commissioning rate averaged 6 new installations per day.

The communications architecture allows data to take two paths to the PI Archive. Operational SCADA data is ported to both DMS *and* the PI Archive simultaneously, while non-operational data goes directly to PI without touching DMS. This design was chosen to keep DMS a 'purely' operational system and maintain high network bandwidth and lower data licensing costs. The data path to PI is achieved by polling each VM RTU using one "collector" server called the "SUBNET PI Server" that aggregates all 35 VMs and creates one interface node to PI rather than 35 separate nodes each interfacing with PI individually.

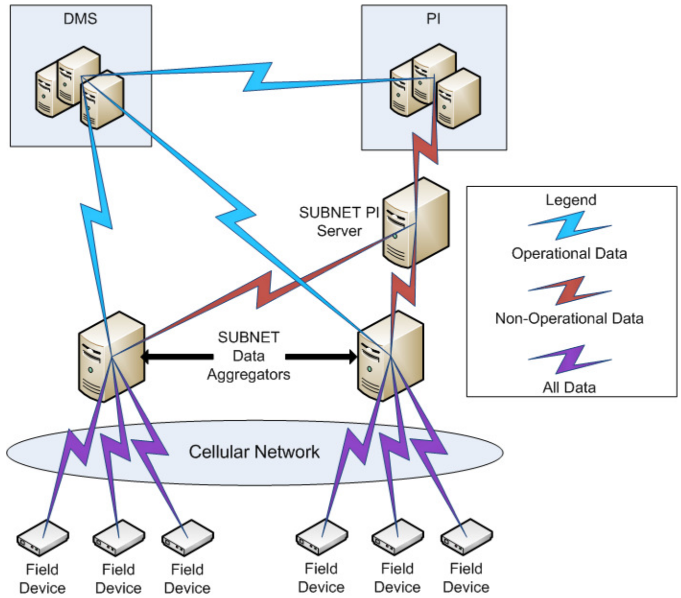

Figure 2: Communications Architecture

This architecture is advantageous because as more automation devices are added in the future and additional VM RTUs are required to collect the data, no additional development work is needed to port new data into PI. The 'collector' server is the sole interface to PI rather than 35+ individual DNP interfaces to PI.

Additionally, the design process incorporated a parallel testing environment where modifications to server architecture or new edge device types can be appropriately tested without interference with the production system.

Build

Converting to virtual RTUs required several months to gradually migrate devices to the new platform without interrupting a significant portion of system data flow at any given time. Old physical RTUs were removed from the datacentre and new Virtual RTUs deployed in stages. All 35 VM installations reside on a singular hardware chassis within the network DMZ, with an identical redundant stack available for primary failover. The collector server was configured as a slave port for the 35 Virtual RTUs acting as master data concentrators. With updated firewall rules, the VM RTUs are highly secure and complied with NERC CIP requirements.

Utilizing the PI Software Development Kit (SDK) and the native SUBNET PI interface, the SUBNET PI Server was interfaced with the PI Archive Server. PI tag data compression was increased as non-operational data granularity is generally less critical than operational data. The new non-operational PI tags were then added as attributes to PI Asset Framework element templates. Once connection testing was complete, users had immediate access to 100,000+ more data points than before, and began to build rich displays and dashboards detailing Distribution Automation performance. In the future, when growth of DA devices exceeds the allotted RTU points, new VMs can easily be added to this scheme by activating an additional slave port on the collector server.

Results

This project resulted in a cost effective and scalable data aggregation system with flexible future scalability. Getting non-operational data into the PI Archive was a big win that allows engineers and asset managers to investigate device performance history and work to improve processes. Engineers are now able to rapidly identify and determine capacitor bank issues by analysing cap bank data, which increased customer power quality satisfaction and decreased diagnostic truck rolls saving over $150,000 per year. Additionally, acquiring recloser wear data allowed the utility to eliminate a physical inspection program and instead perform condition based maintenance by monitoring contact wear thresholds saving $600,000 per year in inspection costs.

IT costs for server and network management was reduced from $50 to $35 per device and greatly reduced the complexity of security patching requiring fewer on-site admins. The overall ease of implementation and maintenance of the virtual RTUs was a success for an already very loaded IT department.

Conclusion

The ultimate goal of any utility flush with data is to improve operational intelligence by transforming it into actionable insights. Often the major challenge that precludes meaningful analysis is getting that data from the edge device into an appropriate data warehouse. This project successfully solved the complex problem of acquiring non-operational data from 7000 overhead devices spread out over a larger geographical area and organizing it in a historian database. The solution simplified the RTU architecture and saved time and money maintaining physical hardware by converting to a fully virtual and secure product. A bit of clever systems design also saved interfacing costs with the PI System and allows for scalable future growth with little added development costs. The resultant analysis yielded a high return on investment by providing information previously unavailable.