28 NREL Stratus - Enabling Workflows to Fuse Data Streams, Modeling, Simulation, and Machine Learning

Preprint

David Rager and Aaron Andersen

National Renewable Energy Laboratory

Presented at the Smokey Mountains Computational Sciences and Engineering Conference
Kingsport, Tennessee
August 24-26, 2021

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-2C00-80214
November 2022

28 NREL Stratus - Enabling Workflows to Fuse Data Streams, Modeling, Simulation, and Machine Learning

Preprint

David Rager and Aaron Andersen

National Renewable Energy Laboratory

Suggested Citation
Rager, David and Aaron Andersen. 2022. *28 NREL Stratus - Enabling Workflows to Fuse Data Streams, Modeling, Simulation, and Machine Learning: Preprint*. Golden, CO: National Renewable Energy Laboratory. NREL/CP-2C00-80214. https://www.nrel.gov/docs/fy23osti/80214.pdf.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-2C00-80214
November 2022

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the National Renewable Energy Laboratory's Computational Science Center. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via www.OSTI.gov.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

NREL prints on paper that contains recycled content.

28 NREL Stratus - Enabling Workflows to Fuse Data Streams, Modeling, Simulation, and Machine Learning

David Rager[1] and Aaron Andersen[2]

[1]NREL, Golden, CO USA, David.Rager@nrel.gov
[2]NREL, Golden, CO USA, Aaron.Andersen@nrel.gov

Abstract. Integrating cloud services into advanced computing facilities provides significant new capabilities and offers several advantages over focusing solely on traditional high performance computing (HPC) workloads. The integration of cloud services is especially potent for workflows that fuse data streams, modeling and simulation ("modsim"), and machine learning. A key challenge to adopting a hybrid edge-cloud-HPC model is aligning optimal capability, data, and user intent on the right resources for each step in a workflow. The National Renewable Energy Laboratory (NREL) Stratus service provides a basis for this alignment. Stratus layers the capabilities needed to make cloud services accessible to a lab-based scientific community on commercial offerings, and currently supports upward of 200 projects, ranging from Internet of Things (IoT) integration to traditional modeling and simulation. This provides a real-world inventory of scientific workflow elements, which enables placing these elements appropriately between the edge, cloud, and traditional HPC. This paper outlines a vision via reference architecture and the application of that architecture in a typical workflow. We highlight multiple components, including sensor data intake, cleaning, and transforming (edge/cloud suitable), generation of synthetic data through modsim, computationally heavy machine learning training and hyperparameter optimization (HPC suitable), and inference and deployment (cloud ideal). Every step in such a workflow involves a cost-benefit analysis of the data movement, computational efficiency, availability, latency, and resource capabilities. The reference architecture and examples outlined in this paper allow for better understanding of new opportunities in the context of emerging workflows that combine IOT, cloud, and HPC to bolster scientific productivity.

Keywords: HPC, cloud, computing, edge, data streams, modeling simulation, machine learning, NREL

1 Introduction

Service cloud providers bring a unique set of capabilities when integrated with traditional high performance computing (HPC) to form large, dynamic, and highly effective computing ecosystems for some of the most challenging research problems. Cloud computing offers, and has successfully commoditized, highly reliable services for receiving data streams from the edge, which is essential for systems dependent on Internet of Things (IoT) devices. These services can be leveraged to optimize throughput, cost, and reliability for complex distributed ecosystems that are not well serviced by traditional scientific computing alone. However, the cloud is not cost competitive for running large, computationally intensive workloads. HPC provides access to highly efficient computing resources for large, computationally intensive workloads. HPC systems typically include single points of failure, and are often heavily subscribed, resulting in scheduling contention, which is not an appropriate basis for servicing real-time or near real-time data systems. However, a hybrid of HPC and cloud computing systems provides a unique basis for handling large scale computational workflows efficiently

At the National Renewable Energy Laboratory (NREL), the Stratus cloud computing environment, which is provided by Amazon Web Services (AWS) enables a wide array of tools and services that work in concert with the Eagle supercomputer in the HPC Data Center. Some use cases include the generation of hundreds of terabytes of data on Eagle, which is transferred to Stratus for data distribution and analysis. In other cases, Stratus manages queues or provides the source data for larger modeling jobs. NREL is at the forefront of building the computational ecosystem needed to manage grid-connected or community-based field studies for large scale energy security or energy efficiency studies. These use cases do not fit the patterns of either HPC or cloud computing alone, so further effort is needed to bring this wealth of resources in alignment to optimize these complex workflows.

2 Advanced Computing at the National Renewable Energy Laboratory

NREL has merged HPC and cloud computing together under a single banner of "advanced computing" to address the compelling need to provide researchers with access to the tools best aligned for the work at hand. Traditionally, cloud computing was focused on outreach and knowledge transfer, whereas HPC did the heavy lifting for research. Over time, however, cloud computing at NREL has become an increasingly valued resource for the laboratory. More analysis and

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

computational workloads are being deployed into the Stratus environment—substantially expanding the role cloud computing now plays at the laboratory—providing tools not previously available to researchers and analysts. As laboratory computing requirements and demands change, NREL's vision for the future of computing is receiving greater focus. The reference architecture and use cases presented here elaborate on NREL's vision for advanced computing [9].

In an idealized scenario, the advantages of agility, performance, and cost- or risk-avoidance would determine which computing environment to leverage. However, optimization and application of complex systems are far more dependent on expertise, experience, and continuity at the leading edge. In order to apply this type of vision to a strategic alignment of resources, it is useful to understand the systems' capabilities, apply them to different workflow scenarios that NREL's Advanced Computing Operations team is currently building support for, and invest in developing supporting expertise.

2.1 ESIF High Performance Computing Data Center

The Energy Systems Integration Facility's (ESIF's) HPC data center, which NREL has operated since 2013, demonstrates best practices in data center sustainability and serves as an exemplar for the community. The facility utilizes a holistic systems approach to the data center, focusing on three critical aspects of data center sustainability:

- Efficiently cooling the IT equipment using direct, component-level liquid cooling with a power usage effectiveness (PUE) design target of 1.06 or better
- Capturing and re-using the waste heat produced
- Minimizing the water used as part of the cooling process.

The Eagle supercomputer, in production since 2019, enhances NREL's leadership role in addressing the nation's renewable energy concerns. Eagle is an 8.9 petaflop Hewlett Packard Enterprise SGI 8600 system with Intel Skylake processors; 2,114 compute nodes; 296 terabytes of total memory; 14 petabytes of high-speed data storage; and a hypercube based 100 Gb/s Enhanced Data Rate (EDR) Infiniband network. Eagle is installed in the ESIF on NREL's campus—which is highly instrumented—and the facility, computing, and built environment are an experiment in energy usage and optimization, currently achieving regular PUE of 1.03 or better.

2.2 NREL Stratus Cloud Computing

The NREL Stratus cloud computing environment was launched over 10 years ago, when cloud computing was nascent. As an early adopter of cloud computing, NREL now provides expertise that continues to benefit a wide array of projects that support NREL's mission [9]. As cloud computing matures, the range of specialized, highly available and scalable services increases, providing compelling functionality that can be applied to specific use cases and needs. NREL's Stratus environment and the Computational Science Center's Cloud Operations Team (CSC Cloud Team) invest time in understanding these services, then help apply them to the needs of researchers throughout the lab. These needs include receipt and processing of data from field studies or equipment outside of the laboratory, big data analysis, publishing websites, and publishing web applications used by industry or the public. Currently, Stratus supports:

- Analysis on large modsim projects
- Data collection from edge devices for
    - field studies
    - power generation experiments
- Data processing workflows
- Public web applications
- Publication of large open data sets.

Recent growth trends in the Stratus environment are indicative of successful implementations in big data analytics utilizing data warehousing and big data management tools; muli-schedule containerized applications, deploying Docker containers at the edge or in serverless functions; and in IoT support for field experiments and grid management studies.

3 Competitive Positioning

HPC, cloud, and edge environments have differing value propositions, cost models, and staffing requirements. HPC traditionally excels at numerically intensive, parallelized, and tightly coupled simulations, with large fixed costs and low

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

variable costs. Stratus cloud computing is composed of specialized services that provide application support and a framework for management and governance, as well as commodity computing and storage services. The edge has developed into an ideal target environment for computing, where very low latency controls, on-site alerts, and machine learning decision support are required. Cloud providers have multiple services to help with edge computing management and receipt of data. NREL is continuing to build on its vision for ideal synergies between HPC and cloud environments to provide efficient and reliable support for our researchers while lowering the barrier for ambitious computing and data intensive research projects. Jobs that require large core counts for hours or days are most efficiently served by Eagle, while jobs with small core counts of high intensity and short duration are often most cost effective when served by the cloud. NREL continues to research both cost and performance optimization for supporting traditional HPC workloads with Eagle and targeted support in the cloud. Recurring scheduled work, analysis, and highly available management of near real-time data streams are typically handled with cloud computing.

3.1 HPC

Eagle, like many HPC systems, can be viewed as a very large, high performance, optimized compute engine that can be molded to fit different workloads. It excels at numerically intensive, parallelized, and tightly coupled simulations. With years of supercomputing experience, the research staff at NREL typically have a higher comfort level with building and scheduling jobs for supercomputing schedulers. However, NREL also has a small but growing staff of highly proficient cloud computing developers, but the overwhelming expertise is with HPC.

Operating an HPC class data center involves significant expertise, labor, and capital investment, not only in the compute resources, but also in power infrastructure, networks, and the power dense data centers and buildings that host the compute resources. The end result is an optimized computing infrastructure that provides a very low cost per floating-point operation (flop) to run compute intensive workloads. As capabilities advance during the lifetime of an HPC facility, these advancements are incorporated into the supercomputing infrastructure with incremental capital investments.

3.2 Cloud Computing

Stratus cloud computing can largely be viewed as a scalable set of reliable microservices, each tuned to process a specific workload, that are accessible via highly available interfaces with a wide array of configurations available. Internally, cloud implementations are approached using best practices for software development. Infrastructure is leveraged as code to standardize how that infrastructure is defined in Python, and then best practices are applied through to launch, via agile software development processes. NREL's CSC Cloud Team manages the Stratus cloud computing enclave, security, application operations, and much of the education for use of the cloud at the laboratory. Pricing of resources in the cloud spans a wide range, depending on what services a customer can commit to using. The published on-demand pricing in AWS can be substantially reduced with pre-purchases or use of excess computing inventory, which AWS makes available as spot instances. The pre-purchased compute instances, or "reserved instances," can provide roughly a 40% savings over on-demand prices, with spot instances providing up to an 80% savings. Spot instances come with the agreement that the excess capacity can still be sold to other customers by AWS at higher prices, and removed from service with short notice, making the reliability of spot instances a challenge [12, 14]. This results in considerable planning complexity. The CSC Cloud Team assists NREL researchers in managing this complexity by working to leverage the most value for the laboratory from cloud services.

The approach to computing in Stratus differs from traditional HPC, as Stratus capabilities generally target a wide array of application support services. The cloud also provides a strong basis for analysis, serving near real-time data with highly reliable, 24/7/365 availability [7]. As an example, AWS provides 12 different database management system services for online analytical processing (OLAP) or online transactional processing (OLTP), and has an ecosystem that comprehensively supports big data management and analysis. The CSC Cloud Team introduced big data capabilities by leveraging cloud services, scaling from terabyte to petabyte datasets for researchers at NREL. This capability became a key enabler for teams to identify patterns of use and to develop cloud capabilities in order to move large datasets that in the past required HPC. Those projects can now be curated within the cloud environment.

Another area where Stratus provides complimentary services to the laboratory (that are not appropriately served in the HPC data center) is through support for IoT environments [3]. The Stratus cloud environment leverages the following advantages for IoT support [8]:

- Highly available endpoints used to receive device data delivered to highly available data streams, which auto-scale to meet variable load demands at minimal cost

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

- Transparent high availability across multiple cloud provider data centers in cases of equipment failure or other communication issues at the receiving end
- Highly available data stores for analysis that have scaling capabilities available to serve highly dynamic or growing workloads
- Highly available data pipelines and workflow systems that are capable of reliably scheduling machine learning pipelines, data processing, and transformations using scalable compute resources.

Cloud providers continue to rapidly expand their capabilities to better support computing at the edge. Recent service additions include placing AWS managed compute into third party data centers enhancing broad geographic coverage, including topologies running inside cellular networks. These computing installations allow research teams to compute in specified geographic or network locations meeting low latency requirements, while utilizing common cloud computing interfaces and support for the deployed infrastructure regardless of where it is located.

To support edge computing requirements in remote locations where cloud providers do not operate, AWS provides software solutions such as GreenGrass or the Real-time operating system for microcontrollers (FreeRTOS) that run on edge devices or gateways. These provide some basic capabilities for the edge systems to leverage and secure connections to highly available endpoints, facilitating transmission of data to the cloud for processing and analysis.

3.3 Edge Computing

Edge computing involves deploying computing resources geographically or topologically close to the IoT devices being monitored and controlled. Edge computing is often an important consideration for applications that have stringent real-time or data reduction requirements at the edge. In most cases, operating IoT environments at the edge have limiting compute capabilities or other constraints, such as space constraints, barriers to installing additional equipment, and prohibitive costs of supporting remote resources. Because of these constraints, latency requirements may need to be relaxed and mitigated so that equipment can be placed in higher latency locations [6, T13].

Edge computing brings a highly dynamic environment with challenges around fault tolerance that are not often encountered in the well-managed HPC and cloud infrastructure environments [5]. Decisions are required when deploying IoT resources in edge environments; these include adoption of proprietary management or data handling systems; appropriate environmental controls; methods for remote management; data retention and collection requirements; investments in network availability; and power management requirements. In many cases, these decisions are dictated by the working environment. In most cases, multiple single points of failure should be expected, which impact operations or communications in remote devices.

4 Hybrid Support of Real-Time Data Vision at Scale

To optimally support the growing demands and new capabilities needed to meet the U.S. government's ambitious renewable energy goal of creating a carbon-pollution-free power sector by 2035 [15], NREL envisions a hybrid environment that optimizes the use of edge, cloud, and HPC computing. This real-time data vision supports several efforts that will be discussed in more detail in the following sections, including data generation/intake, data cleaning, machine learning, and machine learning training.

4.1. Overview of the Workflow

The workflow in Fig. 1 provides a reference architecture for an IoT environment, illustrating the application of a well-supported hybrid ecosystem composed of edge, cloud and HPC components. In this case, the combined capabilities of HPC and cloud are leveraged to host highly available components used to manage IoT devices and the data streams they produce. This architecture supports both very large compute workloads as well as small scheduled workloads within a secured enclave.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

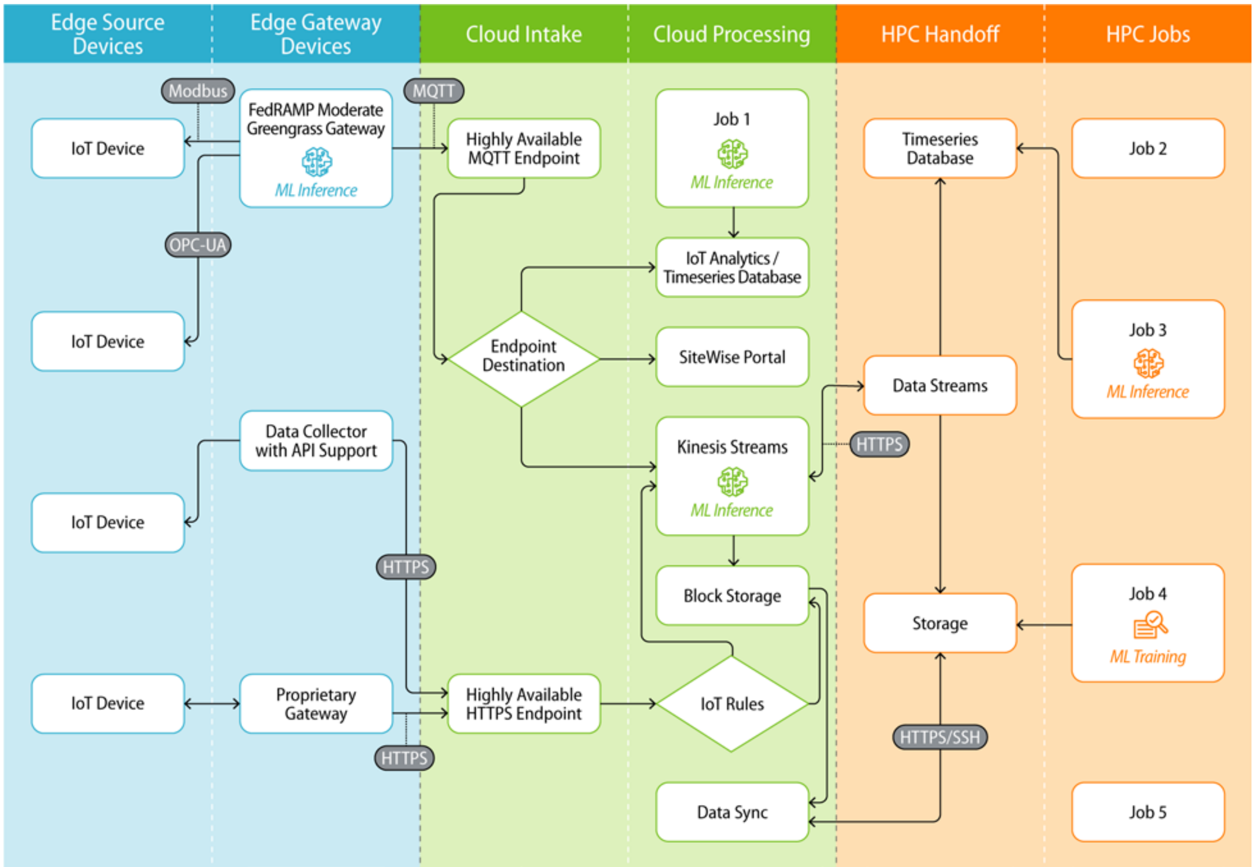

Fig. 1. A reference architecture with edge, cloud, and HPC components.

To best illustrate NREL's vision of an optimized edge, cloud, and HPC ecosystem, we present a theoretical architecture of field studies and large scale energy modeling that leverages the reference architecture and workflow. For the edge, the architecture should support experimental and field studies for the operation and optimization of smart devices in a residential neighborhood, or a "smart neighborhood." Goals include reducing energy consumption and supporting the use of renewable energy on the existing energy grid infrastructure with no or minimal negative impact on the residents. The architecture will support large scale modeling of simulated data ("modsim"), which then, via simulations, propagates the effects of smart neighborhood management across a larger scale—in particular, the surrounding areas affected by changes to power grid management and power consumption. In turn, those effects affect power generation and transmission schemes for typical residential, industrial, and commercial uses of energy. At both ends of this interaction, we will analyze experiments, trends, and effects to best understand how to optimize the larger ecosystem.

Edge devices are part of the system being optimized, monitored, and controlled. At the edge, some of the primary capabilities to be deployed or leveraged include:

- Very low latencies between edge devices and edge operational systems
- Smart aggregation of device data, leveraged to reduce the amount of data the overall ecosystem needs to manage
- Low latency machine learning inference, applied to device or data management
- Low latency data inspection and stream management
- Low latency controls of devices at the edge.

At the edge, some of the primary concerns include:

- Highly constrained compute capabilities, which must be specified to support the workload requirements expected at the edge
- Data segmented by physical location or connected devices, providing a segmented view of activity at the edge

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

- Intermittent network connectivity issues, with multiple single points of failure in line for devices that want to reach out via the internet or cellular networks
- Higher operational and maintenance response times for equipment failures
- Lower levels of security controls around edge devices.

Cloud resources in this scenario represent a highly available collector and processor of data, providing remote control and device management. In this scenario, the primary capabilities of cloud resources are:

- Highly available and secure MQTT or HTTPS endpoints for securely receiving IoT data
- Quickly adjustable compute scale for volatile or batch workloads
- Reliable and secure methods of sending controls instructions to edge devices
- Reliable and secure methods of deploying code or machine learning models to edge devices
- An array of data streaming, data storage, and data analysis tools.

Some of the primary concerns with using cloud resources are as follows:

- Cloud computing introduces complexity into any ecosystem by following a microservices architecture.
- Variable costs introduce risk and unpredictability.
- Complex deployments can introduce security risks.
- Data transfer costs and challenges can occur between the HPC data center and the cloud.

HPC resources in this scenario provide the tightly coupled parallel processing capabilities that require prolonged computation. In the HPC data center, the primary capabilities to be leveraged include:

- Efficient, high performance large scale compute job support
- Large scale modeling support, machine learning training, and inference support
- The capability to work with and generate very large datasets.

Some of the primary concerns with using HPC resources are as follows:

- Work must be scheduled well in advance of run dates.
- Scheduling contention may limit resources available to run the HPC jobs.
- Operational outages may interrupt planned scheduling.
- Wall time for job completion could potentially shut down jobs that have not fully completed.

4.2 Workflow Components and Positioning

The workflow that supports the smart neighborhood and large scale modsim work will be a synthesis of the following components:

- Improving smart controls of devices in the homes within smart neighborhoods
- Measuring satisfaction of smart neighborhood residents
- Measuring changes to energy usage and demand in the smart neighborhood
- Simulating the effects of efficiently run smart neighborhoods on large scale energy consumption and production trends.

The goals of the work include:

- Optimizing energy usage for current and future power generation infrastructure
- Improving the overall energy usage profile for adoption of more renewable energy resources, which are naturally less predictable.

Data Generation and Intake. In this overarching workflow, the real-time data of interest is delivered from smart neighborhoods. These data consist of timeseries from power meters, thermostats, or other sensors. Additionally, state information from different smart devices, such as appliances, heaters, washers, or even solar power distribution infrastructure, is collected. Many smart devices also emit events tied to their status or to sensor readings, such as occupancy sensor events or temperature limit violation events, that need to be collected. In addition to these real-time data streams, historical data on energy usage for the area would be identified and collected.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

To support the modsim work, the smart neighborhood data, and an aggregate representation of the data from the surrounding areas, needs to be available to the HPC data center. In addition, the operators of the smart neighborhood experiment, who are attempting to optimize the controls for the neighborhood within the larger footprint, will want to issue controls by reacting to data in the real-time data streams. The reactions may require low to modest latencies (milliseconds to minutes). At the millisecond level, device alerts could identify possible safety hazards. In other cases, near real-time data streams and latencies of minutes or more may be acceptable—for example, when controlling a schedule for a washer based on current or near-term predicted energy use or cost trends. As such, we need multiple mechanisms in place:

- Edge device consumers of very low latency data to send alerts or activate controls
- Highly available cloud-based consumers of data that feed into decision support for controls informed by larger ecosystem datasets, and supplemental datasets collected form smart device operators
- Data collectors at the edge or in the cloud that aggregate data at a sufficient frequency to then deliver, or make available to HPC workloads, before they are scheduled to run.

The pairing of edge IoT devices and gateways with cloud-based, highly available endpoints for data collection provides a robust mechanism for handling data collection. This pairing works around the low storage and compute capabilities in the smart neighborhood. It also provides researchers with live access to the data streams [6]. Because the data streams are live, the HPC data center is not used to poll the data or publish endpoints for device data collection. Its primary function is to perform batch work efficiently, without the overhead of providing highly available endpoints. Thus, we instead rely on the cloud for this architectural support. Details of the integration of this type of workflow are shown in Fig. 2.

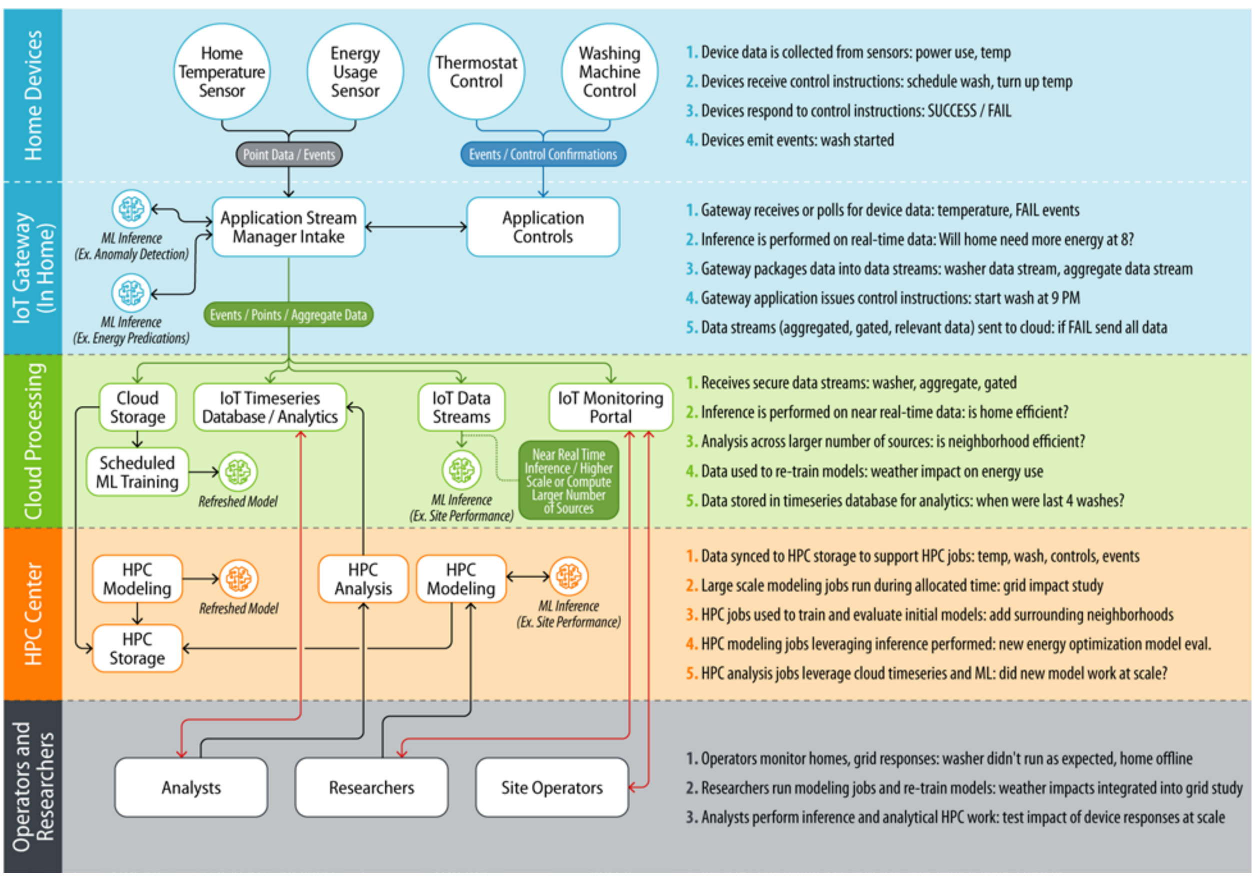

Fig. 2. Workflow illustration for data generation and intake based on a smart neighborhoods example.

Data Cleaning/Transformation. Data cleaning and transformation are a core component of any data processing system, especially when live device data from a disparate range of sources is brought together to form the basis for analysis and decision support. Typically, devices from multiple manufacturers are integrated, and as a result, the data is subject to varying signals, units, metric definitions, and proprietary information, which take some effort to homogenize and make useful. High frequency data may not be needed at all points of analysis and decision support, so aggregation points should be introduced to reduce the

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

cost of data management tasks throughout the system. Aggregation as early as possible in the data flow lowers the impact of costs and complexity. Care does need to be taken, as aggregation can introduce an element of risk into the effectiveness of the analysis and decision support layers as information is lost.

In our scenario, the following data cleaning and transformation items could be configured at or near the edge:

- Simple checks on data quality for expected values
- Simple but smart aggregations to reduce data being operated on in the downstream systems
- Minimum required sample rates at the edge to reduce data being operated on in the downstream systems
- Replacement datasets for missing or corrupted data.

For simpler data manipulation, gateways are provided in line at the edge, as this is the earliest point in data processing. The gateways can assist in providing usable data for decision support at the edge. These devices also provide a software platform for the integration of machine learning or configurable algorithms, which operate on the devices. The options for how to place this software vary from location to location. Ideally, if there is some type of physical collection hub in the neighborhood that is being monitored, a device would be selected or added to handle the gateway software at this location. If that is not available, the next best options should be considered. The following options can help place the gateway system closer to the edge components [16]:

- Smart mobile devices – gateway software that can run on devices at the edge and perform inference or apply custom code/algorithms to data or processes
- Mobile cloud computing – installations of cloud infrastructure for customer use at third party data centers, providing a physically closer computing system to host the gateway
- Mobile edge computing – installations of cloud infrastructure within cellular networks, improving the network topological proximity to the edge systems when cellular networks are used to transmit data
- Cloud computing regions – the cloud provider’s primary data centers, which are distributed across different physical regions; these can improve both physical and topological proximity by selecting the geographical region closest to the edge devices.

In all but the first case, the computing power provided can be significant enough for sophisticated processing and data management, so the requirements for data cleaning and transformation may be expanded. In the case of smart mobile devices, however, the compute power relies on what can be installed or leveraged in the smart neighborhood. To help overcome the limitations of managing edge computing, the AWS-provided GreenGrass mobile device software provides the following flexibility:

- Data streams management, including a cache of data per stream when outbound communications issues arise, preventing data from being handed off to the cloud
- A deployment target for custom gateway code that can be deployed from the cloud
- A deployment target for machine learning inference models that can be deployed from the cloud
- FedRAMP moderate certification, with secure communications back to AWS IoT services [1]
- Multiple methods to deliver data to the cloud for different scenarios, such as operations monitoring, archival storage of data, data streaming environments, and analysis environments.

Often, raw data is retained in a compressed form as a means of capturing the source of truth for the system. While this may result in duplication of captured data in the cloud, the data is handled in a less costly manner, typically being sent directly to block storage as an archive.

Synthetic Data/Modeling. The larger modeling efforts provide a basis for making strategic decisions on energy management that incorporate the future development of smart neighborhoods. This would typically be managed in the HPC environment, where performance and efficiency are maximized, supporting larger scale work at lower costs. The data collected from the edge is a critical input for tuning the larger strategic model. The model itself is used to simulate data, which becomes the basis for the analysis efforts used to determine efficiency gains, strategic direction for energy production, priorities for smart neighborhood capabilities or incentives, and how changes to elements in the larger ecosystem would propagate.

The larger, coordinated modeling effort combines the specialized skills of different groups. These groups generate their simulated datasets, which may then become a basis for the integrated work of other groups. In this environment, compute time is scheduled, and critical paths are developed. Once a data simulation effort is completed, that data is typically transferred up to the cloud for analysis purposes. It is also made available to other HPC jobs or ad hoc/scripted work efforts from the laboratory via big data storage and analysis tools in the cloud.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Machine Learning Inference. There are many opportunities for integrating machine learning capabilities into this workflow, taking into account the scope of the effort, the devices being controlled, the types of data being received from the field studies, and the amount of data being generated via HPC jobs. Opportunities for machine learning include supporting smart controls, data cleaning, data validation, anomaly detection, data simulation, predicting energy use, and scenario simulations. As with other processing work, the placement of inference endpoints and management of the models behind them needs to be determined based on each use case. On this topic, the questions we need to consider include:

- Who is the consumer of the inference endpoint?
- Are there low latency requirements for the inference endpoint?
- Are there availability requirements for the inference endpoint?
- What are the compute requirements for the inference endpoints?
- Are the compute requirements volatile or predictable?

One example of machine learning is anomaly detection at the edge for smart neighborhood field studies. The use of the generated anomalous behavior events will help in determining where the inference is performed. If the behavior is used to clean data before it is sent to the cloud, or if there is a notification mechanism at the edge that is intended to immediately alert people on site, then inference should be deployed to the edge gateway that is collecting the data for transmission to the cloud. In this case, the consumer who is notified of the event is at the edge, there are low latency requirements, the possibility of loss of network availability to the cloud interrupting notification of the event is too risky, and the event, while not predictable, would likely only need low computing requirements to handle, so a gateway device at the edge makes sense.

In these cases, by leveraging the Stratus environment, GreenGrass would be installed on the edge gateway, and AWS IoT deployments would be used to deploy the machine learning model and the code that performs inference on the related incoming data to the gateway [2]. Although the processing requirements go up with the addition of machine learning inference, it is likely that they would fall in line after the edge data reduction efforts and just prior to transmission to the cloud, limiting the amount of data being processed through inference. Depending on the scope of data collection and the number of smart neighborhoods being evaluated, the overall ecosystem may include multiple gateway devices that need to be updated. AWS's IoT Core allows Stratus operators to define a class of devices to receive the model updates as they are released, so that deployments are centrally managed for multiple gateways.

Another use case for machine learning at the edge is handling incorrect or missing data and performing corrective action. One example of this is via use of multiple weather stations at a particular location. It is reasonable to assume that temperature changes would result in somewhat predictable patterns of power usage, such that if two outdoor thermometers' temperature readings significantly disagreed on what the current temperature was in a smart neighborhood, inference on the incoming sensor, modeled in part on meter data, would likely be able to decide which thermometer was correct for use. This has the effect of cleaning the data used and recorded by the active systems and creating an event indicating that the other sensor may need service.

If anomalous behavior is instead based on aggregate data streams or on activity from multiple sources, then the cloud is typically an appropriate host for the inference endpoints. Data from multiple sources at larger aggregate throughputs will be available in the cloud. One key benefit of data processing in the cloud is the ability to implement auto-scaling to meet variable demands at minimal cost. Multiple streams from disparate sources can similarly be processed via a centrally managed endpoint, which is much simpler to update, monitor, and configure. In the Stratus cloud environment, inbound data from the edge is received via IoT Core endpoints, resulting in data streams that are configured in a manner that enables sending the data immediately through an inference endpoint.

To support the modsim evolution, machine learning is included to (1) assist with modeling future behaviors and the effects of changes to the overall ecosystem, (2) detect anomalous data in the simulated datasets, and (3) assist with the creation of simulated data used for analysis. Large scale effects in the simulations would result in recreating the detailed simulated data for the entire ecosystem, so the machine learning environment must be tuned to be responsive while data generation is performed to avoid bottlenecks in the compute work. Additionally, all the near real-time data being generated in the field is collected and available for inspection via the cloud Database Management System (DBMS) systems. The HPC data center at the ESIF is connected via VPN endpoints to the compute and data systems in Stratus, so the ability to inspect and use this data without transferring it all to the data center is supported. In other cases, the data systems may be available via AWS Application Programming Interface (API)s, which allows for direction interaction with data queued in a highly available, scalable, and low-

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

cost AWS Kinesis stream [10], or placed in AWS data services such as DynamoDB or S3. While the rate of processing in the HPC data center can lead to contention in the cloud because of the high throughput rates, the auto-scaling capabilities of the cloud will often be able to correct these issues during HPC job executions. AWS also sets soft limits intended to protect customers from steep rises in costs; these limits also allow AWS to ensure they have provisioned enough resources to meet the limits set by its customers, so a review of these limits is required.

Machine Learning Training. Machine learning models often require retraining and tuning as the environment changes to address model drift. Typically, this is supported by data pipelines, which use a portion of the historical operational data as training and testing data sets. Again, the environment chosen for training will depend on the purpose or scope of the machine learning.

In the case of temperature sensors, the metered power usage will likely vary significantly with the seasons, so as the seasons progress, the historical data can be used to update the models to recognize newer scenarios.

Similarly, with the COVID-19 pandemic dramatically altering how and where energy is used, more recent data will provide the patterns needed to correct the drift in the model.

It is common for regular batch computing jobs to perform training on collected datasets at regular intervals. The jobs themselves are compute intensive and can require access to larger datasets, so this type of work is not performed at the edge, but instead in the Stratus or HPC environment. In the case of the temperature sensors, it is likely that a regular cadence would be beneficial. The cloud computing tools in Stratus make regular scheduling of training jobs simple. In contrast, scheduling repeating work in the HPC data center is less reliable due to maintenance periods and contention with other work being performed. As such, training jobs for these data sets or any other recurring training based on schedule or events will be part of a data pipeline run in the cloud. The pipeline may run the training jobs, publish the updated models, and then deploy those models to inference endpoints in the cloud or at the edge. For large or particularly intensive training jobs that are not reliant on repetitive scheduling, HPC jobs on Eagle would be an appropriate forum. The models could then be pushed to the cloud for further distribution if they are needed in the real-time data systems.

Edge Device Management. Devices at the edge will typically be provided by multiple manufacturers, and many will have proprietary interfaces that do not use industry standard protocols. Some of these devices may publish secure APIs for local access, or in other cases, the manufacturer may collect data that can be retrieved from a centrally managed API available on the internet. That said, others may use Modbus, OPC Unified Architecture (OPC-UA), or other standard interfaces and protocols to interact with the device. In most of these cases, a gateway device should be capable of interfacing with multiple devices installed at the edge. AWS provides GreenGrass with native Modbus and OPC-UA support as a means of interfacing with a large array of IoT devices. In the Stratus environment, AWS IoT supports the creation of shadow devices, which digitally represent the known and desired state of devices in the field [4]. With the understanding that remote connectivity is often unreliable, when controls or configuration updates are issued to field devices, the shadow's state is updated. Devices that are offline during the control or configuration events will receive the event when they come back online, and it is recognized that the desired state represented in the shadow device diverges from the running state of the device. The key components in this scenario are shown in Fig. 3.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

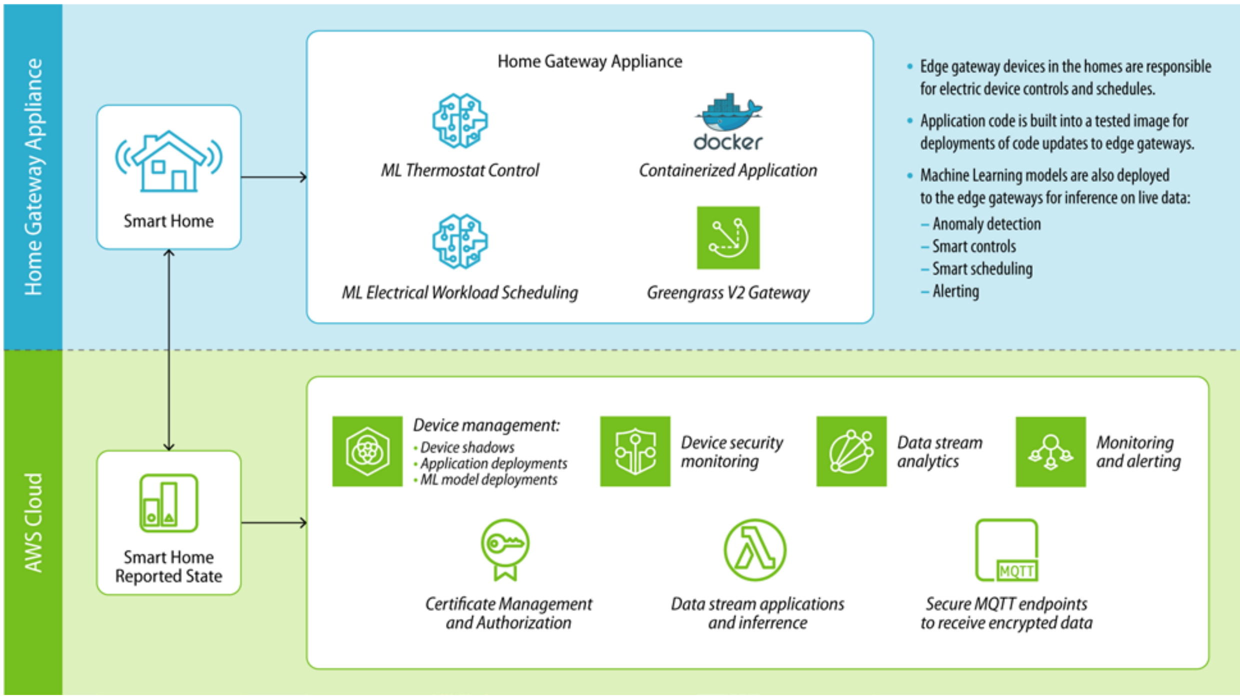

Fig. 3. Details of edge computing and machine learning integration based on smart home data and integration

As noted before, the IoT services in Stratus hosted by AWS also support code and machine learning model deployments to the edge. These deployments can be released to fleets of devices, thus enabling the management of a large footprint from a centralized service. The state of each deployment is available for review.

Another service available in the Stratus environment is provided by SiteWise, an AWS service that provides methods to monitor streams and provide live dashboards of activity. The streams from SiteWise become the source for data streams into other areas of AWS, supporting machine learning and analysis in the cloud.

Finally, cloud providers such as AWS are experts in security automation, and understand the controls needed to certify their service with different security schemes. At the national laboratories, sensitive data (in this case, data that contains details about homes in smart neighborhoods), is often classified as moderate level data, and ideally would require services to be certified under the Federal Risk and Authorization Management Program (FedRAMP) process. In the Stratus environment, the GreenGrass software, IoT Core endpoints and rules engine [11], and IoT Device Management services are FedRAMP moderate certified.[1]

Feedback and Tuning Mechanisms. The Stratus environment provides a very reliable and robust method of logging and auditing systems, both in the cloud and externally. Nearly every activity in the cloud is based on an API call, and AWS provides a robust events system based on those API calls, which can also be augmented with custom events. For instance, a scheduled event could trigger a data pipeline supporting machine learning. Alternatively, that same pipeline could be triggered by an event signaling the arrival of a file in S3, a threshold being hit in a data stream, an alert being received from an IoT gateway, or an event triggered by inference on a related dataset. This flexibility is also built on highly available systems that span data centers supporting highly reliable processes.

Summary. To summarize, the architecture supports the following services at the specified locations:

At the edge, via GreenGrass software, we will have these capabilities:

- Low latency data collection from IoT devices
- Low latency controls of IoT devices, such as adjusting thermostats or setting wash schedules

[1] GreenGrass, IoT Core and IoT Device Management are also certified under the Payment Card Industry (PCI), System and Organization Controls (SOC 2) and Department of Defense Cloud Computing Security Requirements Guide (CC SRG) [1]

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

- Low latency algorithmic or machine learning inference informed decision support, supporting decisions such as when to schedule certain events or when to alert on anomalies
- Low latency notification of safety events to residents or operators of the smart neighborhood
- Integration with AWS IoT to receive deployments of code or machine learning models from the cloud as a means to improve operations of the smart neighborhood
- FedRAMP moderate certified secure software and communications with AWS IoT endpoints for delivery of data [1]
- Data aggregation, validation, and cleaning, based on algorithmic or machine learning informed models, to help ensure downstream systems are operating efficiently
- Delivery of compressed, raw data for archival purposes.

Near the edge, when more compute power is needed, larger and more disparate holistic views of the data processing environment are needed, and latency requirements can be relaxed, AWS WaveLength or Local Zones provide the following capabilities:

- Processing of large data streams from more distributed sources, such as aggregate power consumption profiles across multiple neighborhoods or power distribution events
- Integration of more compute intensive machine learning inference engines that run models developed by HPC systems for larger scale coordination of controls affecting energy consumption in a region
- More intensive data cleaning routines, taking into account variables from surrounding areas, such as temperature readings from sensors in close proximity
- Large storage capacity for in-transit data streams or for archival data preparation.

In the cloud, we have the following services and capabilities:

- Scalable data collection via IoT endpoints or transfers to S3 to support real-time data streams and receipt of archival data
- Rules engines to send different datapoints to different services for processing, such as sending anomaly events directly to an operator via AWS notification services, or triggering a machine learning training when a particular dataset arrives
- Multiple DBMS systems available for analysis, such as timeseries databases for the incoming IoT timeseries data, columnar storage data warehousing for analysis of large datasets generated by HPC, and graph databases for digital things representations and APIs
- Services to manage remote IoT devices and gateways, including use of shadow devices, monitoring for tampering of devices, and deployment mechanisms for machine learning models and custom code
- A robust, highly available events system to reliably trigger services or processes, such as notifying operators when a remote device goes offline
- Secure communication, authorization, and authentication with the HPC data center and data center processes for providing or receiving data.
- Integration of AWS backed capabilities in HPC jobs, such as queue management, inference support, or workflow support.

In the HPC data center, the following capabilities are supported:

- Large processing jobs to simulate modeled energy usage and production data based on scenarios provided by researchers and guided by actual behaviors measured in the field
- Large machine learning jobs leveraging real or simulated data to train models to be deployed to the cloud or edge, to help improve energy efficiency and user satisfaction.

5 Challenges to Supporting the Vision

As with any larger scale ecosystem, the complexity and operational requirements of the deployed infrastructure and code create a significant challenge on many levels. The management of the Stratus environment has been incrementally built upon highly available and scalable services, such that simple services created by NREL or provided by AWS serve single purposes in a highly available and scalable manner. This allows for the development of patterns and practices, largely captured in code, that initially focus on smaller, focused deployments but that can be repurposed for scenarios of much higher scale or complexity. Developers need to be mindful of the potential costs and the AWS limits that will play a part in a large-scale deployment. The ability to fully define the environment their applications run in, as well as the applications themselves, should improve service levels overall, as the infrastructure can be tuned and tested per the software's requirements.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Data flow and data management are also a primary concern for systems such as this, where edge data and HPC modsim data are feeding into processes in the cloud used to augment the system or provide analysis, and edge data is being pulled into HPC jobs to augment the modeling being done. To fully support the environment, Stratus can publish data streams that can be consumed by any consumer connected to the internet with proper permissions, assigned using AWS APIs. It can also make available compressed files to the same consumers, query access to multiple data stores via the same API method or provide VPN access to DBMS systems hosted in the virtual private cloud in AWS.

6 Building Toward the Vision at NREL

The HPC data center at NREL has long been the computing engine driving innovation and research at the laboratory. In the decade since the introduction of cloud computing at NREL, the services provided to NREL by the CSC Cloud Team continue to expand as more specialized services are launched by the team, AWS, and other cloud providers. The expansion of computing at the edge provides additional opportunities and challenges as NREL begins to support power generation experiments, field studies, mobile fleet management, and other IoT supported work, and looks to further enhance and optimize power systems on both the production and consumption sides. As new research opportunities arise that leverage the broad range of capabilities now available to the laboratory, NREL continues to strategically invest in its advanced computing capabilities.

One of the recent steps taken was merging the cloud computing team, who manage the Stratus environment, with the HPC operations team, who manage Eagle and the ESIF HPC data center, under the banner of Advanced Computing Operations. Strategic initiatives have allowed the laboratory to fully leverage the hybrid environments unique to integrated HPC work. The HPC team continues to grow, improve, and expand HPC via dedicated efforts to launch private cloud tools on OpenStack and define and install the next generation of powerful and efficient supercomputers. The Stratus team is expanding support for specialized, highly available and flexible tools for deploying applications, machine learning models, data collection, and data analysis. Projects are underway to investigate the use of cloud computing to improve the efficiency of Eagle by moving inefficient work to the cloud, where it can be run at a lower cost, as well as by developing methods to simplify data transfer and data management between the edge, cloud, and HPC, and by developing methods to make machine learning inference available from multiple environments via centralized deployment mechanisms in the cloud.

The Advanced Computing Operations team continues to support education at the lab on using the various computing systems in support of our mission [9]. This now includes introducing concepts around high availability and scalability best practices to properly leverage cloud computing application support. The Advanced Computing Operations team has decades of cloud computing experience used to help augment the already successful processes in the ESIF data center supporting computationally intensive work. By merging the new cloud computing capabilities with the established, growing HPC capabilities, NREL is suddenly able to successfully support a much wider range of ambitious initiatives that could not be addressed by either computing capability alone.

As these efforts progress, the architectural vision supporting the scenario described in this paper moves much closer to reality. Currently, this vision depends on workloads that NREL is ready to implement; the necessary compute and infrastructure capabilities are currently available. The challenges are now primarily composed of education, process definition, governance and implementation of security controls, and work to support integrations with proprietary interfaces or protocols. While these challenges can still be substantial, they are being addressed with typical risk management strategies and process improvements, and as NREL continues down this path, the risk in those areas will become more manageable and better understood. In essence, the sky, full of clouds, is the limit.

References

1. AWS Services In Scope, https://aws.amazon.com/compliance/services-in-scope/, last accessed 2021/05/26.
2. Bennett, K., Robertson, J. "Remote sensing: leveraging cloud IoT and AI/ML services," Proc. SPIE 11746, Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications III, 117462L (12 April 2021); https://doi.org/10.1117/12.2587754.
3. Borge, S., Poonia, N. Review on amazon web services, google cloud provider and Microsoft windows azure. Advance and Innovative Research, page 53, 2020.
4. Botez, R., Strautiu, V., Ivanciu, I., Dobrota, V. "Containerized Application for IoT Devices: Comparison between balenaCloud and Amazon Web Services Approaches," *2020 International Symposium on Electronics and Telecommunications (ISETC)*, 2020, pp. 1–4, doi: 10.1109/ISETC50328.2020.9301070.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

5. Javed, A., Malhi , A., Främling, K. "Edge Computing-based Fault-Tolerant Framework: A Case Study on Vehicular Networks," *2020 International Wireless Communications and Mobile Computing (IWCMC)*, 2020, pp. 1541–1548, doi: 10.1109/IWCMC48107.2020.9148269.
6. Liu, Y., Yang, C., Jiang, L., Xie, S., Zhang, Y. "Intelligent Edge Computing for IoT-Based Energy Management in Smart Cities," in *IEEE Network*, vol. 33, no. 2, pp. 111–117, March/April 2019, doi: 10.1109/MNET.2019.1800254.
7. Mesbahi, M., Rahmani, A., Hosseinzadeh, M.: "Reliability and high availability in cloud computing environments: a reference roadmap". *Hum. Cent. Comput. Inf. Sci.* **8,** 20 (2018). https://doi.org/10.1186/s13673-018-0143-8.
8. Muhammed, A., Ucuz, D. "Comparison of the IoT Platform Vendors, Microsoft Azure, Amazon Web Services, and Google Cloud, from Users' Perspectives," *2020 8th International Symposium on Digital Forensics and Security (ISDFS)*, 2020, pp. 1–4, doi: 10.1109/ISDFS49300.2020.9116254.
9. National Renewable Energy Laboratory's Mission, https://www.nrel.gov/about/mission-programs.html#:~:text=NREL%20advances%20the%20science%20and,integrate%20and%20optimize%20energy%20systems, last accessed 2021/05/24.
10. Nguyen, D., Luckow, A., Duffy, E., Kennedy, K., Apon, A. "Evaluation of Highly Available Cloud Streaming Systems for Performance and Price," *2018 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*, 2018, pp. 360–363, doi: 10.1109/CCGRID.2018.00056.
11. Pflanzner, T., Kertesz, A. "A survey of IoT cloud providers," *2016 39th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, 2016, pp. 730–735, doi: 10.1109/MIPRO.2016.7522237.
12. Pham, T., Ristov, S., Fahringer, T. "Performance and Behavior Characterization of Amazon EC2 Spot Instances," *2018 IEEE 11th International Conference on Cloud Computing (CLOUD)*, 2018, pp. 73–81, doi: 10.1109/CLOUD.2018.00017.
13. Mohammed Sadeeq, M., Abdulkareem, M., Zeebaree, N., S. R. M., Mikaeel Ahmed, D., Saifullah Sami, A., & Zebari, R. R. (2021). IoT and Cloud Computing Issues, Challenges and Opportunities: A Review. *Qubahan Academic Journal*, *1*(2), 1–7. https://doi.org/10.48161/qaj.v1n2a36.
14. Singh, V., Dutta, K. "Dynamic Price Prediction for Amazon Spot Instances," *2015 48th Hawaii International Conference on System Sciences*, 2015, pp. 1513-1520, doi: 10.1109/HICSS.2015.184.
15. White House FACT SHEET: President Biden Sets 2030 Greenhouse Gas Pollution Reduction Target Aimed at Creating Good-Paying Union Jobs and Securing U.S. Leadership on Clean Energy Technologies, https://www.whitehouse.gov/briefing-room/statements-releases/2021/04/22/fact-sheet-president-biden-sets-2030-greenhouse-gas-pollution-reduction-target-aimed-at-creating-good-paying-union-jobs-and-securing-u-s-leadership-on-clean-energy-technologies/, last accessed 2021/08/31.
16. Zhang, J. et al. "Energy-Latency Tradeoff for Energy-Aware Offloading in Mobile Edge Computing Networks," in *IEEE Internet of Things Journal*, vol. 5, no. 4, pp. 2633-2645, Aug. 2018, doi: 10.1109/JIOT.2017.2786343.