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

CIGRE US National Committee
2018 Grid of the Future Symposium

Design and Implementation of an Extensible Platform for Large Scale, Programmatic Transmission Outage Planning Analysis

K. D. JONES, M. A. PARKER
Dominion Energy
USA

M. JALALPOUR
University of Tennessee
USA

SUMMARY

Driven by increases in renewable energy penetration, capital expenditure by transmission owners, and overall complexity of the grid, the need has arisen for greater analytic capability in transmission outage planning. Existing tools and techniques center on EMS, point-and-click approaches or desktop-applications with grand assumptions for input data – all of which severely limit the dimensionality of the analysis. Furthermore, little has changed in the last decade or more for the EMS as a study platform while the grid continues to grow and evolve. The net effect of these tooling bottlenecks is that existing processes do not scale or adapt to meet the demands for new and better analytic results. These tools turn and processes turn highly knowledgeable and capable engineers into data entry technicians who then may not have time to actually review analysis results and solve problems. The Dominion Energy Transmission System Operations Center began an initiative in September 2014 to build an automated platform for outage planning analysis. The goal of this initiative was to increase analytic capabilities by 10X to 100X or greater thereby improving (1) the quality of the analytic result and (2) the speed at which the organization can address reliability issues for project work. This effort has not been about simple automation of existing techniques but rather enabling the integration of new analytic techniques in the future through extensibility and modularity. The analysis platform achieved its baseline functionality in May 2018. The platform not only meets the goals set out before initiative took off, but it will provide process improvements, cost savings and improved reliability. This paper summarizes the design, implementation, and current performance of this system at Dominion Energy.

KEYWORDS

Outage planning analysis, power flow, contingency analysis, automation, Python, API, PSSE, TARA

kevin.d.jones@outlook.com

1: Foundations of Outage Planning Analysis at Dominion

Within this paper, *outage planning analysis* refers to the analytic technique/process used by engineers in the Dominion Energy Transmission System Operations Center (SOC) to determine the reliability and security vulnerabilities of the transmission system. These analyses specifically consider planned network outages and future projections of environmental and system conditions. Dominion Energy has historically executed a set of standard outage planning analyses assigned to multiple engineers in the Operations Planning team. These classic, standard protocol analyses include 1-day, 10-day, 1-month, 2-month, and 6-month look-ahead windows.

Within a given study window, an engineer will model and analyze each day sequentially – starting at the first day in the window and proceeding to the last day in the window. Within each day, a set of outages is analyzed. Any outage that began prior to the day in question is considered a *base case outage* and is included in the model created by the engineer. All other outages are considered in the analysis. Each outage is analyzed to determine if the outage will cause any security (thermal, voltage, load-loss, or generation-loss) violations during N-1 contingency analysis. If a violation exists for a given outage, the engineer searches for mitigation steps to resolve the violation. If the violation cannot be resolved through reasonable means, the outage must be rescheduled or cancelled.

Historically, the tools used for this process have been the power flow and contingency analysis network applications contained within the Energy Management System (EMS). When performed with these tools, analyses are performed manually (point-and-click through a graphical user interface) – engineers carefully craft network models and system conditions by slowly changing variables one by one until the desired condition is achieved. The speed at which these studies can be manually completed does not scale to the demands expected in the future. It also does not lend itself well to effective planning; if a study requires weeks to execute, a particular outage scenario may only be reviewed once or twice before it reaches the control room for switching. While advanced EMS network applications have attempted to remedy the manual nature of this process, these tools are incomplete and rigid. This is true for several reasons:

- Data ingest into the EMS is very difficult due to compliance and architecture restrictions.
- EMS tools do not provide a flexible model composition layer for base case creation.
- Their outage planning analytic methodologies do not capture violation mitigation.
- Created models and resulting data are trapped within the confines of the network applications and EMS network and are not transferrable to other toolsets.
- The development of new user interfaces is limited to old, inferior technology.
- The development of any new EMS component is prohibitively costly and time consuming both in terms of vendor cost and integration cost.

Therefore, the EMS is not a suitable environment for an application such as this – any benefits from this kind of integration are far outweighed by the disadvantages. A better solution is required for the future.

Dominion's original vision for this problem's solution was articulated as follows: create an automated system that executes an outage planning analysis upon submission of a new, validated outage request to verify the reliability and security of this newly requested outage – this system would be called *Study-on-Submittal*. As Dominion's understanding of the problem improved, it become clear that *Study-on-Submittal* was not a complete vision but rather a single use case within a much broader strategy. A better approach to this problem was not to build a tool that could analyze outages as they were requested, but rather to generalize the problem and build an analytic platform that could perform any combination or permutation of the standard analysis protocols – for any time period, for any set of outages – and do so in a fraction of the time, thereby enabling an exponential increase in analytic throughput of the Operations Planning team. This platform allows engineers to eliminate the manual tediousness of their roles, to ask more numerous and interesting questions about the planning and reliability of power system operations, and to enable new types of analysis to support changing grid conditions such as increased renewable penetration levels. The net benefit of this is that grid reliability and capital project throughput are increased – both with direct impacts to the bottom line.

2: Problem Formulation, Basic Jargon, and Types of Analysis

As mentioned previous, Dominion began the design with the desire to automate something very specific – *Study-on-Submittal*. However, a general formulation of this problem is required for the extensibility needed to enable new analysis techniques in the future. The basis for this assumption is the level of commonality across the various outage planning analysis protocols.

Potential for Automation - Assuming that all of the required input data and simulation tools are accessible, outage planning analysis is entirely procedural and therefore automatable. In fact, much of the analysis protocol is better executed by a computer than by a human; computers are better at data acquisition and data entry required to build the necessary models and computers are better at executing the optimization algorithms required to search for violation mitigation solutions for outages. This would free the engineer to focus her attention on the results of the study and the actions necessary based on those observations and expert conclusions. Potential actions may include rescheduling or cancelling outages based on the results, verifying a potential violation mitigation solution, or re-running the analysis under updated assumptions and input criteria.

Outage Planning Analysis Base Classes - Previously, several different analysis protocols have been discussed. These included the standard Dominion set of analysis protocols (1-, 10-, 30-, 60-, and 180-day look-ahead analyses) as well as the *Study-on-Submittal* protocol which adopts the *start* and *stop* dates of a single outage of interest. Beyond this, a user would ideally establish the boundaries of an analysis for a custom study by specifying a few key variables. The common input criteria amongst each of the various outage planning analysis protocols can be extracted to arrive at the simplest set of base class protocols. For all combinations of time periods and outages, there are only two fundamental types of analysis. These are, as defined in this paper, the (1) *Time Period Study* and the (2) *Incremental Outage Study*. These two classes of studies vary in both methodology (i.e. which outages are considered in the analysis) and the input criteria.

The first and arguably the more fundamental type of analysis is the *Time Period Study,* depicted din Figure 1, which only requires the user or the computer to specify the desired *start date* and *end date* of the analysis. This is possible because all of the input data required for the analysis (network models, load forecasts, weather forecasts, transmission and generation outages, etc.) can all be set by simply specifying the time period, or *start date* and *end date* of the analysis. During this type of analysis, only the outages which began outside the time period of the study are considered in the base case – all other outages are analyzed. All of the standard protocol studies inherit from this base class of outage planning analysis.

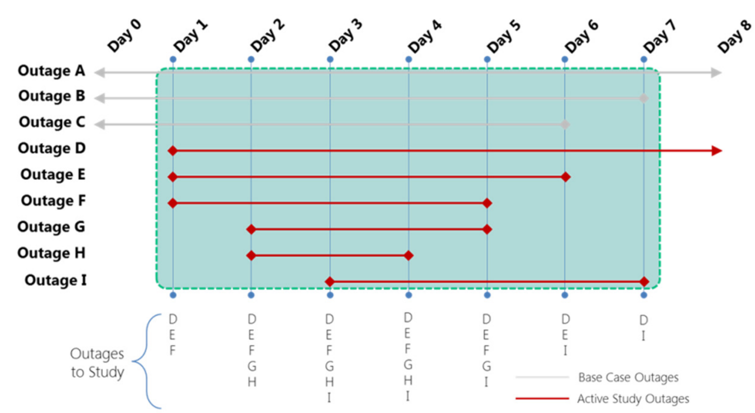

Figure 1: Time Period Study Example

The second type of analysis, depicted in Figure 2, is the *Incremental Outage Study*. There are two key differences between this analysis protocol and the *Time Period Study*. First, the *start date* and *end date* are not directly provided as input. Rather, a single outage is selected as input and its *start date* and *end date* become the time period of the analysis. Second, all outages during the time period of the study are considered in the base cases *except* for the outage specified as input. The purpose of this analysis is to simply evaluate the delta between the system conditions with and without the outage of interest. The *Study-on-Submittal* analysis protocol inherits from this base class of outage planning analysis.

However, given the flexibility of the platform API (discussed in a later section), the input criteria and analysis protocol can be completely specified by the user even if outside the bounds of the two base classes of analysis. Customized studies created through the API underpin the extensibility of the platform. One such example is a naïve implementation of an optimal outage planning analysis. This question is formulated as "When is the best time to take outage X this fall?" This can be translated into a *Time Period Study* with *start* and *end dates* of September 1 and November 30, respectively, and subsequently updating the *start* and *end dates* of outage X to the same dates.

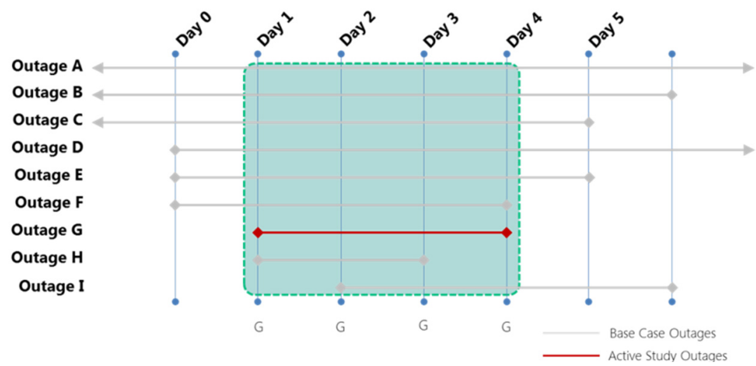

Figure 2: Incremental Outage Study Example

Functional Areas - There are five primary functional areas that need to be addressed for a generic outage planning analysis platform. These include:

1. *Data Acquisition* – The automatic aggregation, persistence, and management of all relevant input data for outage planning analysis. This includes but is not limited to EMS network data, thermal ratings, contingency definitions, load and weather forecasts, transmission and generation outages, and day-ahead market results.
2. *Model Synthesis* – The transformation of raw input data collected through data acquisition into various network models and input files used by the simulation tools
3. *Study Execution* – The handoff of network models and ancillary input files to the third-party simulation tools through their respective APIs or automation capabilities.
4. *Reporting and Visualization* – The post-processing that takes raw output from the third-party simulators and transforms it into human readable reports, displays, and other forms of information to give output and feedback to the user regarding the study.
5. *System Management and Accessibility* – The overall system requires a level of orchestration and ancillary components to complement the engineers' workflow and better extract value from the results.

3: System Design

The following section describes the overall system design including the selection of the technology components, the chosen sources of input data and modeling practices, hardware and software architecture, the API, and the automation techniques. Core design criteria include the following:

1. The system must increase the frequency of the standard outage planning analyses by 10x or greater.
2. The Study-on-Submittal use case must be enabled through end-to-end automation.
3. Customized outage planning analysis (ad-hoc, non-standard) analyses should also be enabled through automation with the integration of an easy-to-use API.
4. The system should be designed as generically as possible to enable additional analytic techniques beyond power-flow and contingency analysis-based outage planning.
5. Technologies, data formats, programming languages should be selected to minimize learning curves for future engineers such that staffing for an in-house solution does not become prohibitive.

The system should enable a flexible, timely execution of standard and ad-hoc analyses to support outage planning and future analytic techniques. It should be composed of as many off-the-shelf tools as possible, integrated through a common, in-house developed automation library.

3.1: Technology Selection

First, this system/platform requires an in-house solution because a fully integrated solution that could meet all of the desired features, data-ingest requirements, and performance requirements does not exist. Others may find themselves considering four alternate approaches:

1. ***Search for an Existing Vendor Solution*** – an existing vendor solution does not exist for this problem. There are siloed technologies which address components of this problem, but no cohesive, flexible, agile integration layer to solve the entire problem holistically. These are the reasons that we did not approach this option.
2. ***Contract an Existing Vendor to Build a Custom Solution*** – Compared to the state of the industry at the beginning of this initiative, today this may be a valid option - with some caveats. However, we did not choose this solution. Instead, we approached the design using a lean/agile methodology that allowed us to discover additional requirements as we proceeded to completely eliminate scope creep. This allowed us to reach a fully functional solution much faster than if we'd spent months or years in an expensive and time-consuming design phase.
3. ***Use the EMS*** – there are no existing EMS tools that completely solve this problem. There are several which appear to address it, but they are incomplete, inflexible, and non-performant.
4. ***Search for a Desktop Application*** – This is a system level problem that cannot be addressed with a simple desktop application.

All of the technology components of the platform were chosen for specific reasons. While these reasons included performance and functionality, additional reasons included customizability, control, transparency, and teach-ability. In several cases, what may appear as an inferior technology in one characteristic was chosen due to its superiority in one of the other categories. Below are some of the core technology components alongside the justification for their selection.

- ***Siemens PSSE*** [1] – the role of this tool is for fine model composition/customization as well as for load and generation loss-analysis (a form of contingency analysis). The benefits of this tool to the platform include its highly flexible Python API and its standard use within Dominion. Drawbacks of this tool include computational speed for lengthier analysis.
- ***PowerGEM TARA*** [2] – the role of this tool is to perform coarse network modeling and bulk power-flow and contingency analysis calculations. This tool was chosen for its computation speed and built-in Outage Reliability Analysis (ORA) module. Drawbacks of this tool include its black-box nature, lack of an API, and terse documentation.
- ***The GE-Alstom EMS*** [3] – this tool was used as the network model data source and as the manual tool used for comparative feedback from the reliability engineers. This tool selection was not optional. The drawbacks of this tool and its environment at Dominion drove the original need for this initiative.
- ***Network Model Exporting with hdbexport*** – rather than export directly into a PSSE readable format or into the Common Information Model (CIM) [4, 5] format, the network model data were exported in plain-text csv. This is a key design decisions for the platform. The reason for this decision is that it grants full control for network modeling and integration of other data sources and provides the transparency required to perform such activities. Direct-to-PSSE and CIM-based exports lack control, transparency, customizability, and most importantly extensibility. CIM also tends to be overly verbose and time consuming to develop adequate parsers.
- ***Windows Task Scheduler and *.bat Files*** –There are many great tools and Python libraries available for task automation [6]. However, the Windows Task Scheduler [7] was chosen for two primary reasons. First, it is incredibly simple to implement and modify compared to more sophisticated solutions. This helps with staffing and training. Second, it is completely acceptable as a solution within the Dominion standard IT policies. More performant approaches would often suffer collisions with Dominion's program whitelist policy as the automation evolves.
- ***HP Superdome*** [8] – This is a high-performance, enterprise server that can scale up to 8 blades of 36 cores each. The benefits and drawbacks of this server hardware are discussed in section *3.3: Architecture* under *Hardware*.

- ***Physically Mounted, Networked HDD*** – An additional hard drive was mounted to the HP Superdome to serve as the shared repository for both server and client operations on the platform. A physically mounted drive will perform better than a true network-attached storage (NAS) device. This is beneficial because the Superdome will be responsible for most of the I/O to the HDD, not the client. To service the clients, the HDD was exposed on the network like a NAS drive.
- ***Python*** [9] – the language of choice for all custom software on this platform is Python. There are two primary reasons for this choice. First and foremost, PSSE has a Python API. Second, but also very important, is that Python is easy to learn and understand. Therefore, it makes it easier for finding and training talent to manage and extend the system in the future.

3.2: Data Aggregation, Model Synthesis, Study Execution, & Reporting

An outage planning analysis requires input from a variety of data sources. Each of these data sources are located in different environments, accessed through different systems and tools, and come in many different data formats. All of this data must be ingested into a common environment by the platform. Once data has been ingested, network models and supporting files that are used by the third-party simulation tools must be synthesized from the raw input data.

Data Sources - Below is a near-exhaustive list of data sources that are used by the platform:

- ***Load Forecast from PJM Oasis*** [10, 11] – The load forecast for the Dominion area is gathered by scraping a text-based web page on Dominion intranet provided by PJM. The resolution of the forecast includes a seven-day forecast with 24-hour resolution.
- ***Weather Forecast from DarkSky.net*** [12] – The weather forecast for the Dominion area is acquired through a free-tiered web service on DarkSky.net. It provides the necessary temperature forecast to establish thermal ratings for up to seven days into the future.
- ***Network Model from EMS*** – The basis for the entire analysis is the network model of the Dominion system. The network model is acquired after every successful state estimator run from the EMS. This provides the required network topology (updated weekly) as well as valid operating conditions (updated every 10 minutes) that are used as a starting point for the model synthesis. Using the network model from the EMS is critical to a successful platform because it maintains a full nodal model of the network. This is required to capture the full detail of contingency analysis. Additionally, the planned outage information maps directly to equipment in the EMS. Therefore, it is also required to properly map outages to the model.
- ***Alternate Limits from EMS*** – The EMS is also the source of all sets of thermal ratings for the analysis. There are 9 separate rating sets across 9 different temperature ranges (32°F, 41°F, 50°F, 59°F, 68°F, 77°F, 86°F, 95°F)
- ***Contingency Definitions from EMS*** – One of the major benefits of working with the EMS data is to take advantage of the nodal model for contingency analysis. As an extension of this benefit, the contingency definitions used by the platform (and thereby the outage planning analysis) are directly modeled as they would be in the EMS and not like a long-term planning contingency set.
- ***Transmission Outages*** – At the center of the outage planning study are the network outages that must be considered. At Dominion this data is created and persisted in an application called iTOA (Transmission Outage Application) [13]. It is a web based application with an Oracle backend. Data is gathered from this system directly from the database through using Python. This data is inherently translatable to the network model data from the EMS.
- ***Generator Outages*** – In addition to the network outages available in iTOA, generator outages need also be considered in the analysis. These outages are provided through downloads from the SDX data sharing service within Dominion's reliability zone. This data shows the outage schedule for each unit and is annotated with MMWG machine names requiring translation into EMS units.
- ***Day Ahead Market Results*** – The day-ahead market results are provided as input to the short-term outage planning analyses as a way to model the desired generation profile. These must be included in addition to the generator outages because (1) a unit may not run even though it is

not on outage and (2) dispatch priority may change based on the day. This information is provided by Dominion's market operator, PJM.

- ***File Based Settings*** – There are additional, slow-to-change input criteria that are critical to the functioning of the platform and outage planning analyses. These are captured as file-based settings and are persisted in a special directory in the study archives (discussed in the following section). These include where to model the swing bus, network voltage limits, long term monthly load forecasts as well as other lookup tables and mapping files to help contextualize input data such as generator outages.

Data Archives - The platform data archives are the structured repositories where all data that is either gathered by or created by the platform is persisted. All of the data sources are automatically aggregated by the platform automation at various refresh rates and saved to their respective data archives on the shared drive. The automated studies execute and create *Study Packages* (see the next section) which are also saved to their respective data archives. The archives not only serve to historize input and output data but also serve as a decoupling mechanism between the various components of the platform. For example, the APIs only query data from the archives and not directly from the data sources themselves.

Study Packages - The *Study Package* concept is one of the most important concepts within the platform. It is a collection of files in a single directory that encapsulates all input and output data for a given outage planning analysis. *Study Packages* are created automatically on the server but can also be created by the user through the platform API. After a study has been completed, the *Study Package* persists in its archive for use and review by the reliability engineers. The API also enables programmatic reconstitution of the *Study Package* so that a user could take an existing study, modifying it, and re-run it to compare the changes. Additionally, an engineer can start the creation of a *Study Package,* walk away from it, and return later to finish because it is persisted on disk and the API enables mid-stream reconstitution of the *Study Package.* Because it is entirely self-contained, it can (1) be moved from one environment to another (server to client or vice-versa), (2) be reconstituted for further analysis, (3) or reviewed historically for process traceability.

Model Synthesis - Input data is gathered synchronously by the platform automation. Downstream from this process, the synthesis of the network models and ancillary input files for study execution must occur. Generally, this only executes on demand (when the API requests the creation of a new *Study Package*). However, the first stage of model synthesis also occurs synchronously as part of the data acquisition process. This first stage is the generation of *seed cases*. The seed case is the *.raw and *.sav PSSE v33 form of the Dominion network model representing the topology and operating condition as it was exported from the EMS. The seed case is directly generated from the CSV output of the hdbexport data from the EMS. Also automatically generated with the seed case are its companion *.con (contingency definitions), *mon (monitored element definitions), and *.sub (subsystem definition) files. These are also based on the data exported from the EMS. When *Study Packages* are created, a seed case and its companion files are automatically and optimally selected from the archive based on the date, time, and loading condition of the snapshot. Now as part of the *Study Package*, the selected seed case must go through several transformations. First, the case is returned to its topological "normal" state – this is called the *all-in case*. Next, a set of derivative cases are created – one for each thermal rating set required for the time period of the study. These are called *altlim cases* or *alternate-limit cases*. After this, dispatch information from the day ahead market results can optionally be used to generate the subsequent set of cases that reflect these desired generation profiles. Finally, the resulting cases are handed off by Python to TARA Automated Module Builder (AMB) to scale loading and apply corrective actions to mitigate any violations that arose in the model synthesis process. This should result in a single *base case* for each day of the study. In addition to the network modeling aspects of the *Study Package* creation, many ancillary input files required by TARA AMB and ORA are also generated. These include control area files, subsystem files, load forecast files, directive files (scripts), and *.bat files. TARA ORA also generates additional network models during its execution that include the series of outages for each day of the study, respectively.

3.3: Architecture

The following section discusses the hardware, software, API, and automation mechanism of the platform.

Hardware - The primary hardware components for the platform consist of (1) a centralized, high-performance server, (2) a large network attached storage (NAS) device, and (3) an array of client workstations.

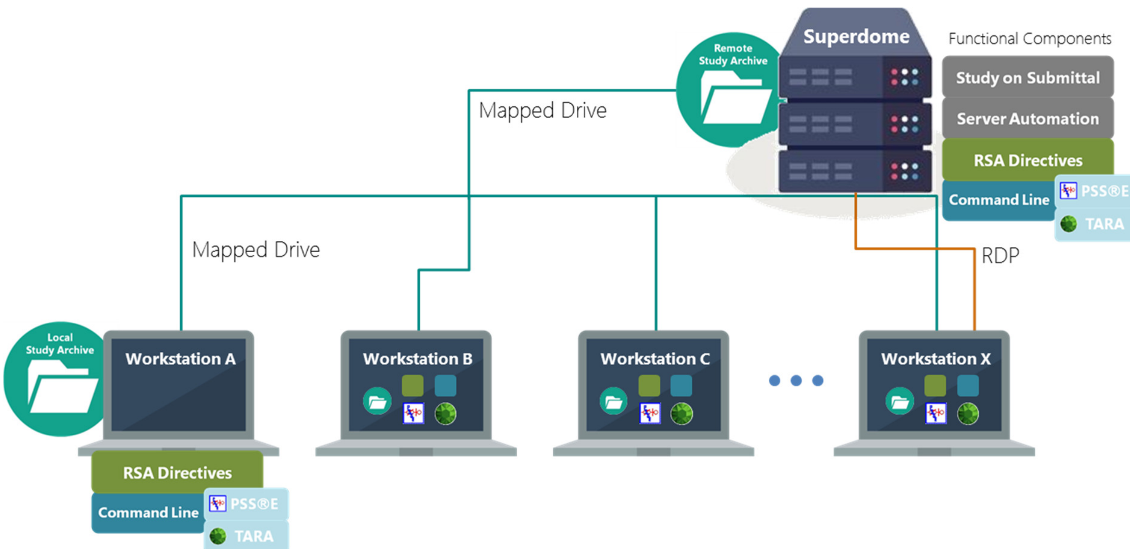

Figure 3: Simplified Platform Hardware Architecture

The roles of both the server and the client workstations are first and foremost as analysis nodes. Each machine contains a copy of the custom platform's Python libraries thus enabling all forms of the analysis to run on the server or the client. In the absence of a web-front end hosted on the server and accessible to the clients, the benefit of study execution on the server is automation and performance whereas the benefit of study execution on the client workstation is analysis customizability and workflow flexibility. This is because the server has much greater compute resources available to it than the client workstations and the users do not directly log onto the server. In the future, clients will have the capability to wield the larger computational resources of the server through the aforementioned web-front-end. The server also plays a dual role in that it is responsible for the automatic aggregation of input data persisted on the NAS. The role of the NAS is simply as the common data repository for the server and the client workstations. All data that is aggregated or generated by the automation (from either the server or a client) resides in this repository.

Alternative hardware approaches for the server potentially include either (1) an array of commodity hardware as opposed to one large, high-performance server or (2) an elastic, cloud-based solution. Dominion Energy opted for a single high-performance machine for several reasons. First, there were other business units at the time that also required a similar computing environment. Second, an array of commodity hardware would have required a custom orchestration layer. Development of this software component would have exceeded the time, budget, and skill constraints of the team at the time. A cloud-based approach was not considered due to data security restrictions, particularly from the EMS, as well as licensing limitations of third party software tools such as PSSE. Basic hardware specifications for the HP Superdome are as follows. Four of eight blades are populated resulting in 144 physical cores and 1TB of RAM. The physically mounted, network exposed drive is 5TB.

Software - All custom software for the platform was implemented in Python. The structure of the platform source code Python package reflects the problem domains within this project (data, modeling, study, reporting) and utilizes a ruthlessly object-oriented design to maximize modularity and conceptual clarity of platform components. Inline documentation generates searchable developer documentation using Sphinx [14, 15].

API - Users of the platform have access to the full capabilities of Python and the study automation libraries installed on the server and client workstations. Any and all functionality is exposed for full customization by the advanced user. However, there is a simplified set of API calls that enable customized yet simple scripts for ad-hoc outage planning analysis.

Automation - There are many technologies available open source for task scheduling. However, the Windows Task Scheduler was used as the primary mechanism for automation due to the desire to minimize the skillset required to maintain and update the platform. The automation mechanism works in the following way. Each task is modeled as a Windows Scheduled Task with its appropriate run schedule and periodicity. Each task is linked to a simple bat file in the platform source code that contains only a single command that is equivalent to "Run the Python script from the same directory of the *.bat file with the same name as the *.bat file". The Python script contains a very terse, high-level script that defines the task to be completed – nearly all of the logic and intelligence is deeper in the platform source code libraries.

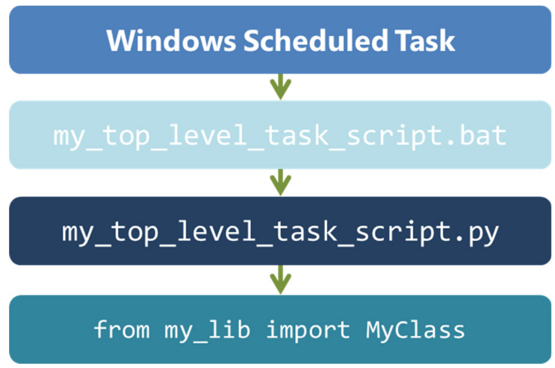

Figure 4: Automation Flow Chart

4: Current Performance and User Traction

In May 2018, the platform began daily operation for the suite of standard analysis protocols. The following analyses are running on the server, producing results for engineers in the Operations Planning team:

Table 1: Automation Periodicity

Analysis Protocol	Periodicity	Average Runtime
Contingency Analysis of Latest EMS Network Snapshot	Every 10 minutes	60-75 seconds
N-1-1 Analysis of Latest EMS Network Snapshot	Every 20 minutes	10-12 minutes
Next Day Study	Every 1 hour	15-20 minutes
Ten Day Study	Every 1 hour	15-20 minutes
One Month Study	Every 1 day	90 minutes
Two Month Study	Every 1 day	90 minutes

The contingency analysis and N-1-1 analysis of the latest EMS snapshots serve as validation feedback loops for the platform. Results of these analyses are compared with results from the real-time EMS system to validate that the model and input data remain of superior quality. Simply put, if these results are consistent with what is observed in the EMS, then Dominion can place a higher degree of trust on the results from other types of analysis.

Initial interactions with the users in Operations Planning have been positive. Results from the outage planning analysis have been observed to be consistent with the results obtained manually from the engineering team. The users have also been active in reporting back to the development team when there have been deviations. In these cases, the underlying issues have been identified and remedied. The next steps for the Operations Planning team are two-fold. First, the engineers will continue to review the platform generated results as they perform their analysis manually. This should happen over an approximately 1 year period to identify any issues and validate the functionality of the platform. However, as engineers begin to see correlation between these two sets of results, they will use them more and more in their daily workflows. The next steps for building greater user traction will be training the engineers on the use of the platform API to enable fully customized studies that can address the more non-uniform investigations they are asked to perform on a daily basis.

Hardware Utilization - The utilization of the server hardware varies substantially throughout a given day. This is because the runtimes of certain analyses are less than the periodicity of the associated run. However, there is a baseline loading of the server that is governed by the execution of the data acquisition, seed case building, contingency analysis, and in particular the N-1-1 analysis. These tasks load the server to approximate 10-11% almost 24/7. When all analysis protocols are running at the same time, the CPU utilization ranges from 60-70%. Therefore, as the team expands the use cases, additional hardware or better orchestration will be required to scale.

Runtime Comparisons - The table below shows runtime comparisons for the standard analysis protocols for traditional, manual execution as well as for the platform automation. Observe that there is a rate increase of greater than 10X for shorter term analyses and greater than 100X for long term analysis. An important aspect of this statistic to note is that while engineers work on a shift or normal business hours and require daily rest and recovery, computers require no such thing. If fully loaded, the server could run jobs for the engineers 24/7/365.

Table 2: Analysis Runtime Comparisons

	Manual Runtime	Platform Runtime	Rate Increase
Next Day Study	4-8 hours	15-20 minutes	**12-32X**
Ten Day Study	8-12 hours	15-20 minutes	**24-48X**
One Month Study	14-21 days	90 minutes	**224-336X**
Two Month Study	14-21 days	90 minutes	**224-336X**

Run Count Comparisons

Performance can also be compared by considering the number of calendar days analyzed in the same time period. This is a proxy for the increase in "coverage" for a particular outage scenario.

Table 3: Calendar Day Coverage Comparisons

	Window Size *(days)*	Manual Rate *(runs per day)*	Platform Rate *(runs per day)*	Manual Total *(days per day)*	Platform Total *(days per day)*	Rate Increase
Next Day	1 day	1 per day	24 per day	1 day per day	24 days per day	**24X**
Ten Day	10 days	1 per day	24 per day	10 days per day	240 days per day	**24X**
One Month	30 days	1 per 30 days	1 per day	1 day per day	30 days per day	**30X**
Two Month	30 days	1 per 30 days	1 per day	1 day per day	30 days per day	**30X**
Total	71 days	N/A	N/A	13 days per day	324 days per day	**24.9X**

5: Next Steps

The purpose of building an *Extensible Platform for Large-Scale, Programmatic Transmission Outage Planning Analysis* goes beyond simply automating existing processes. The true value of this system comes through enabling new types of analysis with existing resources. Therefore, it is interesting and valuable to consider the future use cases that are made possible with this work.

5.1: Additional Analysis Methodologies for Outage Planning

For consideration by the reader, below is a brief synopsis of the new types of analysis that are considered as next steps for this initiative:

Study-on-Submittal – As part of the early vision for platform, studying an outage as soon as a ticket is received provides the maximum opportunity to correct any potential reliability issues before they bottleneck further downstream. In automating generic outage planning analysis, the logic for a Study-on-Submittal approach becomes a thin layer on top of the basic libraries. Data from the transmission outage database is already available to trigger this analysis and the API enables a Study-on-Submittal analysis to follow the same structure as a standard protocol analysis.

Optimal Outage Planning [16, 17, 18, 19, 20, 21] – Due to drastic time reductions in analysis runtimes engineers are enabled to ask more interesting questions. Rather than only questioning, "Does this outage cause reliability issues when it is currently schedule?" they can ask "when can this outage be taken to have minimum effect on reliability?" This kind of optimal outage planning is enabled with this platform and its extensible API.

Stochastic Analysis [22, 23, 24, 25] – One of the key drivers for increasing the dimensionality of the analytic capability of transmission outage planning by 10X/100X is to enable stochastic-based analysis and Monte Carlo simulations rather than considering only a single scenario. The perfect example of this use case is the ability to integrate solar and other renewable energy forecasts as components of the outage analysis.

Dynamic/Transient Stability Constraints for Outage Planning [26] – another benefit of low-cost transmission outage planning analysis is the ability to consider additional constraints in the planning process such as dynamic/transient stability.

Cascading Analysis [27, 28, 29, 30] – even high-dimensionality analyses such as cascading failure studies would now be feasible given the platform data, components, and API.

Synthetic Modeling [31, 32, 33, 34] – with all of the real-time network snapshots and associated loading, weather, and outage data, combined with the results of the outage planning analysis, the team is now outfitted with a data set of high value that can be used synthetic grid modeling for other applications. In particular, this information would be directly transferrable to long-term transmission planning.

5.2: System Level Improvements

In addition to the aforementioned analysis methodologies, Dominion is also working towards some system level improvements to the platform. A sampling of the desired improvements includes:

Developing a Simple Web App – Sophistication levels of users and daily workflows in the Operations Planning team vary substantially. The API provides the ultimate flexibility for customizing an analysis. However, the team will be building a simple web-based application for a customized analysis and reporting workflow. This tool will be used for both creating new studies as well as viewing results from previous studies

Future Technology – Isolated, single-node, desktop applications such as PSSE and TARA prevent easy horizontal scalability. Ideally, there would be a Platform-as-a-Service solution for powerflow and contingency analysis applications with an open and powerful API that are implemented with modern distributed systems technologies. Similarly, large aspects of this platform could be refactored in a similar if a generic data ingest API was designed. Ideally, all of these tools would be open source but managed by a small, nimble company that provides the PaaS to companies for a fee.

6: Conclusion

In order to plan for outages in the transmission system of the future, the dimensionality of the analyses must increase. Dominion Energy Power Deliver, Electric Transmission acknowledged this eventuality in September 2014 by starting an initiative to create an *Extensible Platform for Large Scale, Programmatic Transmission Outage Planning Analysis*. This platform was created in house by a small team of engineers using custom software to aggregate disparate data sources, synthesizing all required models and input files, and integrating all of this with third-party simulation tools. By May 2018, Dominion has demonstrated 10X-100X improvements in analysis runtimes and calendar day coverage and is working towards enabling new types of analysis methodologies and system level improvements for the future.

BIBLIOGRAPHY

[1] Siemens PSSE, https://www.siemens.com/global/en/home/products/energy/services/transmission-distribution-smart-grid/consulting-and-planning/pss-software/pss-e.html
[2] PowerGEM TARA, http://www.power-gem.com/TARA.html
[3] GE-ALstom EMS, https://www.gegridsolutions.com/Software_Solutions/energymanagement.htm
[4] CIM, https://en.wikipedia.org/wiki/Common_Information_Model_(electricity)
[5] CIM, http://www.iec.ch/smartgrid/standards/
[6] Python Task Libraries https://www.fullstackpython.com/task-queues.html
[7] Windows Task Scheduler, https://docs.microsoft.com/en-us/windows/desktop/taskschd/task-scheduler-start-page
[8] HP Superdome, https://www.hpe.com/us/en/servers/superdome.html
[9] Python Itself, https://www.python.org/
[10] PJM, PJM Load Forecast, http://www.pjm.com/
[11] PJM Load Forecast, http://oasis.pjm.com/doc/projload.txt
[12] Darksky.net, https://darksky.net/
[13] iTOA, Sunnet, http://www.sncsw.com/products.html
[14] Spinx, http://www.sphinx-doc.org/en/master/
[15] Sphinx Autodoc, http://www.writethedocs.org/guide/tools/sphinx/
[16] Marwali, M. K. C., and S. M. Shahidehpour. "Integrated generation and transmission maintenance scheduling with network constraints." *IEEE Transactions on Power Systems* 13, no. 3 (1998): 1063-1068.
[17] Anders, George, Gomaa Hamoud, Armando M. Leite da Silva, and Luiz A. da Fonseca Manso. "Optimal outage scheduling—example of application to a large power system." *International journal of electrical power & energy systems* 25, no. 8 (2003): 607-614.
[18] Alam, Aftab, Daniel Johnson, and Abdulrahman Mohammed-Ali. "Integrated optimal outage coordinator: Advancing Operations Planning studies at CAISO." In Power and Energy Society General Meeting (PESGM), 2016, pp. 1-5. IEEE, 2016.
[19] Fu, Yong, Mohammad Shahidehpour, and Zuyi Li. "Security-constrained optimal coordination of generation and transmission maintenance outage scheduling." *IEEE Transactions on Power Systems* 22, no. 3 (2007): 1302-1313.
[20] Fu, Yong, Mohammad Shahidehpour, and Zuyi Li. "Security-constrained optimal coordination of generation and transmission maintenance outage scheduling." *IEEE Transactions on Power Systems* 22, no. 3 (2007): 1302-1313.
[21] Liu, Jenny, Mostafa Kazemi, Amir Motamedi, Hamidreza Zareipour, and Jim Rippon. "Security-Constrained Optimal Scheduling of Transmission Outages with Load Curtailment." *IEEE Transactions on Power Systems* 33, no. 1 (2018): 921-931.

[22] Silva, E. L., M. Morozowski, L. G. S. Fonseca, G. C. Oliveira, A. C. G. Melo, and J. C. O. Mello. "Transmission constrained maintenance scheduling of generating units: a stochastic programming approach." *IEEE Transactions on Power Systems* 10, no. 2 (1995): 695-701.
[23] Wu, Lei, Mohammad Shahidehpour, and Yong Fu. "Security-constrained generation and transmission outage scheduling with uncertainties." *IEEE Transactions on Power Systems* 25, no. 3 (2010): 1674-1685.
[24] Celli, Gianni, Emilio Ghiani, Fabrizio Pilo, and Gian Giuseppe Soma. "Reliability assessment in smart distribution networks." *Electric Power Systems Research* 104 (2013): 164-175.
[25] Adefarati, T., and R. C. Bansal. "Reliability assessment of distribution system with the integration of renewable distributed generation." *Applied energy* 185 (2017): 158-171.
[26] Chen, Luonan, Y. Taka, Hiroshi Okamoto, Ryuya Tanabe, and Asako Ono. "Optimal operation solutions of power systems with transient stability constraints." *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications* 48, no. 3 (2001): 327-339.
[27] Hardiman, Rodney C., Murali Kumbale, and Yuri V. Makarov. "An advanced tool for analyzing multiple cascading failures." In *Probabilistic Methods Applied to Power Systems, 2004 International Conference on*, pp. 629-634. IEEE, 2004.
[28] Baldick, Ross, Badrul Chowdhury, Ian Dobson, Zhaoyang Dong, Bei Gou, David Hawkins, Henry Huang et al. "Initial review of methods for cascading failure analysis in electric power transmission systems IEEE PES CAMS task force on understanding, prediction, mitigation and restoration of cascading failures." In *Power and Energy Society General Meeting-Conversion and Delivery of Electrical Energy in the 21st Century, 2008 IEEE*, pp. 1-8. IEEE, 2008
[29] Vaiman, Marianna, Keith Bell, Yousu Chen, Badrul Chowdhury, Ian Dobson, Paul Hines, Milorad Papic, Stephen Miller, and Pei Zhang. "Risk assessment of cascading outages: Methodologies and challenges." *IEEE Transactions on Power Systems* 27, no. 2 (2012): 631.
[30] Yan, Jun, Yihai Zhu, Haibo He, and Yan Sun. "Multi-contingency cascading analysis of smart grid based on self-organizing map." *IEEE Transactions on Information Forensics and Security* 8, no. 4 (2013): 646-656.
[31] Wang, Zhifang, Anna Scaglione, and Robert J. Thomas. "Generating statistically correct random topologies for testing smart grid communication and control networks." *IEEE transactions on Smart Grid* 1, no. 1 (2010): 28-39.
[32] Birchfield, Adam B., Eran Schweitzer, Mir Hadi Athari, Ti Xu, Thomas J. Overbye, Anna Scaglione, and Zhifang Wang. "A Metric-Based Validation Process to Assess the Realism of Synthetic Power Grids." *Energies* 10, no. 8 (2017): 1233.
[33] Xu, Ti, Adam B. Birchfield, Kathleen M. Gegner, Komal S. Shetye, and Thomas J. Overbye. "Application of large-scale synthetic power system models for energy economic studies." In *Proceedings of the 50th Hawaii International Conference on System Sciences*. 2017.
[34] Birchfield, Adam B., Ti Xu, Kathleen M. Gegner, Komal S. Shetye, and Thomas J. Overbye. "Grid structural characteristics as validation criteria for synthetic networks." *IEEE Transactions on power systems* 32, no. 4 (2017): 3258-3265.