HOW WWT EXTENDS THE AWS EXPERIENCE ON PREMISES WITH AWS OUTPOSTS

According to the 2020 IDG Cloud Computing Study, 92 percent of organizations have workloads running in the cloud while 8 percent still remain on premises. Challenges preventing these on-premises applications from moving to the cloud include regulatory compliance, the latency sensitivity of workloads, dependencies on legacy systems, and local data processing.

THE HYBRID CLOUD & AWS OUTPOSTS

Organizations typically adopt a hybrid cloud model, which allows them to operate both on premises and in AWS. However, most hybrid cloud architectures involve different infrastructure, services, APIs, and application development and deployment tools — both on premises and in the cloud.

Customers are looking for a consistent experience across their on-premises and cloud environments. Many customers also find that the complexity of procuring, deploying and managing on-premises infrastructure slows down both productivity and the pace of innovation.

AWS Outposts is a service that seamlessly extends AWS infrastructure, services, APIs and tools to customers' on-premises environments. AWS delivers, installs, manages and monitors the Outposts rack, relieving customers of the operational burden typically associated with on-premises deployments.

Over the years, World Wide Technology (WWT) customers have engaged our Cloud experts with many different hybrid cloud use cases and issues suitable for AWS Outposts. For example:

- Our **manufacturing customers** are migrating to AWS, but they still have plants in remote locations. These plants run application stacks that need low latency for processes and have resources running locally that cannot be migrated to the cloud. These manufacturers are looking at AWS Outposts to run latency-sensitive applications locally and seamlessly extend to other workloads running in the AWS Region.

- Our **automobile customers** are exploring edge solutions for local data processing for autonomous vehicle testing. They're exploring solutions that include AWS Outposts and AWS Snowball devices.

- Our **financial customers** want to use AWS Outposts as part of their fraud detection and analysis systems for low latency computations and regulatory compliance.

- Our **healthcare providers** are evaluating AWS Outposts as a means to speed on-premises patient analytics while storing long-term data in the cloud.

- Our **enterprise customers** are looking at AWS Outposts to accelerate application refactoring for workloads that depend on on-premises systems.

Across these use cases, WWT's Cloud experts are seeing customers who want to standardize tooling and automation across their on-premises and cloud deployments using AWS CloudFormation and Terraform.

These customers are hoping to integrate into continuous integration/continuous deployment (CI/CD) pipelines.

They are looking to run cloud services — such as database, artificial intelligence (AI) and machine learning (ML), and data analytics — locally on Outposts for data that can't be migrated to the cloud for regulatory reasons.

They are seeking to run container-based workloads using Amazon EKS and integrate with on-premises clusters.

WWT also has customers looking to incorporate Outposts into their edge compute architectures in conjunction with AWS Wavelength.

In other words, there are many situations where AWS Outposts might be right for your business.

TESTING AWS OUTPOSTS IN WWT'S ADVANCED TECHNOLOGY CENTER

WWT has a proven track record of developing and implementing solutions that integrate on-premises infrastructure, complex network environments and enterprise security. As an AWS Advanced Consulting Partner, WWT has a deep understanding of technology and services in the market. We are the rare partner who can assist at any stage of the cloud journey.

Combining expertise in both strategy and execution, WWT offers customers the best of both worlds when designing and deploying hybrid cloud solutions using AWS Outposts.

Since its inception in 2009, WWT's Advanced Technology Center (ATC) has served as a next-gen innovation lab housing cutting-edge infrastructure and software designed to deliver faster, better solutions to our customers. The ATC was the logical place for WWT to deploy an Outposts rack. It allows our architects, engineers and data scientists to seamlessly develop solutions for the hybrid cloud use cases most relevant to our customers.

One of the key benefits of the ATC is the ability for customers to perform like-for-like comparisons and test different architectures and configurations — all with zero impact to their existing environments. We facilitate these proofs-of-concept and vendor bake offs for customers on a daily basis in the ATC.

By working with a team of WWT experts to support the full range of hybrid cloud engagements, our customers can make data-driven decisions faster and more efficiently. We're ready to show you what Outposts can do.

USE CASE: BUILDING A HYBRID ARCHITECTURE FOR TRAINING ML MODELS

In partnership with AWS, WWT has leveraged the ATC to fully demonstrate the value of AWS Outposts for training AI/ML models. Many customers want to leverage their cloud-based AI/ML models to analyze on-premises data that otherwise cannot be moved to the cloud. Building a hybrid cloud architecture with AWS Outposts is the ideal solution for customers with this use case. To validate this premise, we developed a test strategy for customers to use in their R&D efforts, detailed below.

WWT used this framework in an R&D project focused on a Natural Language Processing (NLP) AI developed in AWS. The AI pipeline consisted of a web scraper tool that searched articles on WWT.com and applied an NLP methodology to interpret the articles. Training this type of AI model requires significant computational resources and is best performed on NVIDIA GPU-enabled instances. The data for the pipeline itself, though publicly available, is stored on premises for low latency access.

WWT used an approach known as Machine Learning Operations (MLOps) to address the NLP pipeline. With MLOps, CI/CD style principals are applied to an ML model so that the application, data and modeling components are promoted from development to test to finally production in an agile manner. This MLOps methodology allows the ML pipeline to be continuously evaluated and updated post publication.
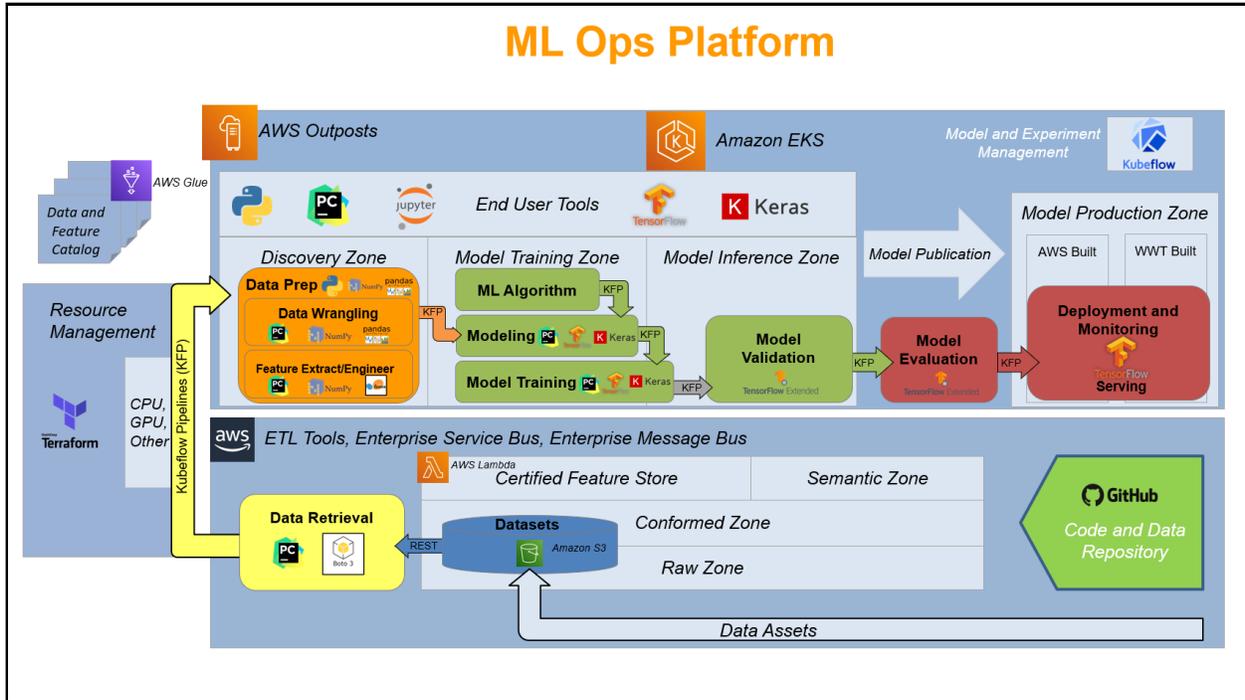
WWT deployed an AWS Outposts rack provisioned with G4 graphics-accelerated Amazon EC2 instances. We used a number of different instance configurations to optimize compute resources. The final configuration featured 4x g4dn.12xlarge instances.

Next, we adapted the MLOps pipeline architecture to work in a hybrid cloud model. The initial model was developed in the AWS Region, then seamlessly deployed on Outposts for training and inference. The model ran on an Amazon EKS cluster deployed locally on Outposts. The model accessed locally sourced data for training and inference with low latency. The model has been consistently evaluated over time and adjusted through the MLOps pipeline.

WWT has successfully demoed the AI with MLOps running on Outposts to several customers across different industries. For example, a large healthcare IDN needed to keep data on premises for regulatory reasons. Using our AWS Outposts architecture, the customer can now offer a consistent AWS experience to their developers and data scientists as they use the same AWS services, APIs, console and tools as in the Region. They're also able to avail consistent and dedicated compute resources, paid upfront with no separate hourly costs.

LEARN MORE ABOUT AWS OUTPOSTS

For more about AWS Outposts, visit AWS today and explore the available AWS Outposts Documentation.

CONCLUSION

As demonstrated above, AWS Outposts can meet the growing demand for hybrid cloud use cases in both public and commercial sectors.

WWT would love to hear about your particular use cases and discuss the many new ones we're busy testing in the ATC right now — they just might be relevant to your organization!

Reach out to a WWT Cloud expert today to discuss how we can incorporate AWS Outposts into your next hybrid cloud solution.