Outcomes, Insights and Best Practices from IIC Testbeds: Deep Learning Facility Testbed

Interviewees:

Brad Klenz
IoT Analytics Architect
SAS
Brad.Klenz@sas.com

Ken Hatano
Chief Specialist
Toshiba
ken.hatano@toshiba.co.jp

Said Tabet
Chief Architect, Emerging Technologies & Ecosystems, CTO Office
Dell Technologies
said.tabet@dell.com

Interviewer:
Joseph Fontaine
Former VP Testbed Programs
Industrial Internet Consortium

To extend the usefulness of the published testbeds in the Testbed Program of the Industrial Internet Consortium (IIC), the Testbed Working Group has developed an initiative to interview the contributors of selected testbeds to showcase more insights about the testbed, including the lessons learned through the testbed development process. This initiative enables the IIC to share more insights and inspire more members to engage in the Testbed Program.

This article highlights the Deep Learning Facility Testbed. The information and insights described in the following article were captured through an interview conducted by Mr. Joseph Fontaine, Former Vice President of Testbed Programs at IIC, with Brad Klenz, IoT Analytics Architect at SAS; Ken Hatano, Chief Specialist at Toshiba; and Said Tabet, Chief Architect, Emerging Technologies & Ecosystems, CTO Office at Dell Technologies. Brad, Ken and Said are active members in the IIC where they serve as co-leads of the Deep Learning Facility Testbed and are key contributors to the Testbed Working Group.

Deep Learning Facility Testbed – From Concept to Reality

When the Deep Learning Facility Testbed began in late 2016, no other IIC testbed dealt explicitly with artificial intelligence (AI) or deep learning. Dell Technologies (EMC at the time) and Toshiba had set a goal to launch a deep-learning-based, technology-based testbed. Fortunately, Toshiba owned and operated the Toshiba Smart Community Center, a state-of-the-art building established in 2013 in Kawasaki, Japan that has more than 30,000 data points and sensors. Toshiba built this brand-new building for the purpose of experimenting, and many Toshiba employees transferred from their Hamamatsu-cho headquarters to Kawasaki. One of the objectives was to optimize maintenance. There was also a need to increase energy efficiency through the adjustment in all power-consuming services in buildings while improving the visiting customer experience and employee-resident comfort. At that time, Toshiba also had a department that focused on deep-learning-specific technology. Based on the resources and technologies from Toshiba and Dell and the perceived needs in the industry, the IIC Deep Learning Facility Testbed was founded to improve building operational efficiency and occupant satisfaction using AI deep learning techniques, and to share the new-found knowledge within the IIC.

In early 2019, SAS joined the testbed in accordance with one of their main corporate objectives: environmental sustainability. SAS has experience with sustainability as a focus area, so they came into the project with prior knowledge on energy efficiency. When starting down the path to make a building more sustainable, many physical infrastructure changes, such as lighting choices, can improve energy efficiency. Progressing down that path leads to more reliance on technology and analytics to improve energy conservation continually. There is also immediate payback associated with more sustainable options, such as money saved on electricity and energy bills.

In addition to a focus on energy efficiency, SAS aligns with Toshiba's mission to focus on equipment, predicting equipment failures and finding equipment malfunctions—a difficult mission analogous to finding a needle in a haystack. IIC's Deep Learning Facility Testbed provides the avenue for these companies to fulfill this vision.

There are areas of experimentation in the Deep Learning Facility Testbed. Modern buildings have many IoT sensor systems for various systems within the building. Sensors monitor large-asset maintenance, granular energy efficiency and occupant usage. AI and Deep Learning can use all of the data from these individual systems to achieve higher-order objectives such as energy sustainability and occupant satisfaction.

The Deep Learning Facility Testbed's Phase 1 target at the Toshiba Smart Community Center is sensor-anomaly detection. The building facility management team is concerned with the shortage of people to monitor, control and manage all the building assets, but AI can support these tasks automatically. Deep-learning technology assists the management team in detecting anomalies in sensors and assets and analyzing data from 30,000 data points.

Another area of experimentation is trying to determine what is computationally feasible and how to apply various high-end technologies to maximize those computation capabilities. For example, autoencoders have been used for anomaly detection. Deep learning needs a substantial amount of time for computing resources, and Dell's rapid storage system helps establish a distributed deep-learning system platform.

In addition to being deployed at the Toshiba Smart Community Center in Kawasaki, the Deep Learning Facility Testbed is also implemented at the SAS Smart Campus in Cary, North Carolina, USA. This corporate headquarters has 24 buildings, ranging from two months to 20 years old. Though primarily an office campus, there are also daycare centers, fitness facilities and restaurants. While the Toshiba Kawasaki building employs greenfield instrumentation, the SAS site also has brownfield. The testbed mostly covers buildings created within the last six years, but there are also projects underway to retrofit some of its older buildings.

SAS uses neural networks on all of the building data. The number of different measurements from the SAS buildings were similar to that of the Toshiba Kawasaki building. Neural networks help create a model that understands building energy usage for monitoring the building to determine whether it is performing as expected in accordance with its design and understanding where and why deviations from those expectations may occur.

The testbed plans to publish implementation lessons learned, best practices and use case results. The algorithmic knowledge and innovations discovered through the implementations of this testbed are also intended to be applied commercially by the participating companies.

Testbed Planning

Prior to the testbed, Dell and Toshiba had formed the basis of a relationship. The IIC helped to formalize this relationship into a partnership as it has done with other testbed teams. SAS was introduced to the team through the IIC ecosystem. SAS had prior knowledge in deep learning and AI analytics, as well as a campus with multiple facilities. Toshiba and Dell needed access to a singular campus with multiple buildings for Phase 2 of the testbed, and SAS was the ideal fit.

From SAS's perspective, the Deep Learning Facility Testbed represented an interesting project that dovetailed with some of SAS's ongoing activity. There was an opportunity to build upon existing commonalities, analyze differences between facilities and expand knowledge into areas to which each company may have had no previous direct connection. No single company has the hardware, software and domain expertise needed to implement the complete project. The IIC ecosystem helps provide the needed capabilities to achieve the testbed goals.

IIC Interactions

When the Testbed was first proposed, the Industrial Internet Reference Architecture (IIRA) and Industrial Internet Security Framework (IISF) played roles in formalizing how the system worked. As a standards-based architectural template and methodology, the IIRA enabled the testbed team to design their system based upon common frameworks and concepts in AI and deep learning technology. Similarly, the security design perspective of the testbed is attributed to knowledge gained from the IISF.

The Industrial Internet of Things Analytics Framework (Industrial IoT Analytics Framework or IIAF) also plays a role in the testbed. The IIAF spells out how companies who want to implement AI for IoT can search through different design patterns and choose those that match their respective cases. In this way, the IIAF has helped guide the design of the testbed, and the testbed provides validation of the framework by showing what works in the field.

Testbed Results

Three primary phases were laid out for the Deep Learning Facility Testbed:

- Phase 1 is centered around the training and analysis of AI for certain usage scenarios in the Toshiba Smart Community Center—specifically anomaly detection;
- Phase 2 will implement the new technology into various different facilities, particularly the SAS Smart Campus; and
- Phase 3 involves moving into a public facility.

The primary usage scenario that has been completed for Phase 1 is the anomaly detection of facility assets. Now the focus is shifting toward energy efficiency and the visitor-experience areas.

One of the major areas of experimentation in Phase 1 is how to incorporate the computational resources for deep-learning activity in an intensive AI environment comprising many data streams and sensors.

A distributed learning system was used to create parallel learning in the multiple servers, enabling shared results of the computations and learnings. Dell storage has a very high network speed for exchanging data, helping the many servers share their data. This was done using computational algorithms to scale the data and prevent saturation that would otherwise be caused by the number of servers.

An autoencoder was used on the Toshiba Kawasaki building for anomaly detection. Because the building is new, there is little defect data to be used to create a model for anomaly detection. The autoencoder therefore only uses data from normal operating conditions, and if unusual data is observed, it can be detected. In one case of an anomaly detection, the autoencoder picked up a dataset that had unusual data. A data scientist was able to analyze the data without any domain knowledge to find which input affects the output of the anomaly, and the building facility management team confirmed which sensor had the abnormal situation or defect.

Phase 2 is underway with the introduction of the SAS Smart Campus into the project. This focuses on two of the newer buildings at the SAS facility due to the large amount of useful data they produce. Neural network models have been running on those buildings for over a year, and they were trained using historical data that dated back about another year. The anomaly detection technology finds outliers, giving facility staff an opportunity to address them. While the number of anomalies cannot be predicted, a target can be set for reducing energy intensity in certain areas. Going forward, sensors have been installed on a couple of older buildings on the SAS campus that employ brownfield instrumentation. In some ways, the older sensor technology is better for targeting specific use cases, but it does not match the breadth that accompanies the newer buildings that have sensors on energy circuits, water systems and some indoor environment systems. Next steps may involve retrofitting the buildings.

At the SAS campus, one industry standard used for measuring energy efficiency is energy intensity, which measures the energy usage per square foot per year within a facility. The SAS facilities are currently running at 15 kWh per square foot per year—a 3.3% improvement over the year before.

Another area of technological achievement is AI-at-scale. It became apparent that employing one large neural network would not give desired results as it would not specify where to take action, nor indicate what action was needed. Therefore, the Deep Learning Facility Testbed employs a large number of smaller neural network models that can give insight into the individual systems and designate on which ones to act.

SAS has some customer engagements underway, one of which has been solidified as a reference.

Key Challenges

From Toshiba's perspective, one challenge was attributed to the domain knowledge of the building facility management team. In typical deep-learning technology, supervised learning is required, and so data must be labeled. Incorporating the autoencoder eliminated the need for this supervision. Another challenge is what the autoencoder cannot do — root-cause which device's behavior is unusual. Toshiba's data scientists did so with their expertise.

Dell perceived several challenges in applying deep learning to very large-scale environments in real buildings, rather than a simulation field in a lab. The real world is a different environment from a lab; the neural networks that are generated must be as precise as possible, and re-training may be required — it is not enough to train an algorithm and deploy it. This requires access to the same data by all compute nodes simultaneously, a notable innovation in the industry.

One challenge SAS encountered was the access and integration of data between the hybrid of different building management technologies. Even within one facility on the campus, there are multiple vendors who install systems there. To get a decent model of the building, the data from those different systems needs to be brought together. Buildings built at different times further add to the complexity of accessing and integrating that data.

Another challenge comes into play when accessing the data from various systems. Often, building managers have a contractor responsible for a particular system. The contractor installs the equipment and collects the data, but they may be hesitant to give out this data if they do not understand why someone is requesting it. The legal agreement must specify that the contractor must share the detail-level data of the system with the data scientists.

Experience

The teams from Toshiba, Dell and SAS have experienced shared learnings and cross-validation in projects through the Deep Learning Facility Testbed. Throughout working on the testbed, it has been useful to reference the IIRA and other IIC resources to enhance learnings. The testbed team looks forward to sharing their experiences with the industry so that others may benefit from their findings and launch their own projects in AI and deep learning technology.

Toshiba, Dell and SAS have been sharing experiences, collaborating and driving learning in terms of what can contribute toward their individual products. The companies have also found that the visibility of their own AI activities has improved due to their involvement in the testbed.

In offering advice to other testbeds or companies embarking on leading-edge IoT implementations, the Deep Learning Facility Testbed team stresses that IoT solutions cannot be realized by just one company. Partner matching is a crucial aspect of reaching a solution, and taking advantage of the IIC ecosystem plays a significant role in setting the stage for a successful testbed.

Closing

While Phases 1 and 2 of the Deep Learning Facility Testbed are well underway, and Phase 3 — transferring the testbed’s technology into a public facility — is on the horizon, there have been a number of key learnings and insights since the testbed’s launch in 2016. The testbed team hopes to build from their learnings and create a more dynamic collaboration within the community to drive the testbed to the next level, leveraging the vast amount of work and innovation that has been accomplished so far.

➢ Return to IIC Journal of Innovation landing page for more articles and past editions